July 2018

Keppel Capital watch

Data Centre

Data Taking Off into the Cloud

Contents

Executive Summary	1
Introduction	2
State of the cloud – all aboard!	2
How the cloud has shaped data centres	4
Runway remains long for data centres	8
Conclusion	10

Executive Summary

We are in the midst of a global data explosion, so much so that projections for future data growth kept being revised up. The biggest driver for data centre demand is the growth of cloud services, which underpins many important technological processes going forward.

Virtualisation has been a boon, instead of being a threat, as it increases cloud adoption by making computing more cost-efficient. The accelerated use of data analytics and machine learning, the mushrooming of smart devices like wearables and smart home assistants, and the advent of 5G mean more demand for data centres.

Hyperscale Service Providers (HSPs) in their endeavours to improve connectivity and service reliability are building more hyperscale data centres, primarily in gateway cities where the bulk of global data centre compute resources, storage capacity and data traffic reside. To increase redundancies, HSPs have progressively grouped their own cloud infrastructures into Regions and Availability Zones (AZ) where the data centres are physically independent, yet interconnected by fibre links within each Region and AZ. This usually translates to more data centres.

To achieve cheaper and lower latency connectivity, HSPs have also ramped up subsea cable investments. The ability to land cables directly inside a data centre as opposed to a traditional cable landing station on the beach will be a service differentiator for cloud service providers and data centre operators.

More smaller-sized data centres are being built in secondary markets as demand increases for computing to be done at the edge or near-edge in closer proximity to sources of data generation or end-users. This is essential for low-latency processes such as autonomous driving. The demand for edge computing will be accelerated by the commercial rollout of 5G mobile technology expected sometime between 2019 and 2020.

As HSPs strengthen their availabilities and resilience both in gateway cities and emerging markets, hyperscale and colocation data centre demand can only grow. Nonetheless, operators must avoid speculative overbuilding, adopt systems to improve energy efficiencies and safeguard against cyber-attacks and technological failures. Investors and operators must also keep abreast of architectural developments and be mindful that growth in key cities may eventually peak. Until then, there is still a long runway for growth.

Authors

Chua Chor Hoon
Director, Research

Wilson Liew
Vice President, Research

Introduction

Data creation growth has accelerated

In our earlier paper on "Data Centres: Still a Long Runway for Growth" in May 2016, we highlighted that the global data explosion was still in the early phases of exponential growth with further leeway to accelerate. Today, the runway for data growth remains lengthy. In fact, the faster-than-expected growth of data has led industry forecasts for 2020 being revised upwards by another c.20% between 2012 and 2017 (see Figure 1).

Figure 1: Global data growth outperforms previous industry forecasts by 20% on average

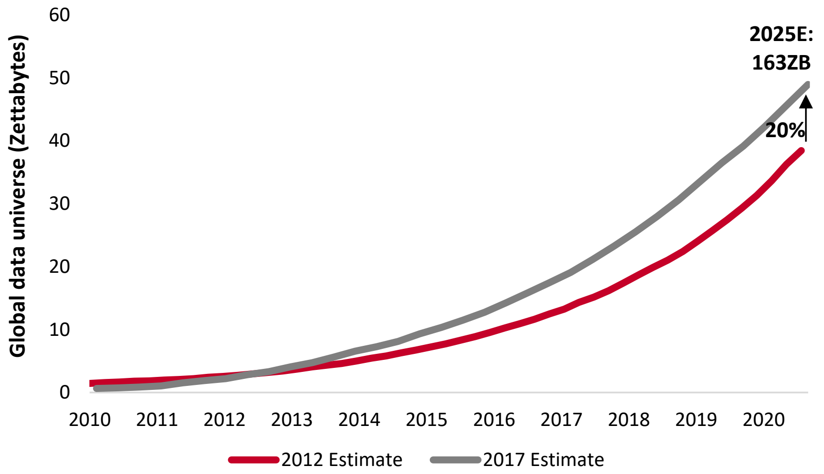

Source: IDC, KepCap Research

State of the cloud – all aboard!

Cloud adoption continues to accelerate, driven by the better reception of cloud-based services by both consumers and enterprises.

<u>1) Enterprises</u>

Enterprises accounted for the bulk (i.e. 76%) of global data centre workloads in 2016. Cisco expects the enterprise segment to grow at 19% CAGR (2016-2021), driven by big-data analytics and the increasing migration of workloads to the cloud.

Virtualisation – net positive for data centre demand. Contrary to earlier concerns that virtualisation may reduce the amount of server space required in data centres, it has instead been a boon for cloud adoption. Cloud service providers create economies of scale through virtualisation of their data centre servers by aggregating customer traffic demand and passing on the cost savings to their clients. The ability for enterprises to rapidly scale their usage in response to business requirements in a cost-effective manner is a key factor behind the migration of work processes to the cloud, thereby driving the demand for data centres.

The Fourth Industrial Revolution. The industrial world is on the cusp of the next phase of industrial revolution. In essence, Industry 4.0 integrates advanced production and operation techniques with state-of-the-art digital processes interconnected by the internet and powering technologies like AI-powered robotics and data analytics. Industrial Internet of Things (IIoT) will also power real-time machine-to-machine and machine-to-human interaction, such as Virtual Reality (VR) or Augmented Reality (AR). With the vast amounts of raw data generated, the cloud will be an integral part of the IIoT backbone in supporting processes such as correlation analyses and machine learning.

Big-data Analytics. Cloud computing has led to the increased adoption of big-data analytics by enterprises. Enterprises such as online retailers can tap on the wealth of information about their customers through their increased digital footprints, and derive insights that help shape business strategies going forward. Enterprises also deploy big-data analytics to help with predictive maintenance. For example, each Boeing 787 Dreamliner generates 0.5 TB of relevant operational data per hour of flight which gets transmitted to the cloud for storage and analysis.

Everything-as-a-Service (XaaS). The world is shifting away from a product-focused economy into a service-based one. For example, Adobe transitioned its Creative Suite product away from a perpetual licencing model to a subscription-based cloud version (Creative Cloud). This shift is not only generating more cloud traffic, it is also creating new monetisation opportunities which were otherwise unfeasible. Examples include Uber and Lyft, which utilise cloud computing to operate transportation-as-a-service applications. XaaS will see increased use cases as more adopters recognise the cost-savings from the "pay-as-you-go" model, therefore driving more cloud-based demand for data centres accordingly.

2) Consumers

Consumers represented 24% of the global data centre workload in 2016, and Cisco expects this to grow at 27% CAGR (2016-2021), driven by social networking and video streaming (see Figure 2).

Figure 2: Social networking and IoT are top drivers of cloud traffic growth for consumers and enterprises respectively

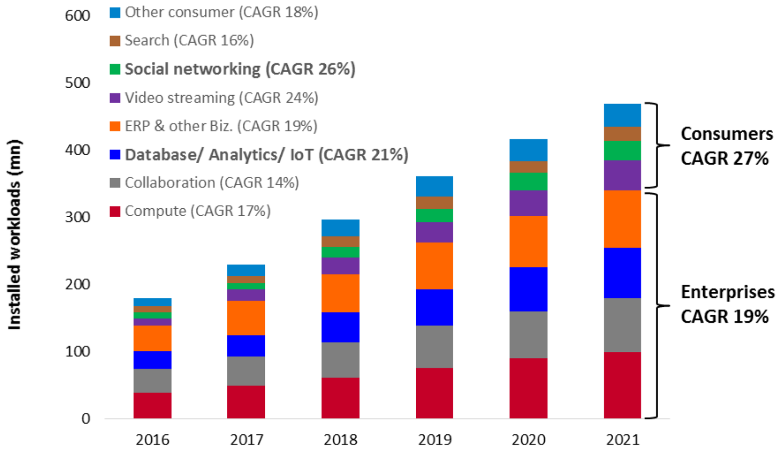

Source: Cisco Global Cloud Index, 2016-2021, KepCap Research

Higher internet media usage. Social media usage has broadened in scope to incorporate video-sharing, while incumbents build out live-streaming capabilities. Consumers are also gravitating towards over-the-top (OTT) subscriptions like Netflix. According to Cisco's Global Cloud Index (GCI), video streaming and social networking are the two main contributors of consumer cloud traffic, and will cumulatively represent over half of all consumer workloads by 2021 as access to the Internet and mobile devices improve globally.

Increased number of smart devices. Consumer IoT has led to an increase in smart devices in various forms and factors, ranging from wearables (e.g. FitBit) to smart home appliances and home assistants (e.g. Amazon Alexa and Google Home). Gartner expects the number of IoT devices to more than double in the next three years to 20 billion units, underpinned by consumer devices growing at 35% CAGR. This greatly increases the frequency of interaction between users and the Internet, leading to a greater digital footprint and enabling cloud-based applications like big-data analytics and machine learning.

How the cloud has shaped data centres

Cloud service providers in hyperdrive. Cloud adoption has been a key driver of demand for data centre space and is projected to keep growing. Consequently, HSPs are enjoying robust revenue growth, with which they are reinvesting by ramping up IT infrastructure capex to expand existing data centres and build out new ones, which may be self-owned but more likely to be leased via colocation.

2017 was an exceptional year as total hyperscale capex hit US$74 billion but is likely to be surpassed in 2018. Total spending for 1Q2018 had already reached US$27 billion[1].

[1] Synergy Research

The ramp-up in HSP capex is not merely to keep pace with strong customer demand. HSPs are also competing against one other to offer the best service levels with regards to uptime reliability, availability, compute power, storage, latency, connectivity and security, which will have positive impact on data centre demand. Not only will HSPs be a key source of data centre demand in the years to come, their actions will also shape where and in what form that demand will be.

Boom in subsea cable investments. 99% of internet traffic today is piped through subsea fibre optic cables which often traverse oceans. Being the internet backbone, existing subsea cables are predominantly owned by consortia made up of telecommunications companies (telcos). As cloud adoption grows rapidly, the need for more transmission capacity becomes more apparent. To gain a competitive edge over one another, HSPs have become major investors in new subsea cable projects. Technological advancements in optical transmission now allow cables to land directly inside a data centre as opposed to a traditional cable landing station on the beach, resulting in cheaper, lower latency connectivity to the cloud. This becomes a service differentiator both for cloud service providers and data centre operators.

Building for resilience through Regions and Availability Zones (AZs). As more mission-critical workloads and data storage take place in the cloud, HSPs continue to improve the resiliency of their networks by increasing redundancies, which usually translates to more data centres. This is evidenced by how the major HSPs have progressively grouped their own cloud infrastructures into Regions and Availability Zones (AZ) where the data centres are physically independent yet connected by redundant network links to improve redundancy (see Box 1).

Box 1: Introduction to cloud Regions and Availability Zones
Pioneered by AWS, a Region is a geographical location comprising at least two but usually more AZs comprising physical data centres on the ground. Regions are physically isolated from and independent of one another in terms of location, power, water supply, etc. The segregation among Regions is critical for workloads with compliance and data sovereignty requirements. Each AZ in turn is a group of at least one data centre having independent power sources, networking and cooling resources. Each data centre belongs exclusively to one AZ and within the AZ, data centres are interconnected by redundant low-latency high-speed private fibre network links. Likewise, all AZs within a Region are interconnected by redundant private network links. This architecture built around redundancy ensures that clients enjoy high fault tolerance and business continuity in the event that one data centre within an AZ, or even an entire AZ within a Region, fails.

Hyperscale data centres to dominate. As the HSPs continue to gobble up market share and grow their AZs and Regions, this drives demand for hyperscale data centres, which are large-scale data centres each housing up to a few hundred thousand servers. Hyperscale data centres are projected to dominate the total data centre market in terms of processing power, storage and traffic. According to Cisco's Global Cloud Index, the number of hyperscale data centres will grow by 85% from 338 in 2016 to 628 by 2021, and represent 53% of the total number of data centres globally (up from 27% in 2016).

With APAC being the region with the highest projected data growth, demand for hyperscale data centres is projected to be the strongest in APAC (see Figure 3). These are also most likely to be located in Tier-1 cities with established power and connectivity infrastructure.

Figure 3: Hyperscale data centres in APAC likely to double from 2016-2021, driven by strong data traffic growth

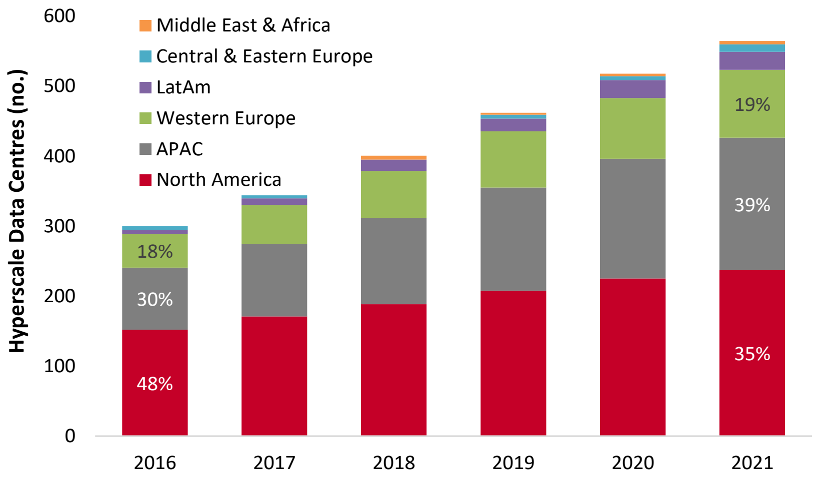

Source: Cisco Global Cloud Index, 2016-2021, KepCap Research

Moving computing resources closer to the source. Data centres are usually located in close proximity to Tier 1 cities due to the greater availability of power and connectivity and denser concentrations of urban populations and economic activities. Before demand in emerging markets reach levels critical enough for cloud and content service providers to justify building out large data centres, the usual practice is to leverage Content Delivery Networks (CDN) to reach out to those markets (see Box 2).

Even as hyperscale and multi-tenanted data centre capacities continue to increase, the sheer scale of data generation and the criticality required to process some of that real-time has spurred the development of edge computing. Edge computing refers to computing that is done locally at or near the source of data creation, thus reducing the need to route all the data traffic to and from the cloud. This means smaller data centres are now being built in secondary cities.

While latency can be reduced with CDNs, the need for higher compute powers at the edge will grow exponentially. The trend of developing smaller data centres at the edge will be accelerated by the commercial rollout of 5G mobile technology expected sometime between 2019 and 2020, with China and South Korea touted as global frontrunners. As 5G provides greater bandwidth and transmission speeds, the volume and velocity of data traffic resulting from IoT will multiply many fold. Edge computing allows some of that data to be processed almost instantaneously, such as for AI-enabled translation, facial recognition, VR, AR and autonomous vehicles (AV) which analyse inputs from a multitude of sensors to make thousands of critical decisions almost instantaneously to navigate through road traffic.

Box 2: What is a Content Delivery Network (CDN)?
A CDN is a network of distributed servers either owned by third-party service providers like Akamai or by content providers themselves. As an illustration, video-on-demand sites like Netflix cache or store popular movies locally on CDN servers across the world, rather than solely on their own servers in the US. When a consumer in Jakarta for example wishes to watch a movie, the movie is streamed from servers located in a CDN nearest to that consumer such as in Singapore, rather than directly from Netflix's host servers in the US. This speeds up the streaming and reduces the overall data traffic on the internet.

Cloud-to-edge continuum – a network of distributed computing
Bringing the concepts of cloud and edge computing together, we see the cloud and edge as opposite ends of a continuum which forms a network of distributed computing.

At the core of the network, hyperscale and colocation data centres undertake cloud workloads like machine learning and data analytics which are less time-sensitive, but require high compute powers and storage capacities.

Moving away from the core data centres and closer to end-users is the layer called the fog or near-edge, where processes requiring low to intermediate latency can be handled by servers in widely-distributed micro-data centres or colocation data centres. The near-edge also serves as the bridge between the core of the network and the edge, interconnected by fibre optics and wireless transmission.

Finally, the edge is where the network architecture will become more varied, especially with IoT sensors and 5G technologies. To handle ultra-low or low-latency processes, more compute powers will move closer to data sources and end-users, potentially residing in multiple micro data-centres next to mobile cell towers or increasingly within devices (see Figure 4).

Figure 4: Cloud-to-edge continuum serves various functional and latency requirements

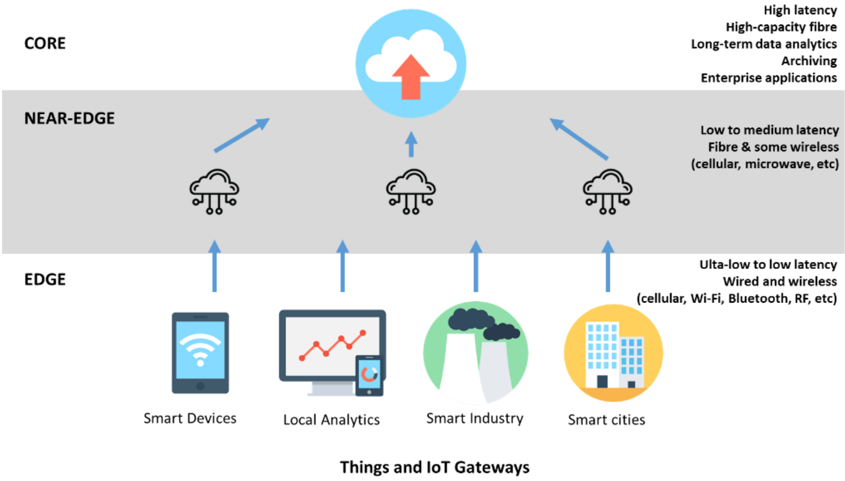

Source: 451 Research, KepCap Research

Some of the benefits to the cloud continuum's distributed network model are:

- distributing computing workloads reduces the burden on the overall network infrastructure preventing data traffic bottlenecks;
- edge computing frees up more resources for the cloud data centres at the core to perform compute-intensive tasks, such as machine learning and big-data analytics;
- the additional data generated from IoT sensors improves the quality of machine learning and big-data analytics performed in the cloud; and
- the distributed nature of the network enhances its robustness in the event of Distributed Denial of Services (DDoS) cyber-attacks, limiting them to the affected node(s) while the rest of the network remains unharmed.

Runway remains long for data centres

Asia Pacific is projected to be the fastest-growing geography for cloud workloads[2], growing at a 27% CAGR (2016-2021) and surpassing North America by 2021. Strong demand for cloud services in Singapore, Tokyo, Sydney and Hong Kong underpin data centre demand from HSPs, as they seek to improve redundancy and reliability by strengthening their AZs.

[2] Cisco Global Cloud Index: Forecast and Methodology, 2016-2021

In addition, Tier 1 cities like Singapore, Tokyo and Sydney typically benefit from strong connectivity (e.g. submarine cables, dark fibre), allowing them to serve as cloud Regions supporting neighbouring markets. For instance, most HSPs use Singapore as a Region for Southeast Asia to reach out to nearby markets like Jakarta, Bangkok and Kuala Lumpur. China remains a market with vast potential as the government pushes for accelerated developments in the internet-based economy. Largely dominated by local players such as Alibaba Cloud and Tencent, AWS and Microsoft have made some progress collaborating with local operators.

However, this does not preclude the HSPs from establishing some presence directly in edge markets via colocation to meet requirements such as data sovereignty and catering to burgeoning internet populations. Vietnam and Indonesia stand out as markets where demand for reliable data centres continue to grow.

Cloud workloads in Western Europe are projected to grow at 20% CAGR from 2016 to 2021, slightly outpacing the 19% CAGR expected in North America[3]. Frankfurt and Amsterdam remain important data centre markets as providers leverage the presence of two of the world's largest internet exchanges, namely Deutscher Commercial Internet Exchange (DE-CIX) and Amsterdam Internet Exchange (AMS-IX). London retains its position as one of Europe's largest data centre markets and despite uncertainties over Brexit, local demand from HSPs and enterprises remain keen. Meanwhile, other European data centre markets drawing the attention of HSPs include Madrid, Milan, Hamburg and Dublin. These smaller markets tend to serve as edge locations for the HSPs and Dublin also presents an alternative location to the UK in the event of Brexit.

Stay vigilant about emerging trends. While demand for more data centres is evident, operators must stay disciplined and avoid speculative overbuilding especially in the short run. Operators must also safeguard against large-scale event risks such as cyber-attacks and technological failures which may result in data loss and/or downtime.

Moore's Law, which suggests that processing power doubles every two years, may have slowed down, but operators must continue to improve power density and energy efficiency to stay competitive. This means that data centres themselves will become smarter as more operators deploy IoT and AI-enabled data centre infrastructure management (DCIM) systems to improve efficiencies.

Operators and investors looking to capture edge demand must also stay updated on architectural developments as the industry continues to explore which data centre form factors work best in the IoT age.

Over the longer term, investors must bear in mind that growth of hyperscale data centres in key cities will eventually peak. Some emerging markets will develop into new data centre hubs, but these will likely come with higher execution risks in the early years.

[3] Ibid.

Conclusion

Globally, data is being generated, consumed, and processed faster than ever, and mainstream adoption of cloud computing has only just begun. The vast volumes of useful data traffic congregating in the cloud for compute-intensive processes or to be stored will spur further demand for hyperscale data centres, especially in key gateway cities. The acceleration of IoT and exploding number of connected devices will spur more smaller-sized data centres in secondary markets for computing to be done at the edge or near-edge. Hyperscale cloud providers continue to compete with one another through the provision of better connectivity via subsea cable investments and more Availability Zones and Regions. These will be supportive of strong demand for both hyperscale data centres and colocation data centres.

Looking ahead, the runway for data centre demand not only remains intact – it looks primed to take off.

Keppel Capital is a premier asset manager in Asia. It is the asset management arm of Keppel Corporation Limited, a multinational conglomerate with key businesses in Offshore & Marine, Property, Infrastructure and Investments, providing innovative solutions for sustainable urbanisation.

With assets under management of approximately S$29 billion, Keppel Capital has a diversified portfolio that includes real estate, infrastructure and data centre properties in key global markets.

Keppel Capital aims to create value and deliver sustainable returns for institutional and retail investors through a range of products including REITs, business trusts, private funds investing in real estate and infrastructure, separate accounts and pooled investment vehicles.

The asset managers under Keppel Capital include Keppel REIT Management Limited, Alpha Investment Partners Limited, Keppel Infrastructure Fund Management Pte. Ltd., Keppel DC REIT Management Pte. Ltd., and Keppel-KBS US REIT Management Pte. Ltd.

Keppel Capital
Trusted Investment Partner

Members of Keppel Capital:
Keppel REIT Management Limited
Alpha Investment Partners Limited
Keppel Infrastructure Fund Management Pte. Ltd.
Keppel DC REIT Management Pte. Ltd.
Keppel-KBS US REIT Management Pte. Ltd.

Our Offices • **Singapore** • **Australia** • **China** • **Japan** • **Korea** • **Taiwan**