Hyperscale meets hyperlocal: Value maximization in the data centre boom

February 2026

The better the question. The better the answer. The better the world works.

EY Parthenon

Shape the future with confidence

Contents

EY Parthenon

1

Executive summary

Data centres supporting hyperscale public cloud in key hubs are a premium return asset class, expected to experience strong organic rent growth with limited customer churn

Hyperscale public cloud data centres as an asset class

Hyperscale public cloud data centres are an attractive asset class expected to experience strong organic rent growth with limited customer churn for four key reasons:

1. Data centre ecosystem and public cloud: **Data centres housing critical public cloud infrastructure in key hubs are fundamental for economic growth**, catering to rapidly growing compute, network and storage needs.
2. Hyperscale profit incentive: **Public cloud services generate outstanding economics for hyperscalers** – large cloud services providers that operate massive cloud infrastructure. These companies earn high margins from a sticky end customer base, supporting robust share price performance and incentivising continued capacity expansion. Outsourced data centre rent remains a modest cost for hyperscalers, driving organic rent growth for data centre operators.
3. Hyperscale public cloud as fixed infrastructure: **Hyperscale public cloud data centre architecture is a fundamentally fixed infrastructure with significant switching costs and risks;** public cloud data centre locations are driven by latency needs (physically limited by the speed of light), competitive dynamics, and existing multi-billion-dollar networks and ecosystems which are intrinsically difficult to replicate.
4. Market conditions and supply constraints: **Public cloud data centres in key hub locations are positively differentiated because of power and land constraints**, as structural limitations on new power grid infrastructure and planning restrict additional supply in these high-barrier markets.

Data centres are physical buildings, housing servers that run applications essential to daily life, including mission-critical, cloud-hosted business and government operations

1 **Data centre ecosystem and public cloud:** Ecosystem overview

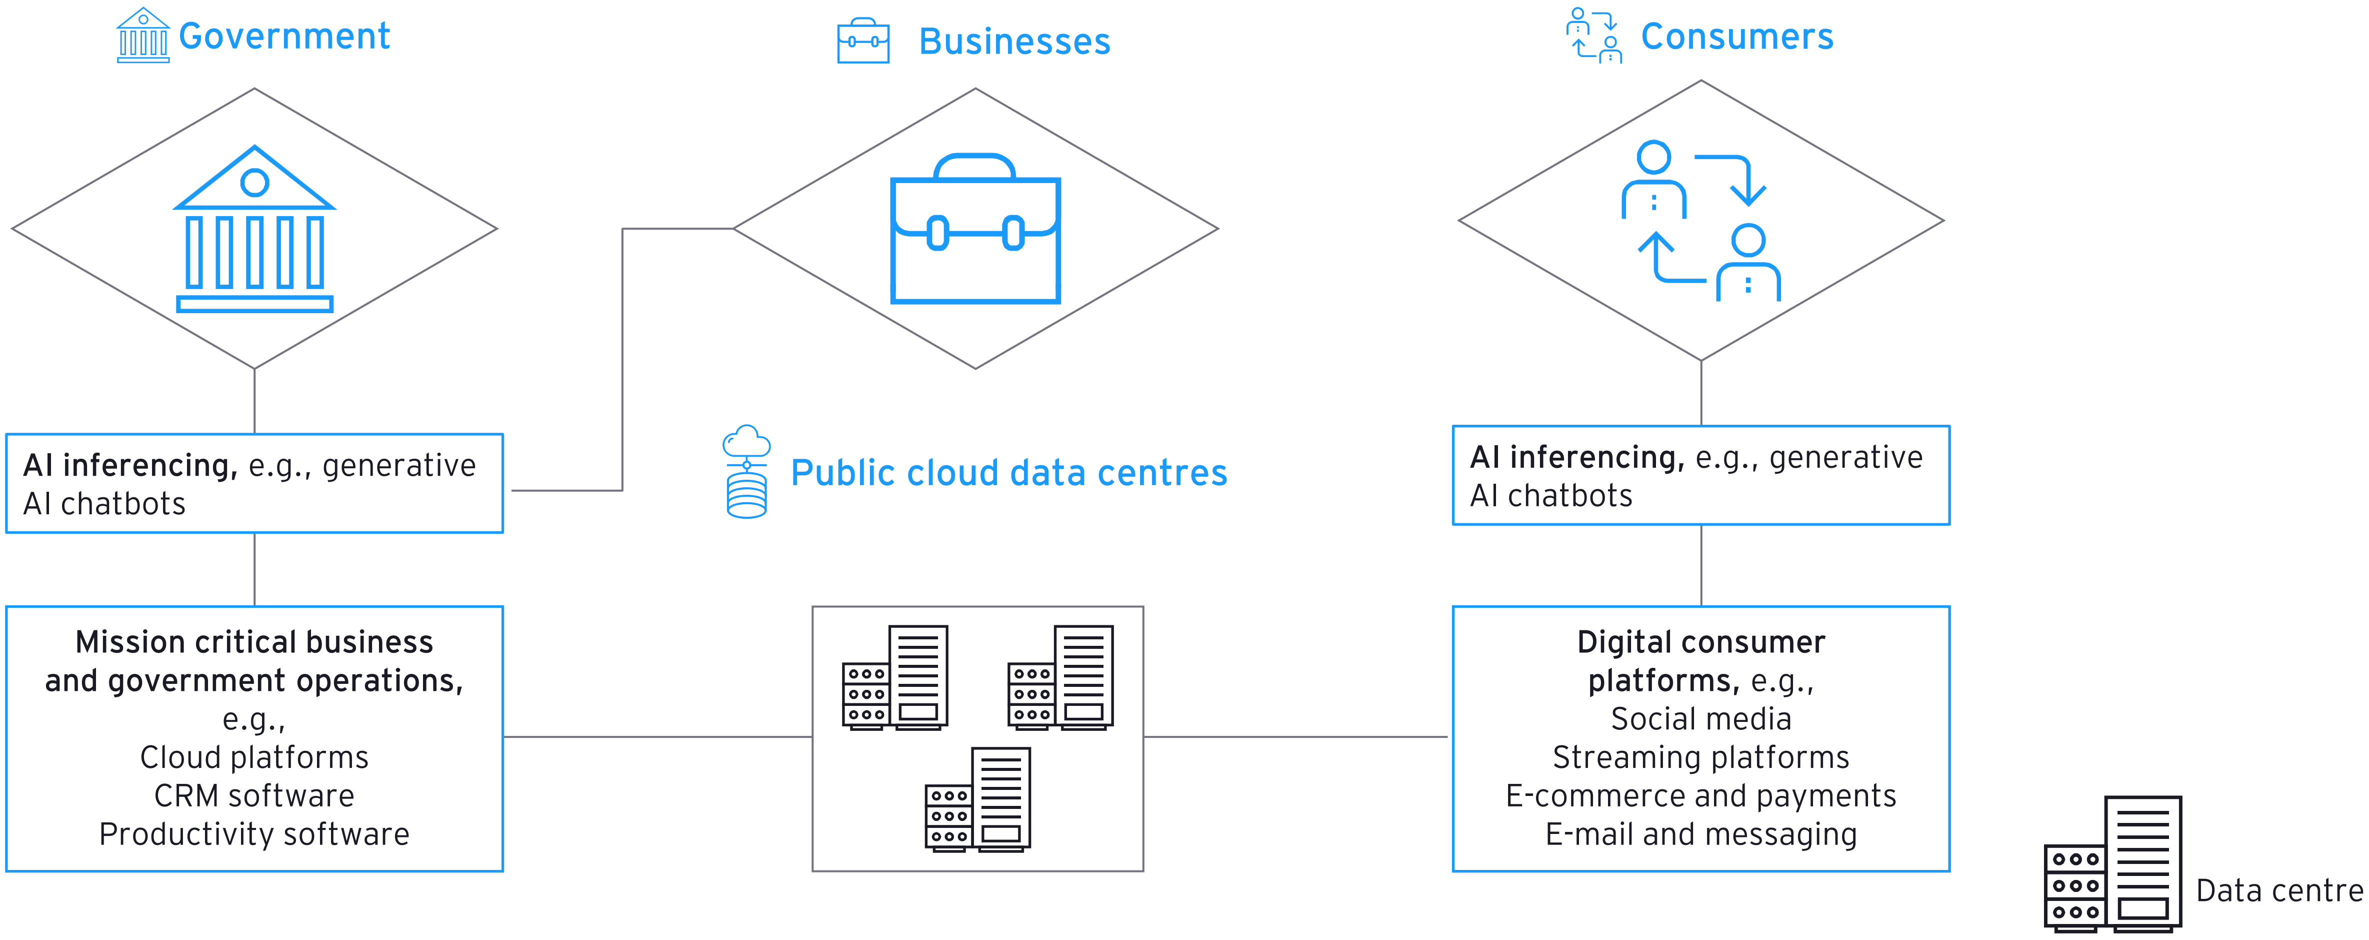

EY Parthenon

Mission-critical public cloud and AI workloads are typically hosted in hyperscale infill facilities

1 **Data centre ecosystem and public cloud:** Types of data centres

	Hyperscale Hyperscale refers to size (>5MW and recently 100MW+)		Interconnection-focused	Enterprise-focused
Description	**Large data centres** selling space and power to cloud/content companies with large capacity needs **The largest cloud service providers (CSPs)** are known as hyperscalers, for the multi-MW facilities they operate		Data centres where **customers interconnect with each other** and/or require low latency connectivity	Typically multi-tenant data centre facilities focused on **smaller enterprise customers**
Example operators and owners	Cloud service providers	Large colocation data centre operators	Telcos and ecosystem data centre facilities	Multi-tenant data centres
Key sub-categories	Infill facilities	Regional facilities	–	–
Workloads served	▪ Public cloud ▪ AI inference ▪ AI training	▪ Data storage ▪ Batch processing ▪ AI training and inference	▪ Communication ▪ Financial transactions ▪ Real-time (e.g., online gaming)	▪ Internal applications ▪ Internal databases ▪ Financial transactions

Mission-critical public cloud runs only in hyperscale infill sites. AI training and inference typically occur in cloud environments but can also run in regional facilities.

Source: EY-Parthenon research and analysis

EY Parthenon

The volume of data created, stored, and consumed continues to grow exponentially, driving the need for scalable IT architecture offered by the public cloud

1 **Data centre ecosystem and public cloud:** Compute, network, and storage volume projections

Compute

Compute workload growth generates demand for scalable public cloud hosted resources

Cloud workload growth, index 2018 = 100 workloads (#, 2018 – '28F)

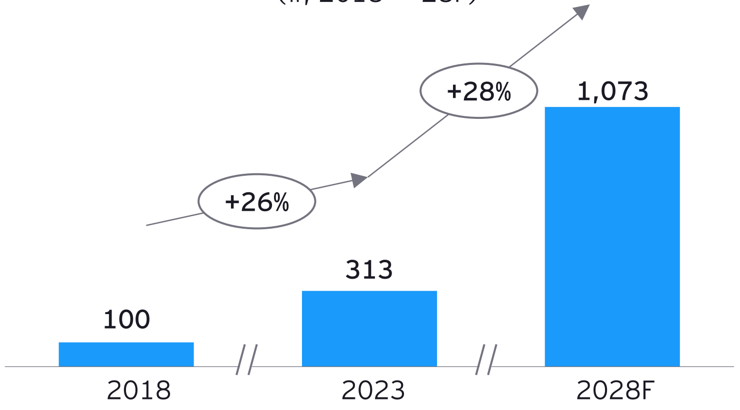

Networking

Demand for reliable real-time information drives cloud networking infrastructure demand

Global IP traffic (Exabytes, 2018 – '28F)

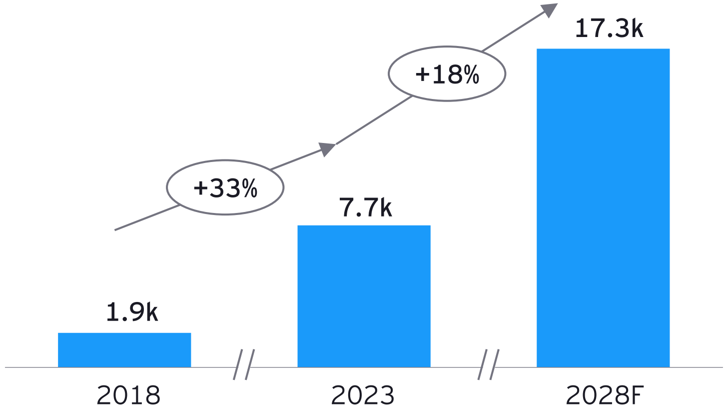

Storage

Data volume growth drives the need for scalable storage available over public cloud

Data created, captured, replicated and consumed worldwide[1] (Zettabytes, 2018 – '28F)

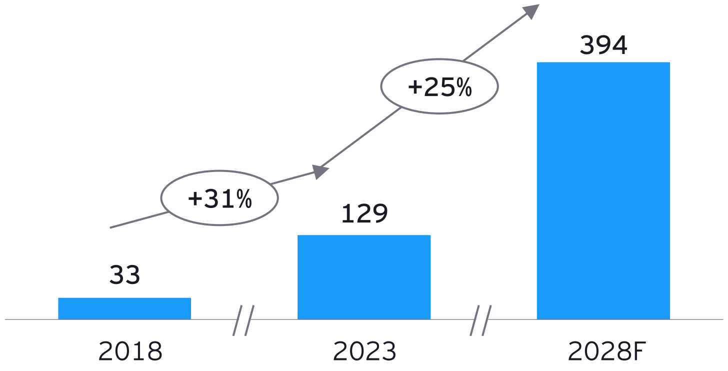

Example use cases and enterprise customers

- AI training and inference
- Analysing consumer behaviour
- Compiling code
- Web hosting

Large enterprises, CPSs, software firms

- Stock trading
- Connecting multiple devices
- Running a video conference call
- Collaborating on a shared drive

Social media, telcos, productivity software

- Media storage
- Data encryption
- E-commerce payments
- Processing payrolls and billing

Payment platforms, database platforms

[1] Based on IDC Worldwide Global DataSphere, measuring how much data is created, captured, replicated and consumed each year. This includes data stored on private and public IT infrastructures, utility infrastructures, and personal computing devices, including laptops, tablets, smartphones, and IoT devices

Sources: IDC Worldwide Global DataSphere; Cisco; Nokia Bell Labs; EY-Parthenon research and analysis

Hyperscalers have seen strong share price gains fuelled by high-margin cloud growth, outsourced data centre rent makes up a fraction of their cost base

2 **Hyperscale profit incentive:** Financial overview

Hyperscalers' share price growth is driven by cloud growth …

	Cloud revenue (Growth FY19-24)	Share price (Growth FY19-24)
Player 1	+207%	+137%
Player 2	+124%	+234%
Player 3	+385%	+183%
Player 4	+37%	+132%

Cloud represents ~$280bn revenue (FY24) for the largest hyperscalers, highlighting its strategic importance

… which is a high margin business …

Operating margin before tax (indicative)

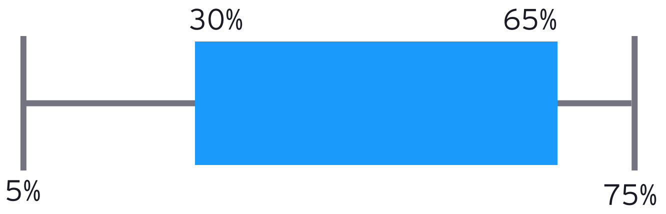

… where outsourced rent is a fraction of the cost base

Outsourcing operating costs (% of total)

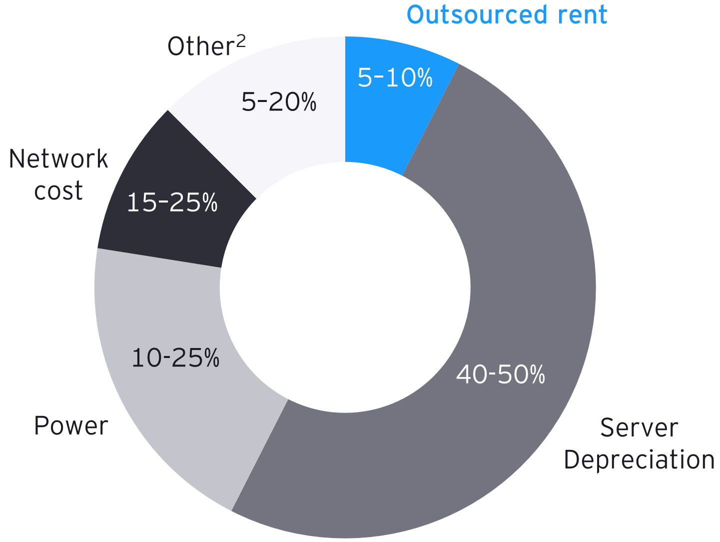

Insufficient capacity risks foregone revenue, earnings growth, and ultimately share performance

[1] Annualised revenue for CSP 1 and CSP 2, based on YTD run-rate
[2] E.g., maintenance, remote hands services
Source: Annual reports; Analyst reports; Earnings calls; EY-Parthenon research and analysis

Recent organic rent growth has been driven by demand intensity, supply shortages, and tenant price flexibility, which are expected to persist

2 **Hyperscale profit incentive:** Organic rent growth drivers

Accelerated organic rent growth drivers

Growth driver	Demand intensity	Market supply shortage	Hyperscale tenant price flexibility
Description	In markets where international hyperscalers compete for capacity, **operators have greater bargaining power**	Key constraints include **power and land availability**	Hyperscalers are **willing to absorb higher rent** increases due to the importance of this high margin product and the operational complexities of migrating workloads

"Hyperscalers continue to deploy record levels of capital into infrastructure expansion, with AI workloads accelerating demand for high-density capacity. This is driving a structural uplift in leasing velocity and pre-commitment levels across Tier 1 markets.

European Broker

"In key European hubs, the availability of power and land is increasingly constrained, with moratoriums in place. These constraints are expected to persist, reinforcing pricing power for incumbent operators.

European Private Equity Principal

"Speed-to-market is essential for hyperscalers, for whom a two-year delay in revenue is more significant than a rent increase, since this constitutes a very small proportion of their total cost stack.

US-based Hyperscale Data Centre Operator

[1] A cloud service provider (CSP) is a third-party company that delivers on-demand, scalable computing resources – such as compute, storage, platforms, and applications – over a network

Source: EY-Parthenon research and analysis

EY Parthenon

Public cloud architecture is inherently sticky due to high switching costs, regulatory constraints, high capex, and latency requirements

3 **Hyperscale public cloud as fixed infrastructure:** Overview

Switching barriers

Switching facilities adds operational complexity and may delay capacity expansion for new demand.

New facility	Operational complexity	
$1.4b deployment	Direct cost $20–40m	+ Material & uncertain impact !

~40MW facility

Multi-billion-dollar sunk investments

Availability zones require significant capital ($6b+) as well as development of a multi-billion-dollar modern infrastructure ecosystem over decades

~$5b	~$1b	+$?b
DC servers & facilities	Fibre & interconnection	Power, transport, etc.

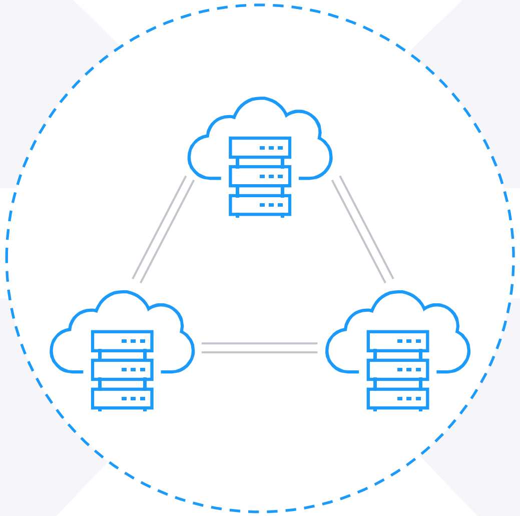

Regulatory standards

Data localisation regulations restrict cross-border data transfers, driving location specificity of public cloud facilities

Public service data must remain within the UK (aligned with EU GDPR)

Company financial records must be stored in Germany

Latency requirements

Locations are constrained by 1) technical intra-AZ latency, and 2) customer demand for low latency leading to cloud proximity competition

<100 km	10–60km	<10 km
Optimal AZ-to-Enterprise distance	Between AZs distance	Between DCs in an AZ requirement

Many cloud services rely on synchronous communication between data centers, which require short distances for ultra-low latency

Source: EY-Parthenon interview programme; EY-Parthenon research and analysis

EY Parthenon

CSPs select specific site locations to solve for high availability, low latency, fault preventing, and resilient load balancing within their distributed architectures

3 **Hyperscale public cloud as fixed infrastructure:** Availability zones (AZs)

Typical cloud region and availability zone architecture

Cloud region: 2+ AZs in a geographically distinct area

Availability zone (AZ): 1+ independent DC with separate power, networking and cooling

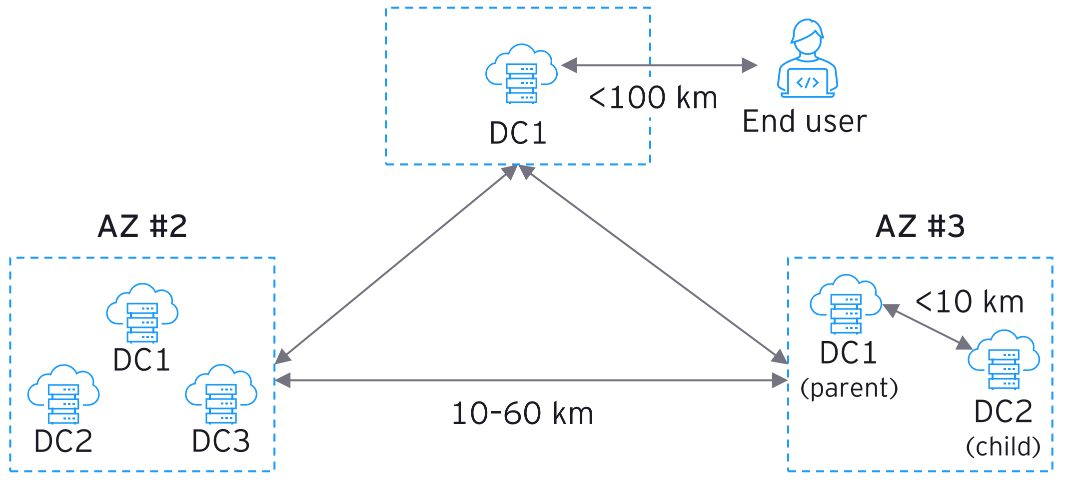

Availability zone design criteria

Distance to End User: Low latency

Close proximity to user populations through well-distributed regions/edge locations

→ Minimizes transmission delays, improves responsiveness, throughput, and reliability

Distance between AZs: Fault tolerance and disaster recovery

Multiple, physically separate data centers with proximate location for low-latency. 3 distinct fiber connections.

→ Mitigates single points of failure, enables fault tolerance and rapid disaster recovery

Distance across AZs in a region: Effective load balancing

Distribution of workloads across multiple AZs with real-time traffic routing and dynamic resource allocation

→ Balances system load, avoids congestion, and enhances overall performance and uptime

AZ locations should be near customers, near other facilities in the AZ, and in reach of the cloud region architecture, limiting the set of potential locations and driving stickiness

Source: Company websites; EY-Parthenon interview programme; EY-Parthenon research and analysis

Supply in key markets is constrained by structural limitations on new power grid infrastructure, available land, and permitting processes

4 **Market conditions and supply constraints:** Overview

Power availability

Utilities are struggling to adjust to growing demand for power and are expected to continue to lag data centre needs

"

In Northern Virginia, which is the largest cluster of data centres, **there is no power available.** All new power coming online in the next 3 years is already pre-leased … the grid company cannot keep up with the demand

Hyperscaler

Northern Virgina faces constraints due to challenges around transmission infrastructure where the utility, Dominion Energy cannot guarantee power delivery

Land availability within acceptable zones

As demand grows, plots of land available which meets the needs of hyperscalers become increasingly limited

"

West London … has the highest number of individual data centres and the largest deployed capacity in Western Europe … there is **almost no near-term available land** to support further development, as it's completely planned.

Broker

Amsterdam has implemented a series of data centre moratoriums. As of April 2025, the municipality of Amsterdam has imposed a full ban on new data centres and expansions within city limits

Permitting processes

In some markets, local authorities are introducing permits or restrictions to make data centre construction more challenging

"

If hyperscalers aim to build on a brownfield site in the UK, the planning permit authorities often react negatively once they realise it's for a hyperscale data centre. This reaction can make it **more difficult for hyperscalers to build.**

Operator

Dublin data centre development has faced permitting constraints from planning rejections, appeals backlogs, and environmental pressures

Supply constraints can create a seller's market, increasing pricing power. Constraints are expected to be sustained in the future.

Source: EY-Parthenon interview programme; EY-Parthenon research and analysis

2

Data centre ecosystem and public cloud

Data centres form the core of the digital ecosystem. They enable secure storage, scalable cloud infrastructure, and resilient workload placement

1 **Data centre ecosystem and public cloud:** Overview

Data centre ecosystem and public cloud: **Data centres housing critical public cloud infrastructure in key hubs are fundamental for economic growth,** catering to rapidly growing compute, network and storage needs.

A Data centre ecosystem

Data centres are physical buildings housing servers that run applications essential to daily life, including mission-critical, cloud-hosted business and government operations

- **Data centres house servers** that store information, host applications and compute large amounts of data; they are the **backbone of the digital economy**
- Data centres can be classified as **hyperscale**, **interconnection-focused**, and **enterprise** focused facilities; mission-critical cloud and AI workloads are typically hosted in hyperscale facilities
- Investment in data centres lays the foundation for the **digital economy** and **drives economic growth**
- The **volume of data** created, stored, and consumed continues to **grow exponentially**, fuelling the continued need for scalable IT architecture offered by public cloud

B Public cloud

Global data centre demand is forecast to grow ~13% p.a. over the next 10 years, with public cloud and AI workloads each contributing roughly half of this growth

- The **cloud** is a **virtual environment** of servers, software, and databases; workloads running on public cloud infrastructure are the **foundation of modern business models,** with **no comparable alternative** in terms of scalability and cost effectiveness offered, driving end customer stickiness
- **Generative AI** applications are commonly **deployed within existing cloud infrastructure**, providing a predictable workload transition from training to inference and further cloud for hyperscale facilities operating in local specific availability zones

Data centres are buildings that house servers for businesses; these servers store information, host applications, and compute large amounts of data

1 A Data centre ecosystem: Introduction

Mobile networks

- Data centres support mobile networks by processing and storing data

Hyperscale public cloud

- Hyperscale public cloud is delivered through large data centres organised into availability zones[1], providing mission-critical enterprise computing
- Governments, businesses, and consumers rely on hyperscale data centres to utilise cloud-hosted applications and artificial intelligence (AI) workloads

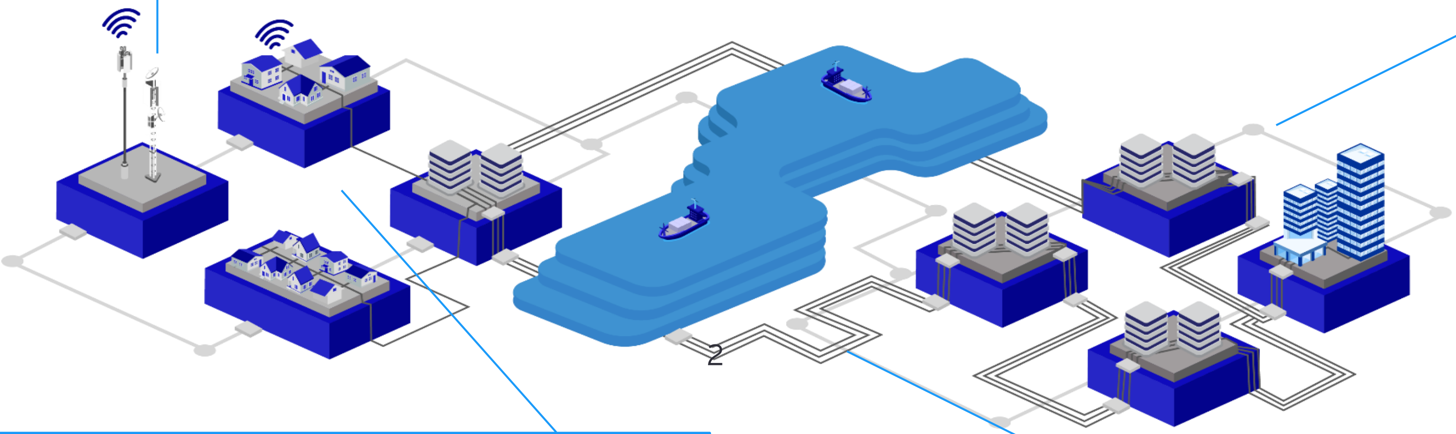

Enterprise computing

- Data centres support enterprises by providing robust infra. for data processing, storage, and management; data centre offerings include cloud services, data analytics, and disaster recovery

Retail wireline networks

- Data centres underpin home internet by managing and storing data, ensuring fast and stable connections for streaming, remote work and more

[1] An availability zone is a distinct, isolated data centre location within a cloud region, designed for high availability and fault tolerance

[2] Subsea cables connect to subsea landing point DCs, which process and route data, enabling seamless international data exchange and supporting global internet infrastructure

Source: EY-Parthenon research and analysis

Mission-critical public cloud and AI workloads are typically hosted in hyperscale infill facilities

1 **Data centre ecosystem and public cloud:** Types of data centres

	Hyperscale Hyperscale refers to size (>5MW and recently 100MW+)		Interconnection-focused	Enterprise-focused
Description	**Large data centres** selling space and power to cloud/content companies with large capacity needs **The largest cloud service providers (CSPs)** are known as hyperscalers, for the multi-MW facilities they operate		Data centres where **customers interconnect with each other** and/or require low latency connectivity	Typically multi-tenant data centre facilities focused on **smaller enterprise customers**
Example operators and owners	Cloud service providers	Large colocation data centre operators	Telcos and ecosystem data centre facilities	Multi-tenant data centres
Key sub-categories	Infill facilities	Regional facilities	–	–
Workloads served	▪ Public cloud ▪ AI inference ▪ AI training	▪ Data storage ▪ Batch processing ▪ AI training and inference	▪ Communication ▪ Financial transactions ▪ Real-time (e.g., online gaming)	▪ Internal applications ▪ Internal databases ▪ Financial transactions

Mission-critical public cloud runs only in hyperscale infill sites. AI training and inference typically occur in cloud environments but can also run in regional facilities.

Source: EY-Parthenon research and analysis

Investment in data centre infrastructure enables a strong digital economy, unlocking economic growth

1 A **Data centre ecosystem:** Data centres, digital economy, and economic growth

- Future economic growth relies on a digital economy powered by an advanced data centre ecosystem
- Data centres are estimated to have contributed **>$500b[1]** to US GDP in 2023

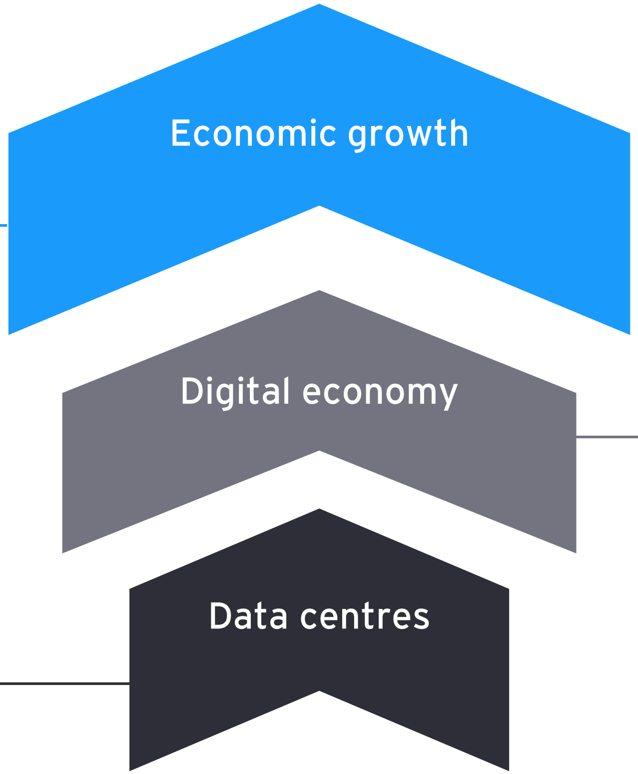

- The OECD defines the digital economy as 'all **economic activity reliant on**, or significantly enhanced by the use of **digital inputs**, including digital technologies, digital infrastructure, digital services, and data'
- The US digital economy is worth **~$2.5tn** (~10% of GDP)

- Data centres form **the backbone of the modern digital economy**, providing infrastructure to process and store vast amounts of data
- The largest 4 CSPs are projected to spend **>$320b on data centres in 2025**, up from $246b in 2024, as AI investment accelerates

[1]The Data Centre Coalition's dataset covers contributions from 2017 to 2021. Contributions for 2022 and 2023 were estimated using the average CAGR of 8% observed from 2017-21
Source: Bureau of Economic Affairs; Linklaters; Nutanix; OECD; Data Centre Coalition; Financial Times; EY-Parthenon research and analysis

The volume of data created, stored, and consumed continues to grow exponentially, driving the need for scalable IT architecture offered by the public cloud

1 A **Data centre ecosystem:** Compute, network, and storage volume projections

Compute

Compute workload growth generates demand for scalable public cloud hosted resources

Cloud workload growth, index 2018 = 100 workloads (#, 2018 – '28F)

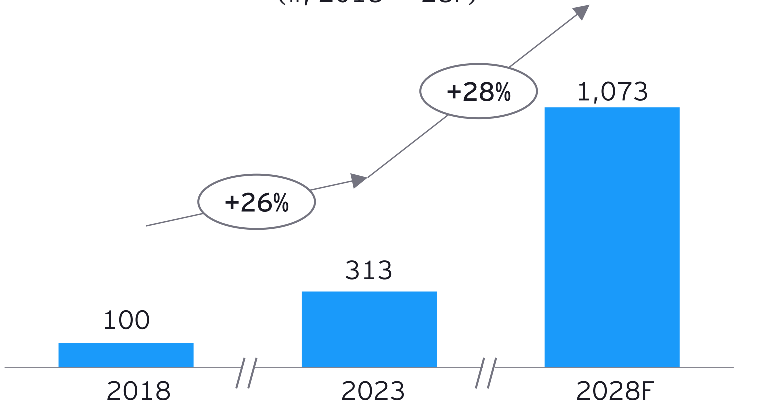

Networking

Demand for reliable real-time information drives cloud networking infrastructure demand

Global IP traffic (Exabytes, 2018 – '28F)

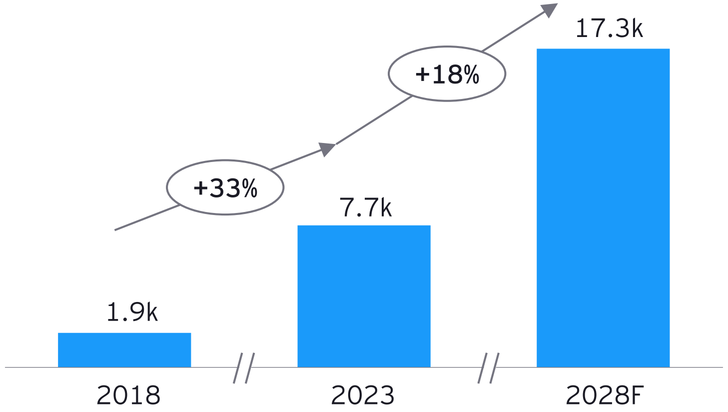

Storage

Data volume growth drives the need for scalable storage available over public cloud

Data created, captured, replicated and consumed worldwide[1] (Zettabytes, 2018 – '28F)

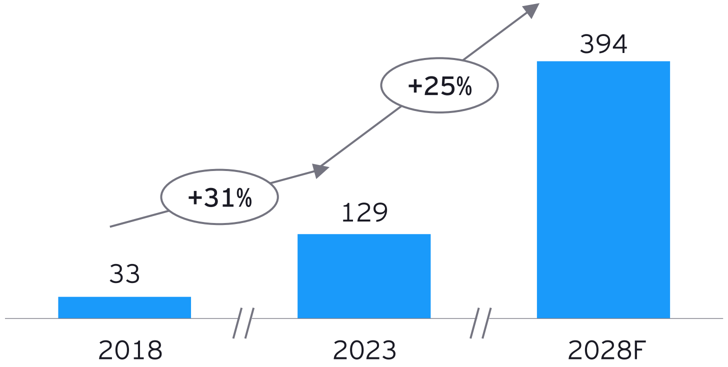

Example use cases and enterprise customers

- AI training and inference
- Analysing consumer behaviour
- Compiling code
- Web hosting

Large enterprises, CPSs, software firms

- Stock trading
- Connecting multiple devices
- Running a video conference call
- Collaborating on a shared drive

Social media, telcos, productivity software

- Media storage
- Data encryption
- E-commerce payments
- Processing payrolls and billing

Payment platforms, database platforms

[1] Based on IDC Worldwide Global DataSphere, measuring how much data is created, captured, replicated and consumed each year. This includes data stored on private and public IT infrastructures, utility infrastructures, and personal computing devices, including laptops, tablets, smartphones, and IoT devices

Sources: IDC Worldwide Global DataSphere; Cisco; Nokia Bell Labs; EY-Parthenon research and analysis

Global demand for hyperscale infill data centre capacity is forecast to grow ~13% p.a. to 2034, with public cloud and AI workloads each contributing half of this growth

1 B **Public cloud:** Global demaand forecast and drivers

Global hyperscale infill data centre capacity demand, by workload (GW[1], 2024-34, excl. China)

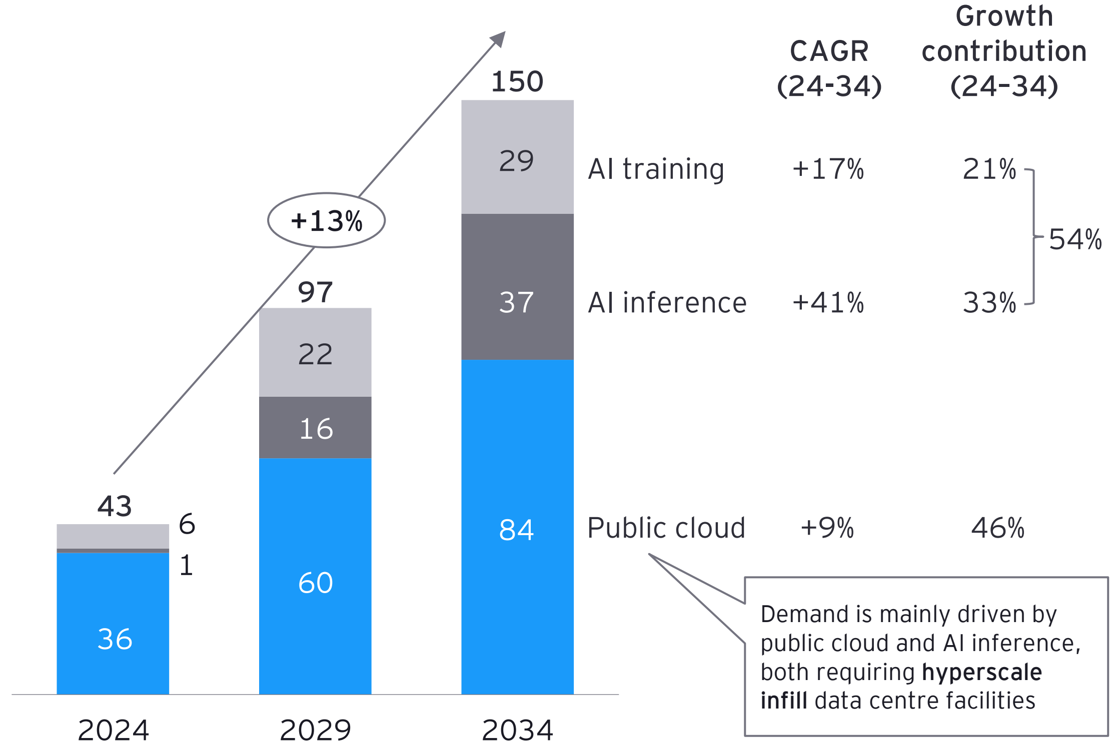

Demand drivers

AI

- Surging AI workloads are accelerating infill and new data centre builds, with **operators securing sites ahead of future demand**
- **AI training** requires compute-intensive infrastructure and less latency sensitivity, driving demand in more **cost-effective, higher capacity locations**
- **AI inference** needs low-latency proximity to users and scalable deployment, driving **increasing demand for new builds** in **established cloud regions**

Public cloud

- Rising demand is driving hyperscalers to scale capacity in established cloud regions
- Requirements to comply with data residency regulations, minimize latency, and tap into new markets are driving creation of availability zones **in developing regions**

Historical order of deployment

[1] GW stands for gigawatt, a unit of power used to express overall power demand or capacity requirements
Source: EY-Parthenon Digital Infrastructure proprietary data centre supply database; EY-Parthenon research and analysis

EY Parthenon

The public cloud supports mission-critical business and government operations, as well as digital consumer platforms

1 B **Public cloud:** Hyperscale data centre ecosystem overview

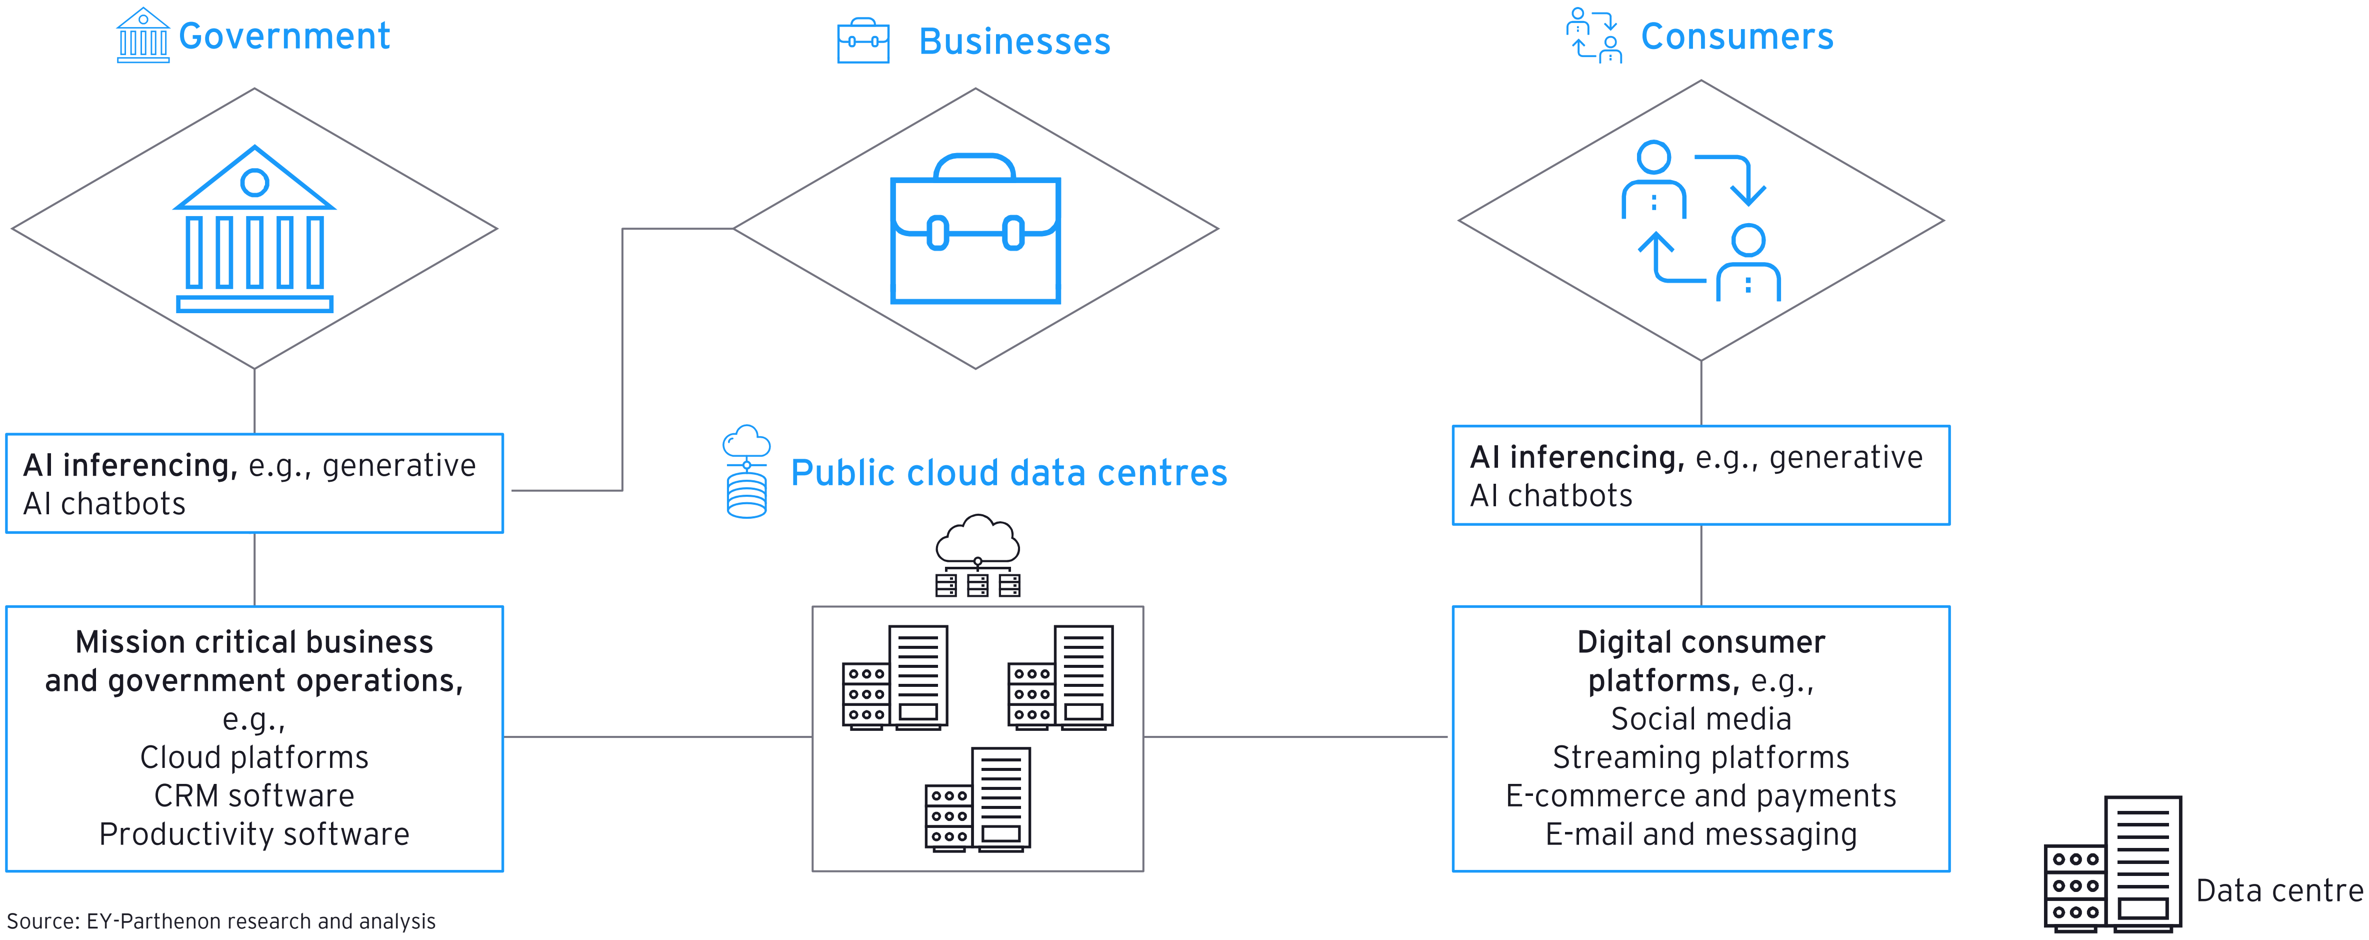

Source: EY-Parthenon research and analysis

The personal and professional day-to-day life in developed economies relies on the public cloud

1 B **Public cloud:** Data centre example workloads served

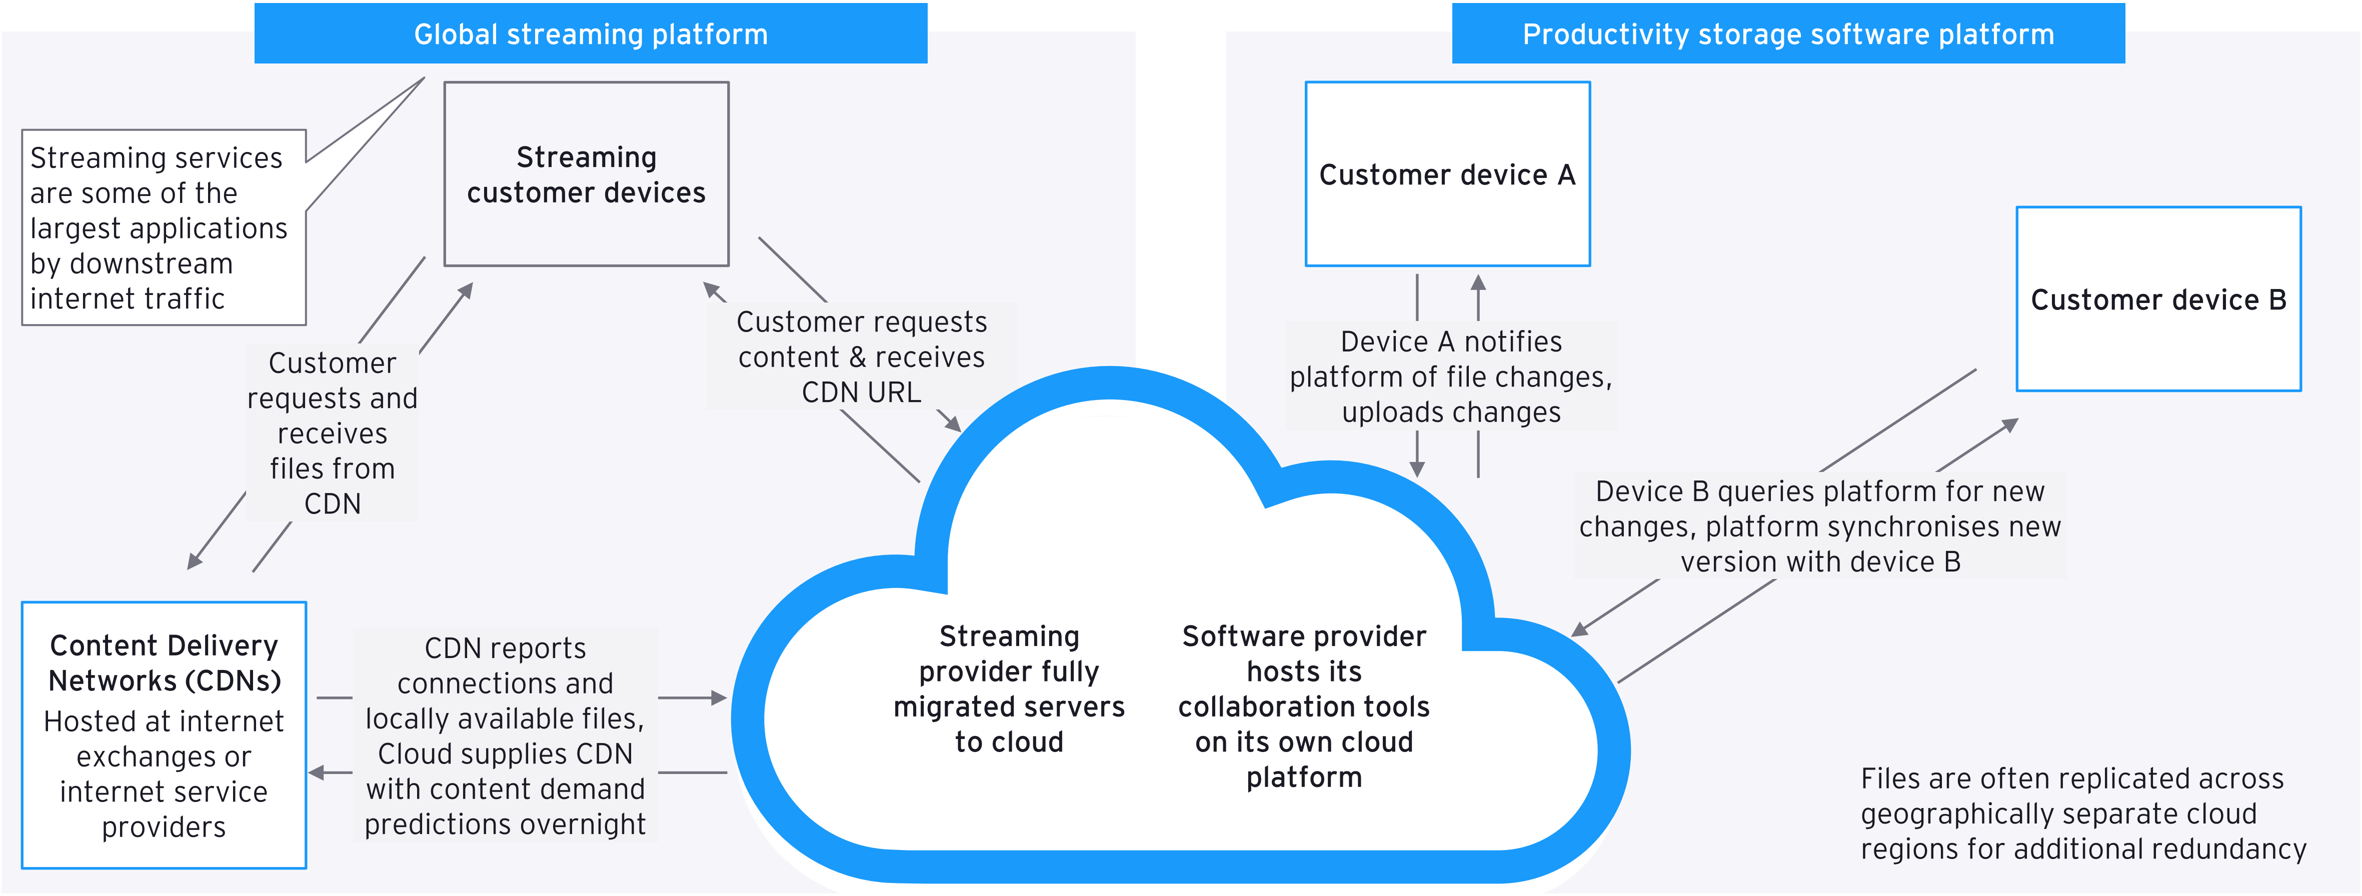

Source: TechTarget; DgtlInfra; Data Center Dynamics; Sandvine; EY-Parthenon research and analysis

EY Parthenon

The distributed architecture of public cloud operators allows multi-national customers to scale operations, optimize cost and redundancy, and serve a global customer base

1 B **Public cloud:** Geographic scalability and enterprise cloud adoption

Cloud regions by hyperscaler

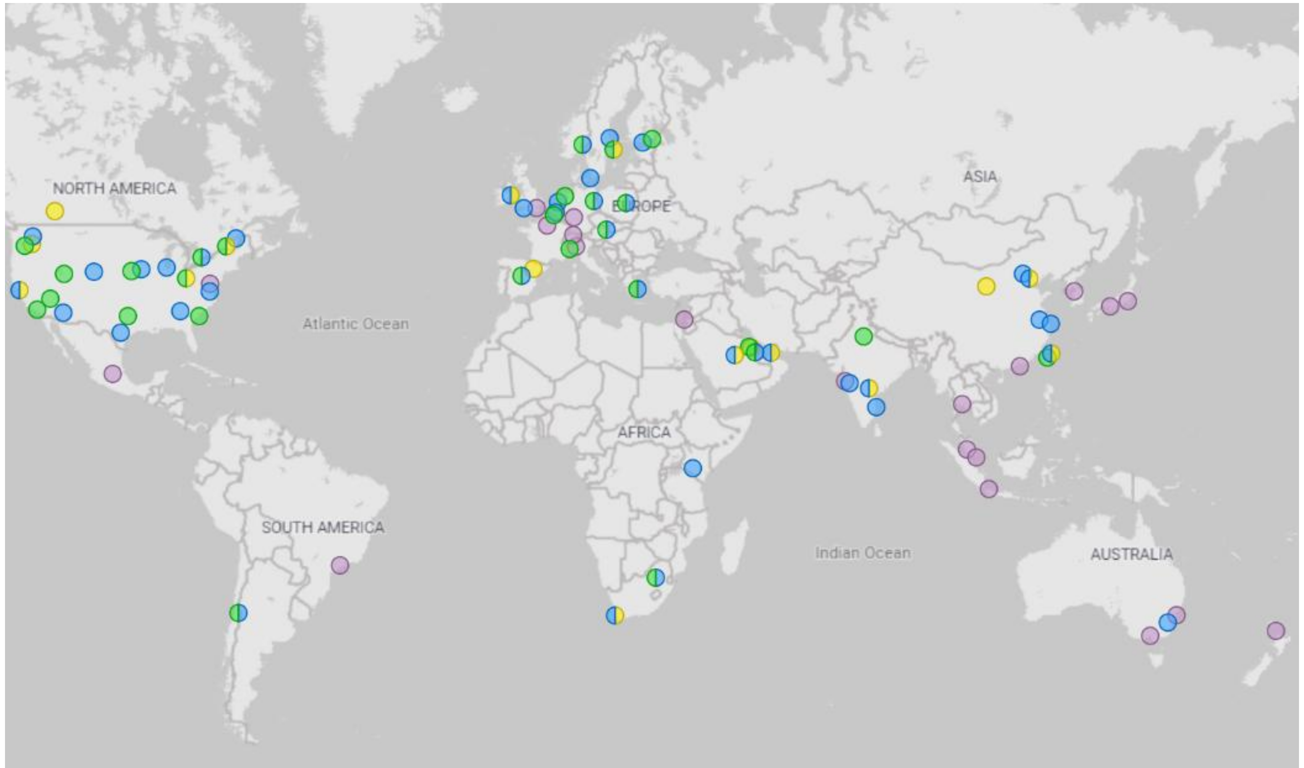

Key: CSP A CSP B CSP C All three

Source: Cloud Infrastructure Map; Company websites; TomTom; EY-Parthenon research and analysis

Enterprise public cloud adoption

Enterprises are increasingly adopting the public cloud as an alternative to on-premises servers …

Public cloud services spend, including both cloud costs and applications ($b, 2020-28)

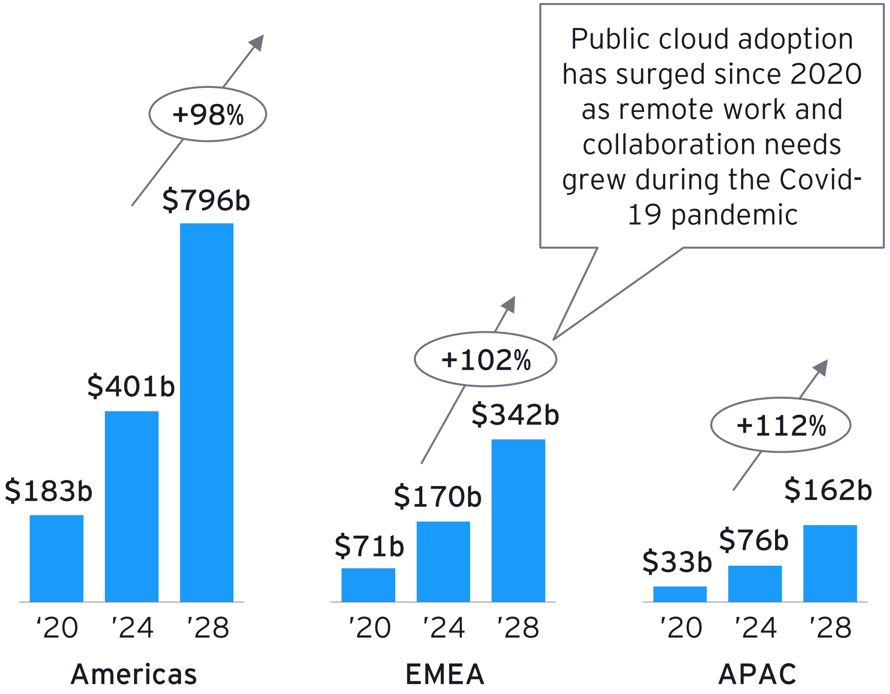

Public cloud is a strategic enterprise asset, with long-term stickiness driven by business-critical adoption drivers and complex migration journey

1 B **Public cloud:** Adoption drivers and journey

Key enterprise public cloud adoption drivers

Design flexibility and scalability

- Enterprises choose outsourcing levels to matches their service needs
- Public cloud offers flexible solutions for ensuring disaster recovery and business continuity
- Flexibility allows businesses to rapidly scale operations across regions

Access to new technology

- Cloud facilitates various applications, ranging from content to AI
- Cloud platforms enable access to next-generation chips for high-compute workloads, bypassing procurement challenges
- Some services, such as advanced cybersecurity, are cloud exclusive

Cost optimisation

- Cost optimisation, such as a pay-as-you-use model, eliminates upfront CapEx
- Maintenance costs are minimised, since suppliers handle infrastructure maintenance and updates

Cloud adoption journey

	Steps	Description	Duration
	Plan adoption strategy and deployment model	▪ Choose workloads use cases ▪ Select cloud approach (single, multi, or hybrid)	~2-3 wks.
	Cloud platform selection	▪ Choose a cloud provider, validate Proof-of-Concept (PoC)	~2-8 wks.
	Cloud adoption planning & design	▪ Assess resource needs, develop risk management plans, and establish KPIs	~3-4 wks.
	Transferring to the cloud	▪ Set up production environment and transfer data continuously	~4-22 wks.
	Final testing & optimisation	▪ Performance testing, compliance checks, and security optimisation	~1-8 wks.

Transferring from on-premise to the cloud is a mission-critical process involving high costs and a significant time commitment, creating a stickiness for users, attracting further investment and organisational integration

Source: EY-Parthenon research and analysis

EY Parthenon

Generative AI is a subset of AI with applications usually deployed on public cloud-hosted IT infrastructure; computing workloads include training and inference

1 B **Public cloud:** Generative AI overview

Context

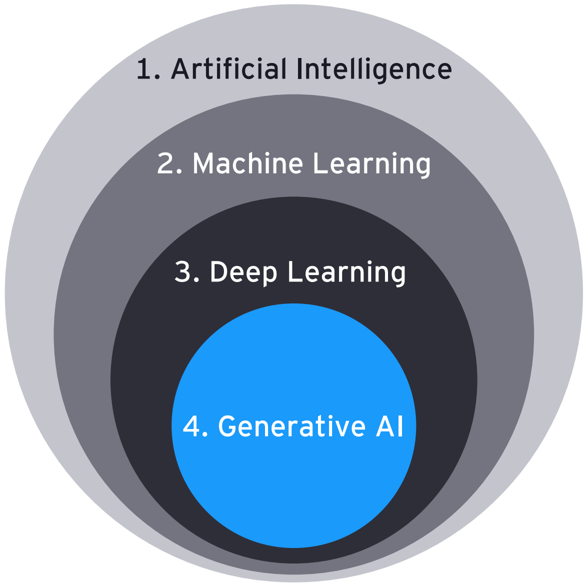

1. **AI:** Simulation of human intelligence tasks
2. **ML:** Algorithms to detect patterns in large data
3. **DL:** Neural networks for in-depth data processing
4. **GenAI:** Content generation using learned patterns

Source: Epical; NVIDIA; ResearchGate; EY-Parthenon research and analysis

Workloads

GenAI workloads can be classified into two main categories:

Training

Learning a new capability from existing data

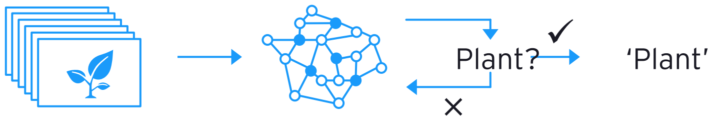

Inference

Applying a new capability to new data

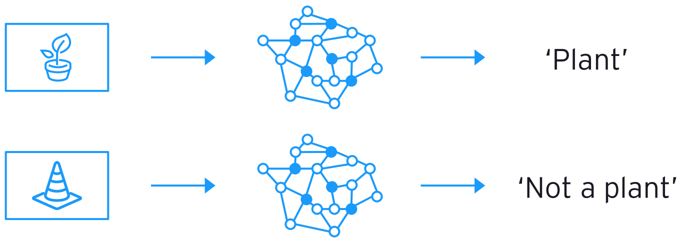

Fit into existing infrastructure

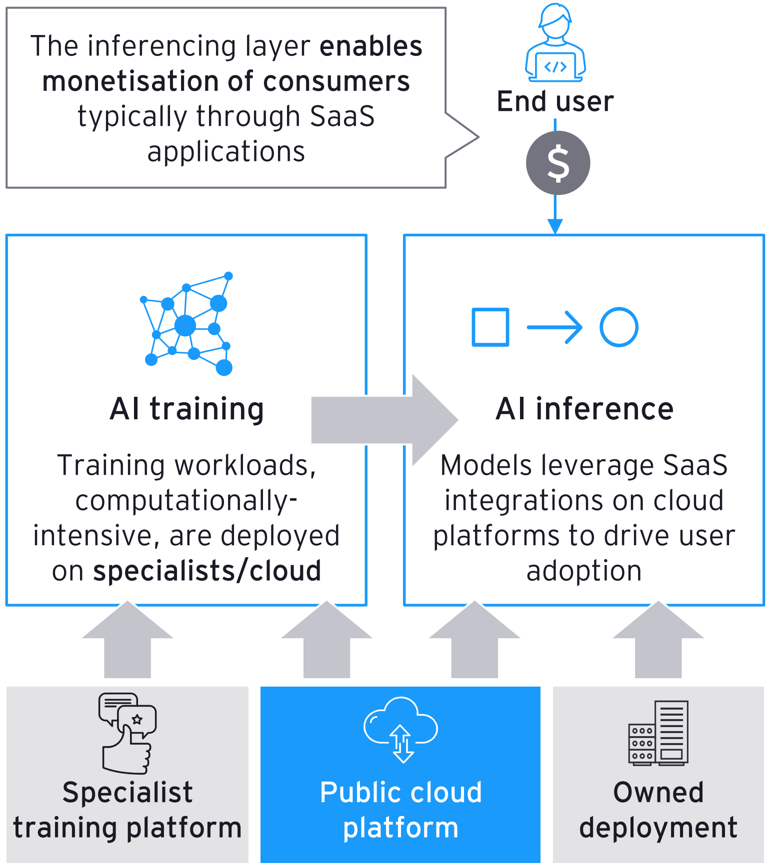

EY Parthenon

3

Hyperscale profit incentive

EY Parthenon

Public cloud services are both the fastest-growing revenue segment and the highest-margin business for cloud hyperscalers, driving strong shareholder value

2 **Hyperscale profit incentive:** Overview

Hyperscale profit incentive	**Public cloud services generate outstanding economics for hyperscalers**, who earn high margins from a sticky end customer base, supporting robust share price performance and incentivising continued capacity expansion. Outsourced DC rent remains a modest cost for hyperscalers, driving organic rent growth for data centre operators.

A Public cloud profitability

Hyperscaler equity stories depend on public cloud growth

- Cloud is a high revenue growth, high-margin business (**15-65%**)
- Cloud hyperscaler **share price growth** is predominantly attributable to their **growing cloud businesses** (>20% CAGR)
- Hyperscalers therefore directed budgets towards **cloud capex**, e.g. server equipment investment to maintain and accelerate growth

A Low lease cost sensitivity

Lease costs represent a small share of hyperscaler data centre costs (~5-10%)

- **Outsourcing** facilities **reduces** hyperscaler upfront **CapEx** burden and **provides** operational **flexibility**
- The **speed-to-market** compensates for minor differences between insourcing/outsourcing TCO
- Outsourcing propensity is expected to increase despite historical regional variances, e.g. **Europe has historically seen more outsourcing** than North America

A Operator pricing

Tight market conditions are expected to drive organic rent growth

- High-demand, low-supply markets command **premium rates due to location specificity** and associated supply constraints
- Exceptional **logistics rent growth** from supply and demand imbalances **hints at hyperscale public cloud data centres' future**

Source: EY-Parthenon research and analysis

Cloud products are driving profitability for hyperscalers, incentivizing them to continue investing in data centre infrastructure

2 A **Public cloud profitability:** Revenue, operating margin, Capex, and FCF

Profitability metric		Perspectives
	Hyperscaler cloud revenue	▪ Hyperscaler cloud revenue are significant and growing, accounting for c.$40b-$100b of revenue for each hyperscaler, with AI products expected to keep driving revenue up ▪ Hyperscaler cloud revenue has grown at c.20%-40% annual rate over the past half decade
	Cloud operating margin	▪ Cloud products have high operating margins, c.35%-50% on average across key players ▪ Cloud platforms are generally comparably more profitable than the business as a whole
	Hyperscaler capex	▪ Hyperscalers have spent a cumulative c.$190b in FY 2024 on capex, c. $90b of which have gone towards servers for their data centre facilities ▪ Hyperscalers have signalled they plan on continuing to accelerate spending on data centre and cloud infrastructure
	Hyperscaler free cash flow	▪ Free cash flow has grown to c. $200b across hyperscalers, growing at 10%-20% growth rate over the past 5 years ▪ Cloud margins are recognised as driving FCF and having a positive EPS impact

Cloud products are driving profitability and growth for CSPs across key metrics, eventually driving a positive EPS and share price impact

Source: Annual reports; Analyst reports; Earnings calls; S&P Cap IQ; EY-Parthenon research and analysis

Outsourcing primarily provides hyperscalers with faster speed-to-market, higher return on assets, and greater operational flexibility

2 B **Low lease cost sensitivity:** Data centre outsourcing benefits

Outsourcing is perceived by hyperscalers as having 3 main benefits

Return on assets

- Owning data centres can act as a drag on balance sheets
  - Data centre operations have relatively low-return vs. public cloud services
  - Operating data centres is not the core business model of hyperscalers, outsourcing it frees up capital to utilise elsewhere
- Outsourcing mitigates the risk on asset depreciation as, over time, the hardware will require continual upgrades

Speed-to-market

- Third-party operators possess the expertise and resources to deploy facilities rapidly, **with key benefits including:**
  - Faster land & power access
  - Established local relationships
  - Knowledge of local regulations
  - Availability of skilled labour

> **Time to market is a key consideration for hyperscalers. Often times, it's about beating the competition to the market.**
>
> Senior Director, Broker

Operational flexibility

- Outsourcing provides hyperscalers with significant operational flexibility
  - Hyperscalers can begin with an "on-ramp" to test demand and establish a foundational infrastructure
  - Once demand justifies it, hyperscalers can then scale to a full Availability Zone (AZ)

> **Hyperscalers have increasingly taken a modular approach to data centre development, which aligns well with the outsourcing model.**
>
> Director, Data Centre Operator

Source: EY-Parthenon interview programme; EY-Parthenon research and analysis

Hyperscalers outsource most of their capacity outside of North America; outsourcing is expected to grow over time due to demand and constraints

2 B **Low lease cost sensitivity:** Outsourcing geographical variations

Region	State of market	Market perspectives
North America	Hyperscalers have **historically relied less on outsourcing** given the relative homogeneity of the market	"I think less than 40% of hyperscaler facilities in North America have been outsourced." — Senior DC Manager, Hyperscaler
Europe	**Step change** in outsourcing behaviour relative to North America. Outsourcing rates are 50-60%	"Europe experiences significantly higher outsourcing due to varying regulations in each country." — Senior Manager, Hyperscaler
LATAM	LATAM has a **balanced mix** of insourcing and outsourcing, with the latter largely driven by skilled labour shortages	"Overall, I think the split between insourcing and outsourcing in LATAM is roughly 50-50." — Senior Manager, Hyperscaler
APAC	Overall, **outsourcing is prominent** in APAC, driven by difficulties with land and power access	"I think the split of outsourcing is ~70% in Vietnam and Malaysia, but closer to 50% in Singapore and Japan." — Senior Manager, Hyperscaler
MENA	**Outsourcing is the go-to model** for hyperscalers, notably due to strict data sovereignty laws	"Hyperscalers have an implicit dependency on outsourcing in the UAE, KSA, and Qatar." — Data centre Manager, Hyperscaler

Expected future situation[1]

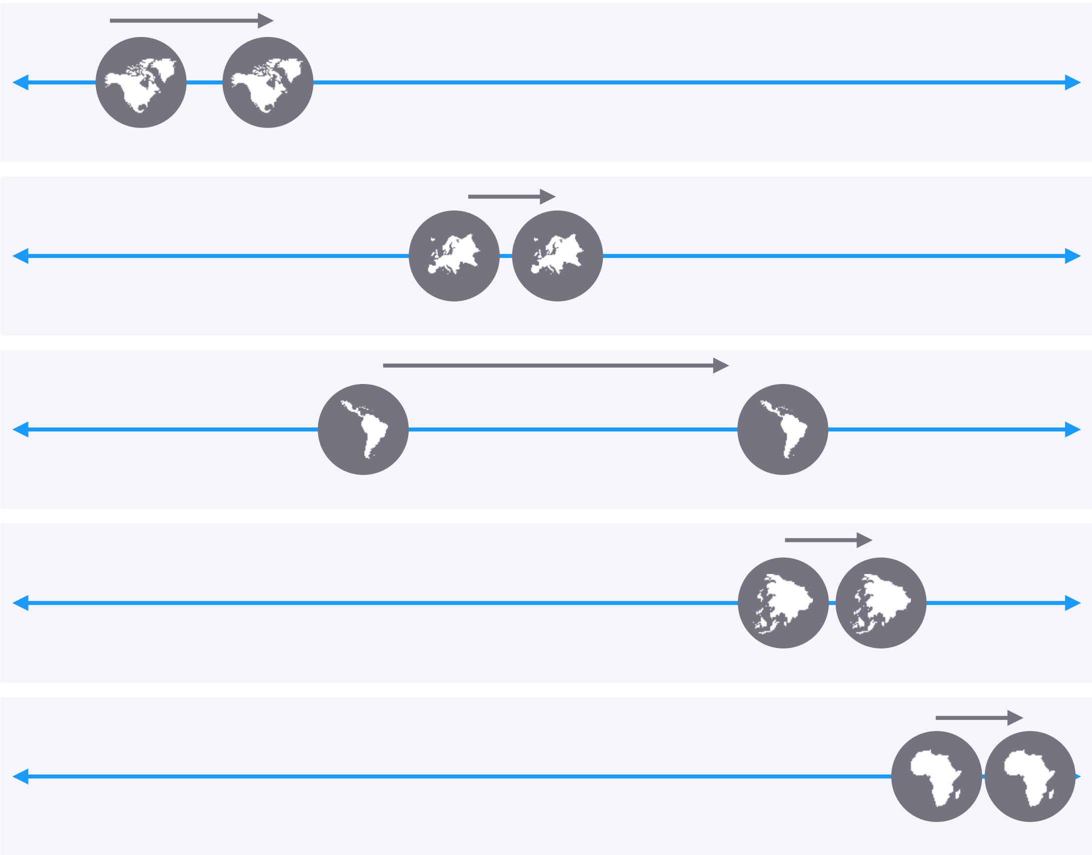

Less outsourcing 10-20% ⟷ More outsourcing 80-90%

[1] Next 10 years

Source: EY-Parthenon interview programme; EY-Parthenon research and analysis

EY Parthenon

Hyperscaler contract pricing is primarily determined by the rack space and power available, but also the additional services offered and connectivity options

2 C **Operator pricing:** Overview

Relevant DC types

- Scale/Hyperscale
- Interconnection-Focused
- Enterprise-Focused

Rack Space or Power

- Volume is measured by maximum power capacity
- Power costs are typically passed through to tenants
- Occupied space in a data centre is measured in racks
- Pricing unit: $/kW/month

Connectivity

- Connectivity is a significant factor when customers are choosing data centres, having crucial implications for improving latency
- Pricing unit: Connectivity tiers tied to data speed (Gbps or Tbps)
- Examples: cross connect, direct connect, carrier hotels

Other Services

- Data centres may provide additional services on top of colocation like security, network services (SD-WAN, DNS), and other managed data centre services to assist in deploying in the DC or in the cloud
- Typically charged per service

Final Pricing Bundle

- Some clients will only rent space/power, while others will buy a full bundling package
- The final pricing will depend on the combination of these components
- Typically, if a data centre offers bundles, the overall price is lower than if ordered piecemeal

Pricing Factors

Market Dependent → Dependent on Type and Speed of Connectivity → Dependent on Type and Extent of Service = Combination of Services (potentially with a bundling discount)

Pricing Analysis Focus

Source: Nuday Networks; Datacenterhawk; EY-Parthenon research and analysis

Recent organic rent growth has been driven by demand intensity, supply shortages, and tenant price flexibility, which are expected to persist

2 C **Operator pricing:** Organic rent growth drivers

Accelerated organic rent growth drivers

Growth driver	Demand intensity	Market supply shortage	Hyperscale tenant price flexibility
Description	In markets where international hyperscalers compete for capacity, **operators have greater bargaining power**	Key constraints include **power and land availability**	Hyperscalers are **willing to absorb higher rent** increases due to the importance of this high margin product and the operational complexities of migrating workloads

"Hyperscalers continue to deploy record levels of capital into infrastructure expansion, with AI workloads accelerating demand for high-density capacity. This is driving a structural uplift in leasing velocity and pre-commitment levels across Tier 1 markets.

European Broker

"In key European hubs, the availability of power and land is increasingly constrained, with moratoriums in place. These constraints are expected to persist, reinforcing pricing power for incumbent operators.

European Private Equity Principal

"Speed-to-market is essential for hyperscalers, for whom a two-year delay in revenue is more significant than a rent increase, since this constitutes a very small proportion of their total cost stack.

US-based Hyperscale Data Centre Operator

[1] A cloud service provider (CSP) is a third-party company that delivers on-demand, scalable computing resources – such as compute, storage, platforms, and applications – over a network
Source: EY-Parthenon research and analysis

EY Parthenon

Operators are securing price increases at contract renewals to mitigate rising costs and align with elevated market rates in the current supply-constrained environment

2 C **Operator pricing:** Renewal pricing case studies

Contract re-pricing bridge

Case study 1: Supply deficit Atlanta

Case study 2: Additional services Phoenix

Case study 3: Market rate matching Northern VA

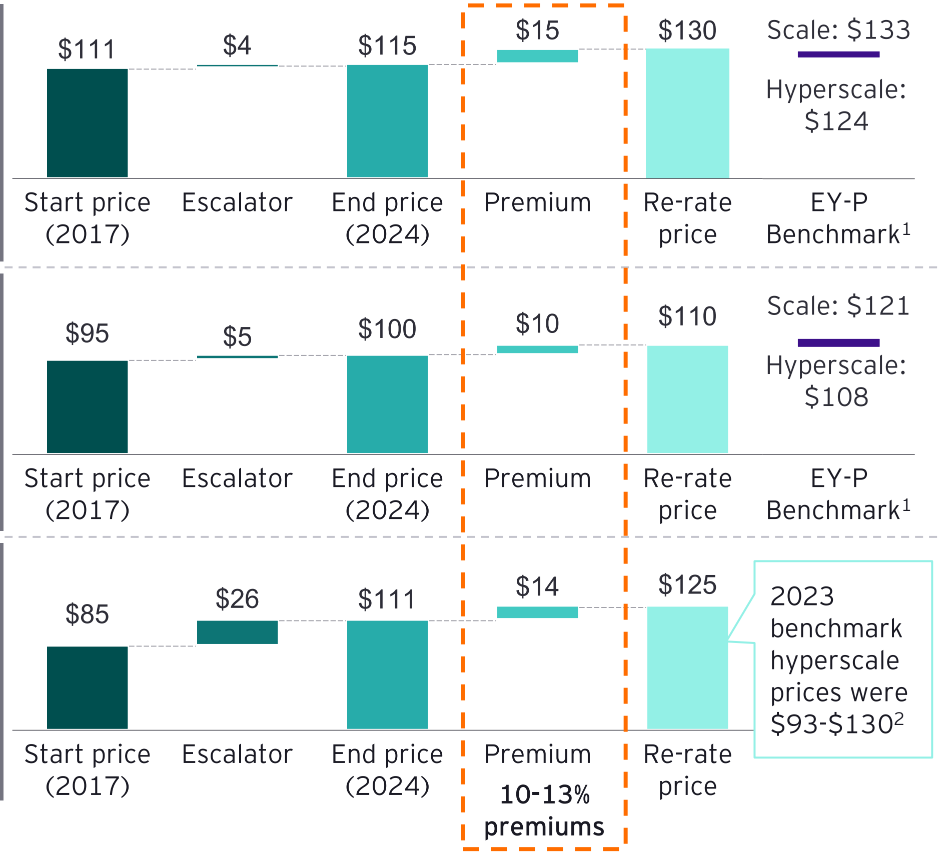

Context and market commentary

Case study 1

- Existing customer (4 MW) sought more capacity, faced 13% price increase (7 year contract, mid term) due to low-supply

> "We were able to **negotiate an uplift** in terms of how costs and interest rates have risen and there's very little supply in this market.
>
> Director Strategic Accounts, Operator

Case study 2

- Existing customer (9.5 MW) required onsite support in Phoenix, operator raises prices with bundled additional services

> "**I don't think any operators are looking to price gouge**, but they also have to account for cost. Some contracts factor in inflation.
>
> Director Strategic Accounts, Operator

Case study 3

- A hyperscaler renewed a **12 MW** deployment at a facility in Ashburn at the **$125 market rate[2]**, giving the operator a **13% uplift upon renewal**

> "We are willing to **tolerate higher prices** and more constrained capacity in **primary markets** like Northern Virginia.
>
> Hyperscaler

[1] EY-Parthenon hyperscale spot pricing benchmark in 2024, $/kW/month, [2] 2024 market prices have increased to $140-$170, largely due to market constraints

Source: EY-Parthenon interview programme; EY-Parthenon research and analysis

The infill data centre market exhibits similar characteristics to the logistics market; locations are constrained and competitive while being key to the business model

2 C **Operator pricing:** Logistics and data centres comparison

	Focus on demand hubs	Exchange with other businesses	Fixed geographic scopes
	 	 	 
Logistics	▪ Logistics hubs are developed in regions with **rising shipping volumes** and **predictable consumer behaviour** to optimise operational efficiency and profitability	▪ If a single logistics operator were to exit a hub, they would lose their **connections to other logistics operators**, reducing the overall value of their product	▪ Logistics hubs are extremely **location-constrained**, these must be within tight travel times to end-customers and other parts of the logistics infrastructure
Data centres	▪ Data centre operators prioritise locations with **growing enterprise and consumer data demand** to maximise long term ROI	▪ CSPs connect not only to their own networks and customers, but also **interchange with other CSPs** within hubs	▪ Similarly, data centre locations are dictated by **latency**, **redundancy** and **supply constraint** factors

We can expect infill growth in existing data centre regions to be located in close proximity to existing facilities, similarly to growth seen in the logistics market

Source: EY-Parthenon research and analysis

Exceptional logistics rent growth from supply and demand imbalances hints at data centres' future, where differentiating attributes indicate a positive spread vs. logistics

2 C **Operator pricing:** Commercial real estate and logistics rent escalation trends

Rent, by industry (Indexed 2013 = 100, 2013-2023)

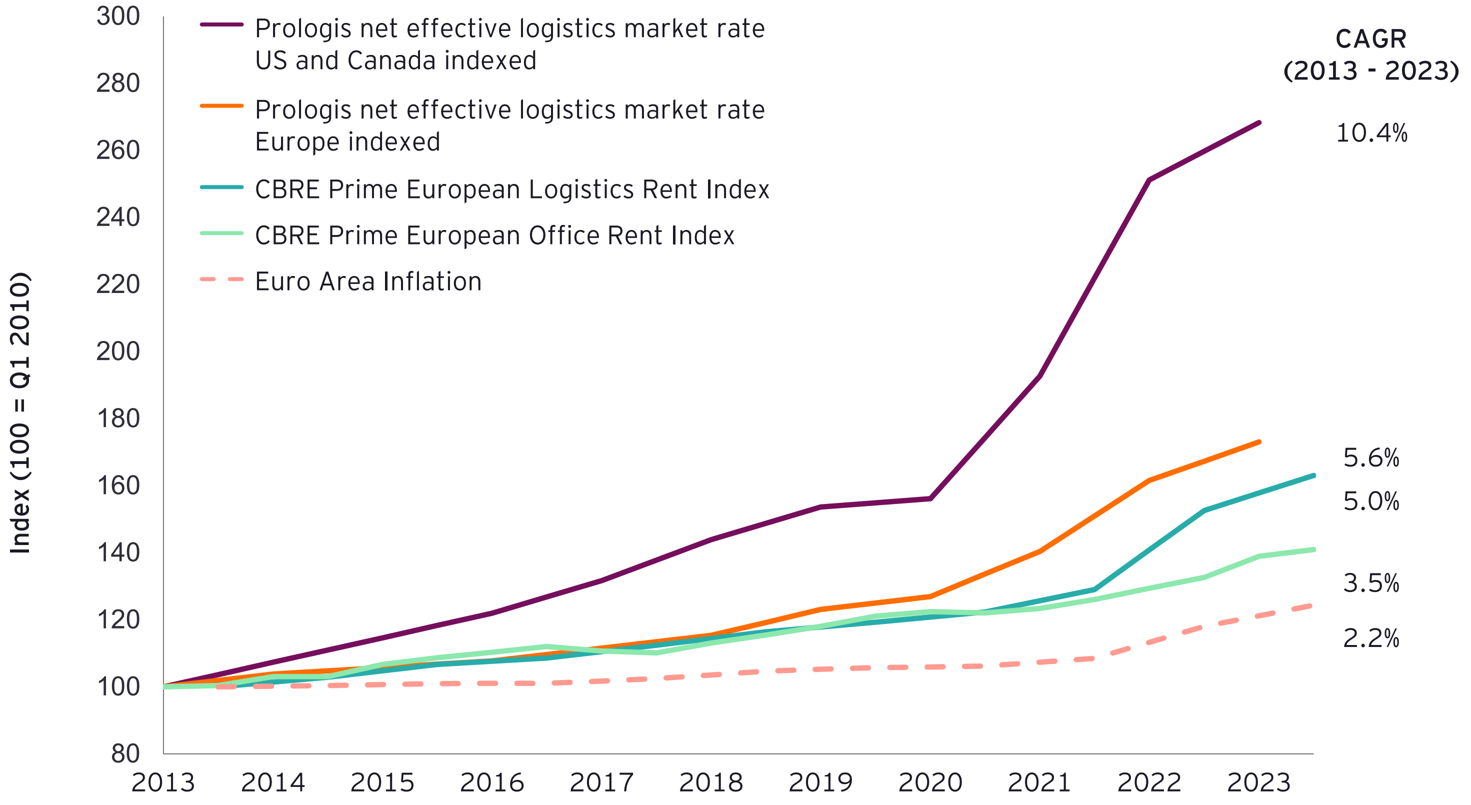

Source: CBRE; Prologis; EY-Parthenon research and analysis

Commentary

Logistics rent has seen significant growth, especially in the US and Canada

- Logistics rents have **consistently outperformed** office rents over the past 10 years due to supply and demand imbalances
- These rents were **robust** to the effects of Covid-19 in 2020-22, and have grown above inflation
- Data centres' hub-and-spoke network model is similar to logistics and means locations need to be **positioned in specific markets**
- Differentiating attributes, such as the higher costs of deployment, the higher costs of migration, and the hyper-local location requirements mean that **data centre returns are expected to be at a positive spread to those of logistics**

4

Hyperscale public cloud as fixed infrastructure

EY Parthenon

Public cloud architecture is fundamentally fixed infrastructure due to switching costs, regulatory standards, multi-billion dollar sunk investments, and latency requirements

3 Hyperscale public cloud as fixed infrastructure: Overview

Fixed infrastructure: Hyperscale public cloud data centre architecture is a fundamentally fixed infrastructure with significant switching costs and risks; DC locations are driven by latency needs (physically limited by the speed of light), competitive dynamics, and existing multi-billion-dollar networks and ecosystems which are intrinsically difficult to replicate.

Switching barriers

Switching facilities adds operational complexity and may delay capacity expansion for new demand.

New facility	Operational complexity	
$1.4b deployment	Direct cost $20-40m +	Material & uncertain impact !
~40MW facility		

Multi-billion-dollar sunk investments

Hyperscalers require significant capital ($6b+) for their distributed infrastructure footprint which they rarely move away from

- ~$5b — DC servers & facilities
- ~$1b — Fibre & interconnection
- +$?b — Power, transport, etc

Regulatory standards

Data localisation regulations restrict cross-border data transfers, driving location specificity of public cloud

Public service data must remain within the UK (aligned with EU GDPR)

Company financial records must be stored in Germany

Latency requirements

Locations are constrained by 1) technical intra-AZ latency, and 2) customer demand for low latency leading to cloud proximity competition

- <100 km — Optimal AZ-to-Enterprise distance
- 10-60km — Between AZs distance
- <10 km — Between DCs in an AZ requirement

Many cloud services rely on synchronous communication between data centers, which require short distances for ultra-low latency

Source: EY-Parthenon interview programme; EY-Parthenon research and analysis

Hyperscalers face direct and indirect costs which act as barriers against switching, many indirect costs are unquantifiable with potentially catastrophic consequences

3 A **Switching barriers:** Direct and indirect costs of switching

	Switching barriers		Financial impact
Direct costs	Planning and implementation	▪ Costs incurred in the planning and implementation of the migration including location identification, site testing and physical migration	**$20-40million** 1-2 months of duplicated operating costs, plus various once of planning and implementation costs
Direct costs	Parallel workloads	▪ Hyperscalers must temporarily maintain both legacy and new IT environments to ensure business continuity and minimise the risk of service disruptions	
Direct costs	Stranded IT costs	▪ If existing servers and equipment cannot be relocated/repurposed, the remaining book value of the IT assets would need to be written off ▪ However, customers usually time the migration with IT refreshes	***Sunk cost*** *– Noting that servers can be repurposed or sold*
Indirect costs	Unplanned downtime - outages	▪ There is potential downtime risks during the migration process, causing temporary service disruption to customers ▪ If severe, could result in customer churn	**Material and uncertain financial impact**
Indirect costs	Penalties	▪ Penalties can emerge if migration breaches the contractual SLAs between the hyperscaler and customer	
Indirect costs	Customer perception damage	▪ Reduced trust or satisfaction can result in damaged reputation in the industry, impacting future customer wins	Public outages can cost tenants millions in a matter of hours or days (e.g., major airline reported to incur £80m loss during a three-day outage)

Source: EY-Parthenon research and analysis

EY Parthenon

End-to-end switching could take up to ~3 years, introducing various risks and creating an administrative burden for hyperscalers

3 A **Switching barriers:** Planning and implementation timeline

Step	Identify new location	Assess state of existing infrastructure	Procure connectivity[1] for new site	Install in new facility	Migrate relevant equipment	Test post-Migration
Time	2-5 months	1-3 months	12-24 months		1-3 months	1-2 months
Description	▪ Hyperscalers need to be within AZs (~60km radius) ▪ Other customers may have other location requirements	▪ Assess what hardware to physically migrate and what to replace ▪ Identify existing applications and their connections outside of the DC	▪ Assess previous connectivity links and determine how these can be replicated ▪ Design and implement cable routing and management	▪ Install new servers and connectivity to the extent possible in new facility	▪ Hardware needs to be moved very carefully	▪ Systems and applications need to be tested once installed ▪ Customer needs to run the new site in parallel with while testing
Actions	▪ Site surveying and assessment ▪ Contract negotiation with new operator	▪ IT department project hours	▪ Significant planning by IT department ▪ Negotiation with carriers	▪ Labour	▪ Physical labour ▪ Transport equipment	▪ IT project hours
Risks	▪ Regulatory constraints ▪ Limited availability of suitable sites	▪ Legacy systems hard or costly to migrate ▪ Overlooked dependencies	▪ Long lead times for connectivity ▪ Network performance risks	▪ Compatibility or configuration errors ▪ Installation errors	▪ Hardware damage ▪ Downtime or data loss	▪ Configuration errors ▪ System failures ▪ Inadequate monitoring

[1] This assumes migration will require laying new fibre at the new site
Source: EY-Parthenon interview programme; EY-Parthenon research and analysis

EY Parthenon

Hyperscalers and their customers sign strict contracts; switching increases the risk of breaching these commitments, with potentially significant cost penalties

3 A **Switching barriers:** Breaching contracts

Contracts and SLAs

- Customer contracts are signed between hyperscalers and their customers and include Service Level Agreements (SLAs) that covers the key metrics the hyperscaler must maintain
- Hyperscalers sign wholesale contracts with operators when outsourcing capacity
- Any new facility must meet the same requirements the previous facility, with key feature across 4 categories:

Services	▪ Detailed outline of services provided ▪ Detailed exclusions where the contract doesn't apply
Performance metrics	Includes: ▪ Uptime ▪ Latency ▪ Throughput
Incident management	Includes ▪ Response time ▪ Disaster recovery
Security	Includes: ▪ Data security ▪ Physical security ▪ Compliance (e.g., GDPR, ISO 27001)

Switching risks

- Hyperscalers migrating DCs are at risk of breaching SLAs
- Potential risks include

Location-specific contracts	**Downtime** above that allowed in the SLA
Current inter-connections	**Connectivity** with other DCs

- Breaching these contracts or SLAs can lead to **significant financial penalties** for the hyperscaler and loss of customer trust, making them more reluctant to take the risk of switching

"In the event of an issue enterprises would penalise the cloud provider, and the hyperscalers would in turn penalise the co-location provider. **The SLA negotiations and agreements specify exactly who has to pay what.**

Hyperscaler

Sources: Industry publications; EY-Parthenon interview programme; EY-Parthenon research and analysis

EY Parthenon

Data sovereignty regulations restrict cross-border data transfers, requiring CSPs to deploy capacity across continents, regions, and countries

3 B **Regulatory standards:** Global data regulations

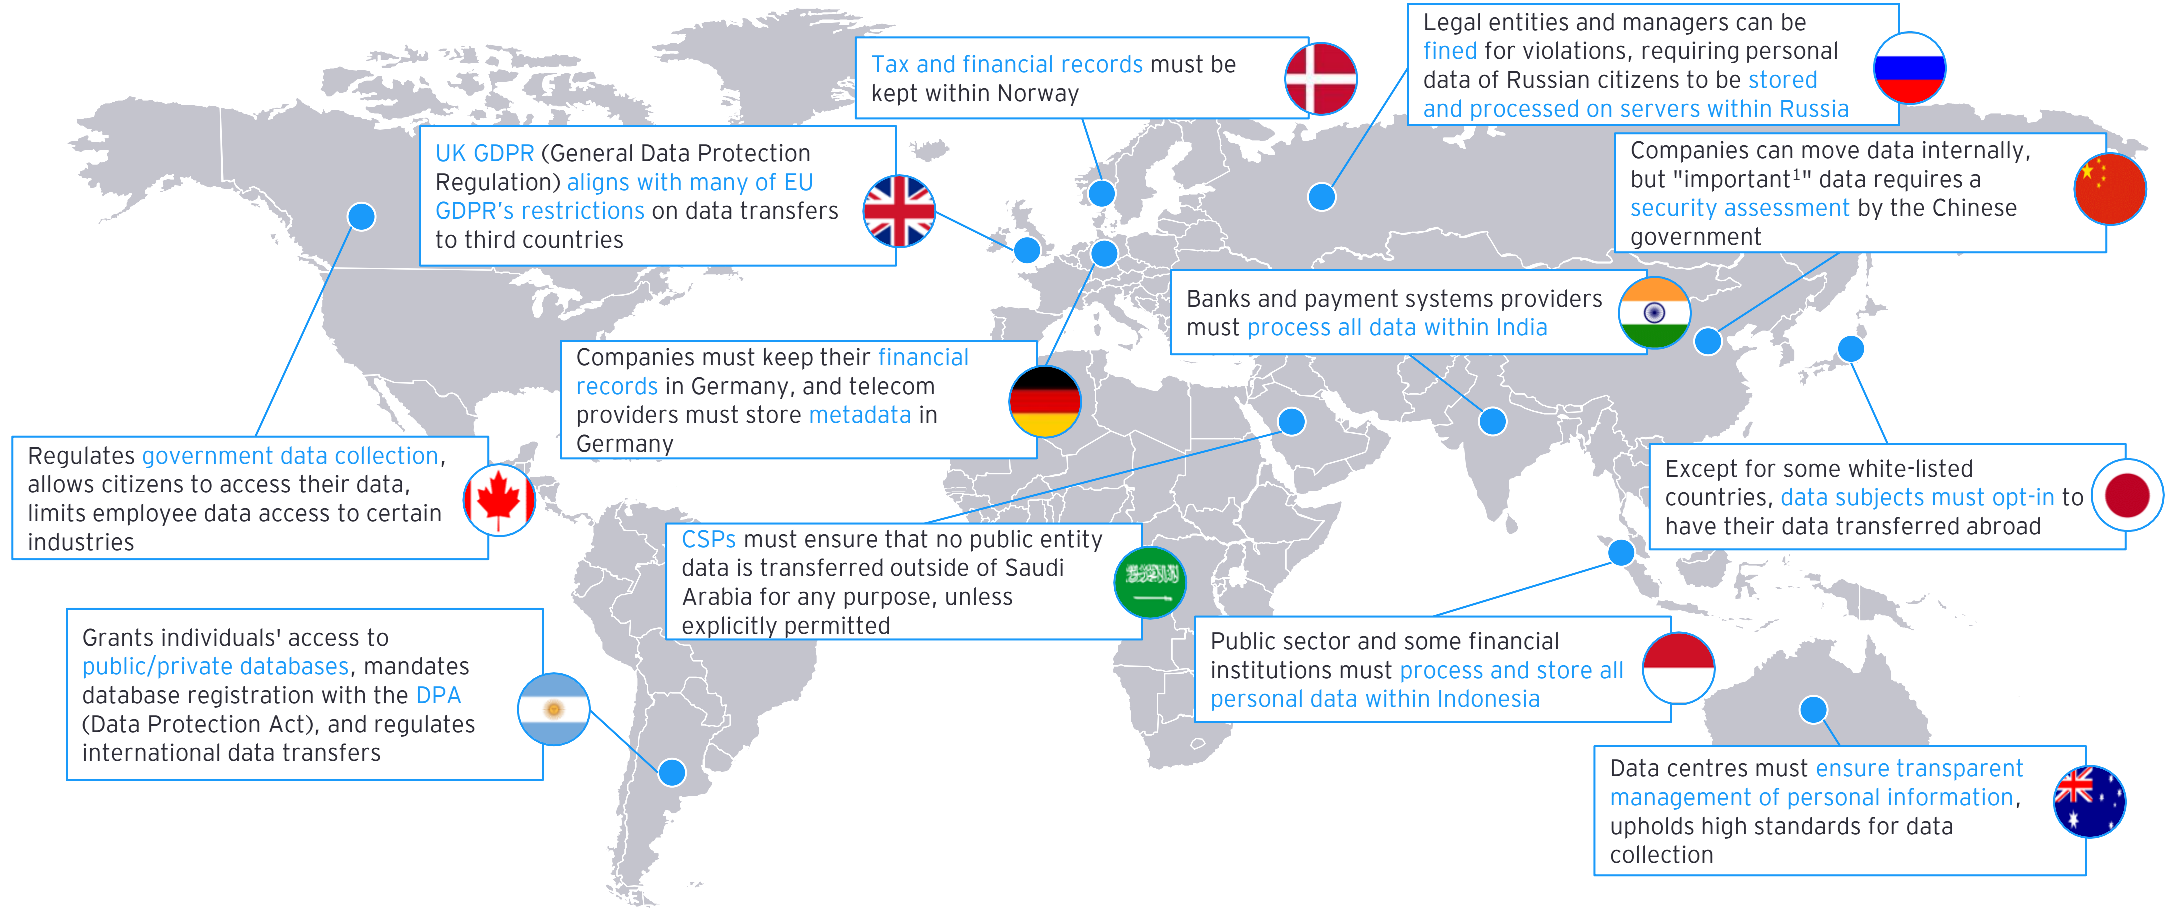

[1] Refers to information related to national security or that can identify Chinese citizens
Sources: Arizton Advisory & Intelligence; Baker McKenzie; EY-Parthenon research and analysis

Cloud availability zones require significant capital ($6b+) in addition to leveraging multi-billion-dollar existing power, transport, and other infrastructure

3 C **Multi-billion-dollar sunk investments:** Overview

	Data centre servers and facilities	Fibre and interconnection networks	Power transmission and transport infrastructure
Investment drivers	Servers ▪ Hardware (e.g., CPUs, GPUs) and software (e.g., operating system) ▪ **Estimated cost:** $10-30m/MW	Fibre networks ▪ Connecting data centres to each other and to end users ▪ Proximity to existing networks to reduce costs ▪ **Estimated cost:** $1-5m/MW	Power transmissions ▪ Need for proximity to existing grid, developed over decades with multi-billion sunk costs ▪ **Estimated cost:** N/A
Total sunk investment per AZ[1]	Facilities ▪ Building construction, mechanical systems, and electrical infrastructure ▪ **Estimated cost:** $10-15m/MW		Transport infrastructure ▪ Need for logistics & workforce mobility ▪ Public-sector spending to support broader accessibility ▪ **Estimated cost:** N/A
	~$5b	~$1b	+$?b

~$6b+ estimated direct sunk costs into a typical AZ + **multi-billion wider infrastructure** ecosystems

[1] A typical availability zone consists of three data centres of 40 MW each
Sources: Thunder Said Energy; EY-Parthenon research and analysis

Availability zone infrastructure placements are an extension of the multi-decade multi-billion-dollar modern infrastructure ecosystem

3 C **Multi-billion-dollar sunk investments:** Advancement of infrastructure over time

Previous State → Current State

Infrastructure hubs	Strategic factors	Relevance for data centres	Examples
Cities	▪ Access to water and fertile land sustained large populations, fostering the development of cities that became hubs for agriculture and trade	▪ Located in urban areas to serve dense populations ▪ Ensure proximity to business hubs	▪ One Wilshire (Downtown LA) serves as an interconnection between North America and Asia
Roads & waterways	▪ Roads and waterways were built to connect cities, enabling trade, communication, transport and urban growth	▪ Locations close to major roads offer easy transportation access ▪ Proximity to water sources for cooling	▪ Frankfurt, once key trading hub along the Rhine and now linked to major highways, hosts one of Europe's largest DC hubs
Railroads	▪ Railroads quickly moved goods and people, connecting industries, ports, and resources, which increased trade, boosted the economy, and grew cities	▪ Close to transport networks offering power supplies and easy access	▪ Union station technology centre (Indiana), has been transformed from a train station to a data centre
Telegraph and telephone lines > fibre optic	▪ Copper and fibre cables enabled long-distance communication and created dense city hubs. Utilising railway land, they formed a dual network of transport and communication, supporting industry and trade	▪ Strategically located along communication routes ▪ Shared infrastructure and high-speed connectivity	▪ Ashburn (Virginia), previously a telecom corridor, is now one of the largest data hubs in the world
Power grids	▪ Power grids began in dense areas and spread to urban and rural regions, supplying electricity for industry, technology, and connections with other grids	▪ Locations near power sources ideal for continuous operations	▪ A Hyperscaler has located a data centre next to the Bonneville Power Administration grid
Data centres	▪ Data centres are located near cities to reduce latency, close to power grids for reliable energy supply, and adjacent to roads for improved accessibility	▪ Often clustered in data centre hubs to leverage existing infrastructure and access skilled talent	▪ Slough is one of Europe's largest data centre hubs due to its proximity to London and role as a logistics and industrial hub

Often, DC hubs emerge due to a combination of strategic factors such as proximity to population hubs and major business centres, availability of industrial land, reliable power infrastructure, and strong transport connectivity via road and rail networks

Source: EY-Parthenon research and analysis

EY Parthenon

CSPs select site locations to solve for high availability, low latency, fault preventing, and resilient load balancing within their distributed architectures

3 D **Latency requirements:** Availability zones (AZs)

Typical cloud region and availability zone architecture

Cloud region: 2+ AZs in a geographically distinct area

Availability zone (AZ): 1+ independent DC with separate power, networking and cooling

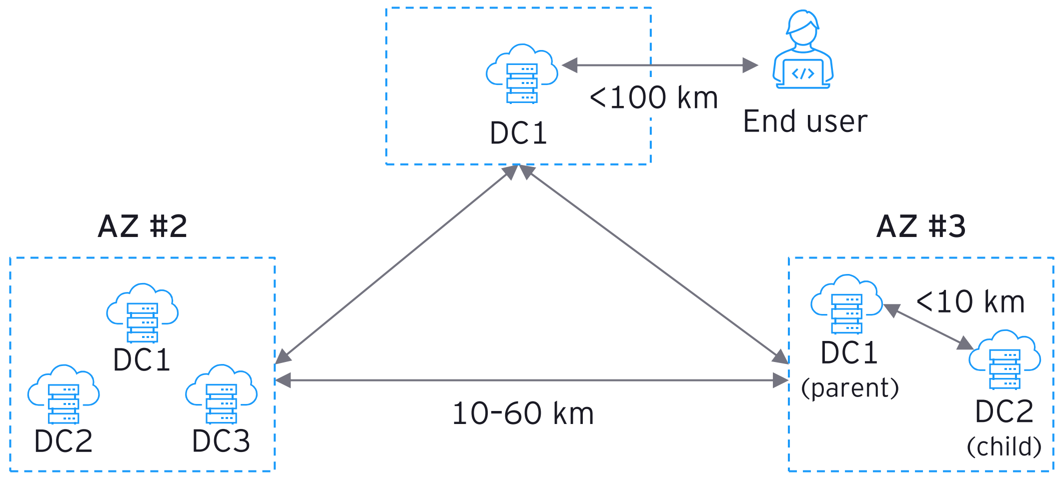

Availability zone design criteria

Distance to End User: Low latency

Close proximity to user populations through well-distributed regions/edge locations

Minimizes transmission delays, improves responsiveness, throughput, and reliability

Distance between AZs: Fault tolerance and disaster recovery

Multiple, physically separate data centers with proximate location for low-latency. 3 distinct fiber connections.

Mitigates single points of failure, enables fault tolerance and rapid disaster recovery

Distance across AZs in a region: Effective load balancing

Distribution of workloads across multiple AZs with real-time traffic routing and dynamic resource allocation

Balances system load, avoids congestion, and enhances overall performance and uptime

AZ locations should be near customers, near other facilities in the AZ, and in reach of the cloud region architecture, limiting the set of potential locations and driving stickiness

Source: Company websites; EY-Parthenon interview programme; EY-Parthenon research and analysis

Hyperscalers expand their presence to match increasing demand, entering new markets with on-ramp facilities before deploying local zones and then full cloud regions

3 D **Latency requirements:** Hyperscaler cloud region development pathway

Peering
(Typical size per hyperscaler: <1MW)

- Initial market entry with peering setup to monitor user demand
- Typically small deployments (under 1MW)

Example markets: Budapest, Rome

On-ramp
(Typical size per hyperscaler: 1-3MW)

- Growing demand prompts hyperscalers to deploy a cloud on-ramp
- Cloud providers manage user access to direct connectivity, driving further cloud service adoption

Example markets: Barcelona, Dusseldorf

Local zone
(Typical size per hyperscaler: 5-10MW)

- Hyperscalers launch edge/local zones in high-demand regions
- Local zones typically provide a focused set of services

Example markets: Prague, Lisbon

Cloud region
(Typical size per hyperscaler: 20-50MW)

- Infrastructure is designed based on network, power, and data centre capacity
- Initial deployments in Tier-2/3 markets are typically 4-8MW

Zurich and Madrid's initial deployments were limited but have subsequently expanded

Example markets: Zurich, Madrid

Key data centre hub
(Typical size per hyperscaler: 250+MW)

- Cloud regions provide full compute, storage, and networking services in high-demand areas
- Multiple availability zones (AZs) are integrated for redundancy

Example markets: Frankfurt, Northern Virginia, London

Source: EY-Parthenon research and analysis

EY Parthenon

If AI demand declines, data centres in AZs can be readily repurposed for alternative cloud workloads, while those outside AZs face greater transition challenges

3 D **Latency requirements:** AI workload future proofing

Key factors	Hyperscale data centre within AZ	Hyperscale data centre outside of AZ
Positioning	Data centres within existing availability zones, close to end customers, serving workloads with low latency requirements	Data centres outside of availability zones, serving low-cost workloads with less strict latency requirements
Workloads served	Cloud; AI Inference; AI training	Cloud; AI Inference; AI training
Investment rationale	Cloud providers make investments into DC infrastructure based on a **probability weighted pipeline of cloud service sales**. "AI" DCs in AZs can be repurposed for cloud, derisking the investment	Finance teams face risks with DCs outside AZs that serve **AI, a newer technology with less predictable sales services**. As AI develops, sales may stabilise and incentivise further investment
Future proofing	✓ Given their positioning, data centres in AZs are flexible to serve alternative workloads	? DCs outside AZs cannot be repurposed for public cloud as they will not meet the design considerations

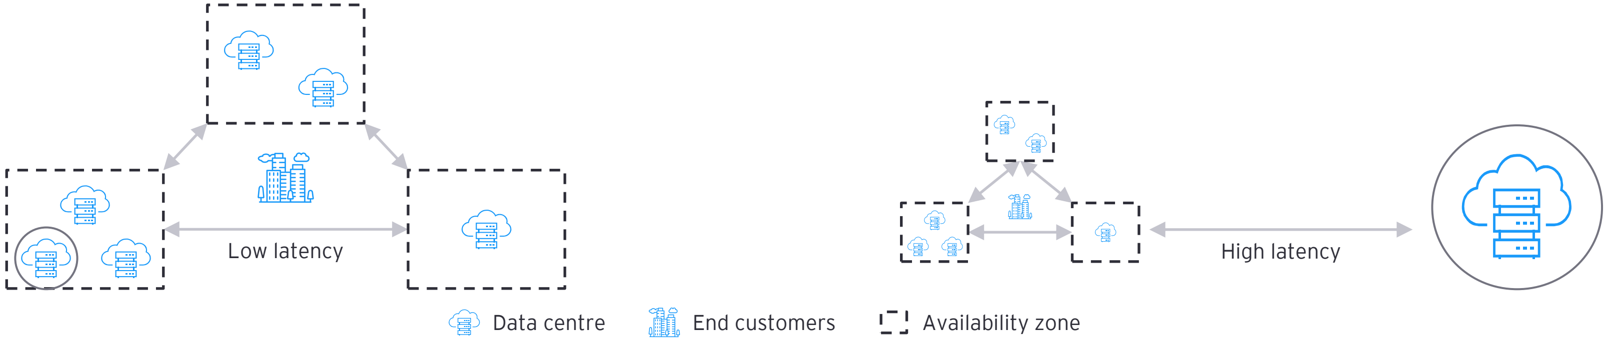

Source: EY-Parthenon research and analysis

EY Parthenon

5

Market conditions and supply constraints

EY Parthenon

Hyperscalers encounter significant challenges due to power, land, and permitting constraints when expanding in key hubs

4 **Market conditions and supply constraints:** Overview

Market conditions and supply constraints	**Data centres in key hub locations are positively differentiated because of power and land constraints**, as structural limitations on new power grid infrastructure and planning restrict additional supply in these high-barrier markets.

Power availability

Utilities are struggling to adjust to growing demand for power and **are expected to continue to lag data centre needs**

"In Northern Virginia, which is the largest cluster of data centres, **there is no power available**. All new power coming online in the next 3 years is already pre-leased… the grid company cannot keep up with the demand."

Hyperscaler

Northern Virgina faces constraints due to challenges around transmission infrastructure where the utility, Dominion Energy cannot guarantee power delivery

Land availability within acceptable zones

As demand grows, there is **increasingly limited plots of land available which meets the needs of hyperscalers**

"West London… has the highest number of individual data centres and the largest deployed capacity in Western Europe… there is **almost no near-term available land** to support further development, as it's completely planned."

Broker

Amsterdam has implemented a series of data centre moratoriums. As of April 2025, the municipality of Amsterdam has imposed a full ban on new data centres and expansions within city limits

Permitting processes

Some **authorities are introducing permits** or restrictions to make data centre construction more challenging

"If hyperscalers aim to build on a brownfield site in the UK, the planning permit authorities often react negatively once they realise it's for a hyperscale data centre. This reaction can make it **more difficult for hyperscalers to build**."

Operator

Dublin data centre development has faced permitting constraints from planning rejections, appeals backlogs, and environmental pressures

Supply constraints can create a seller's market, increasing pricing power → constraints are expected to be sustained into the future

Source: EY-Parthenon interview programme; EY-Parthenon research and analysis

Power constraints are generally becoming more difficult to navigate; sophisticated operators who can execute in constrained markets will gain competitive advantage

4 A **Power:** Constraints and implications

Power Constraint Driver	Description	Degree of Constraint: Short term	Degree of Constraint: Medium term	Degree of Constraint: Long term	Process and Current Constraint Outlook
Power generation	▪ As grid operators and utilities move away from carbon-generating power sources, they will need to be replaced with wind, solar, hydroelectric, geothermal, nuclear, batteries, etc.	L	L	M	▪ Power generation is generally not the key constraint in most markets; sufficient power to support new data centres in the near-to-medium term
Transmission capacity	▪ Transmission capacity is insufficient to meet rising power needs ▪ Grid operators tend to be monopolies, with limited incentives to accelerate deployment ▪ Transmission upgrades take 2–3 years, including planning and approvals	H	H	M	▪ Key constraints in expanding and extending transmission capacity from the power source to the DC, especially in Europe and U.S. due to permitting delays
Transformer/ substation equipment	▪ Lead times for substation equipment are increasingly extended, first due to Covid, now driven by surging demand from data centres and clean energy projects	H	M	L	▪ Lead times for substation equipment (e.g., transformers, switches) have grown over the last few years, and the supply chain remains constrained

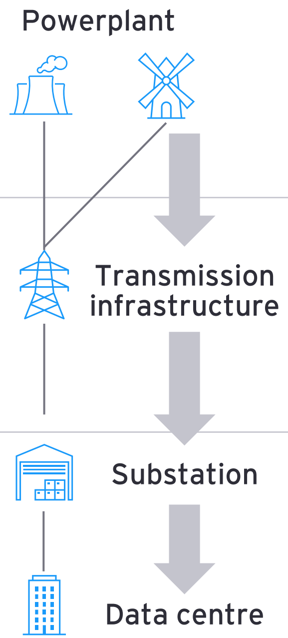

As power constraints create competitive barriers to entry, operators with a track-record of navigating power constraints in challenging markets are likely to accrue substantial benefits vs. less experienced/adept competitors

H High M Medium L Low

Source: EY-Parthenon research and analysis

Power constraints in Dublin have historically led to construction delays and investment in alternative power sources, increasing the value of existing capacity

4 A **Power:** Power constraints in Dublin case study

Market overview

Dublin is one of Europe's primary DC hubs – growth historically driven by its favourable regulatory[1] and tax environment

Dublin serves as a critical cloud region, with CSPs hosting their European HQ in the city

Hyperscalers' future ambitions are expected to grow in the market as demand grows

Current power constraints

Dublin is struggling to meet the increasing energy demand from DCs with their existing grid

"There is an evolving, significant **risk to electricity security** of supply in Ireland"

Commission for Regulation of Utilities

"The greater Dublin area is constrained and will remain so **for the foreseeable**."

EirGrid

Irish electricity consumption, split by consumption use case

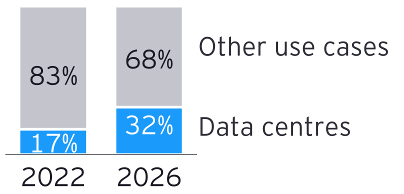

Regulatory changes

Dublin's de-facto moratorium

Introduced in 2022 with new DC assessment criteria:

- Facility location
- Onsite dispatchable generation
- Demand flexibility when requested

Limited new approvals were expected in Dublin until 2028, unless conditions change under the newly established government

Impact for operators

Case study 1: Operator investment in new technologies that avoid strain on local grid

Context	A Hyperscaler requested a 54.3km² expansion to a campus with two new DCs
Barrier	Finding suitable energy sources for permission
Outcome	Approval granted, conditional on construction of 170 MW gas-fired power plant, operating 8 hours daily, with build costs of €100m

Case study 2: Operators face delays in constructing new data centres

Context	A Hyperscaler requested for a 72.4 km² third phase of its Grange Castle DC campus
Barrier	Insufficient energy on the grid
Outcome	Dublin rejected the request due to "insufficient capacity" and "lack of detail of how it will impact power supply once operational"

[1] Relevant regulations: Commission for Regulation of Utilities (CRU/21/124), Government Statement on the role of data centres in Ireland's Enterprise Strategy

Source: IEA; Arizton; Data Centre Dynamics; Company websites; EY-Parthenon research and analysis

EY Parthenon

Power and land supply constraints in key US and EMEA hubs are expected to further drive rent growth and benefit experienced operators who can navigate these challenges

4 A B **Power | Land:** Case studies of constraints

US power constraints

Future nuclear reactors within key hubs

- US utility providers are **unable to keep up with demand for localised power** in key hubs, such as Northern Virginia
- Hyperscalers are partnering with future **Small Nuclear Reactor** providers to alleviate this constraint

Existing nuclear power plants

- Hyperscalers are also building data centres directly beside **existing nuclear reactors**, buying power capacity through Power Purchase Agreements
- These data centres are **far from major population centres**, so face high latencies

Amsterdam power & land constraints

Faced with grid congestion and a shortage of suitable land in Amsterdam, **regulators have placed restrictions on the construction of data centres:**

From 2019 to 2020, Amsterdam imposed a **moratorium on new data centres**

Since 2022, with the exception of certain rural regions, **new hyperscale data centres have been banned** by the government

Data centres in the Netherlands must have a **PUE below 1.3, or below 1.2 in Amsterdam**

Dutch multi-tenant total data floor surface area (1000m²)

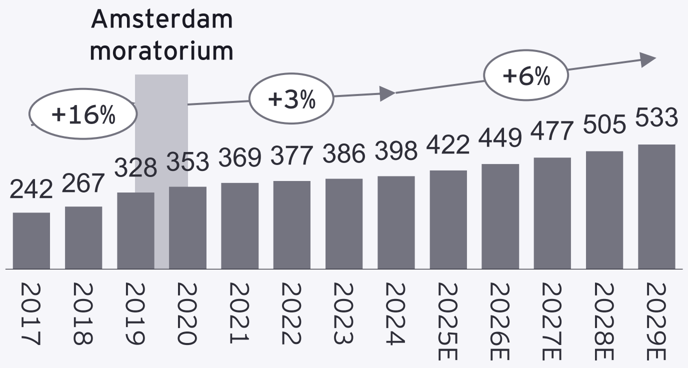

Frankfurt power & land constraints

Land suitable for data centres is scarce in Frankfurt

- **7 hectares of land** are needed each year for new data centres
- Local councillors are concerned about data centres **pricing out other industries**

Power is constrained by the local electrical grid

- Local utility companies are struggling to expand the grid to meet the rate of data centre demand

In response to the land and power constraints, data centres in Frankfurt:

Must be constructed or retrofitted to comply with **maximum PUE requirements**, of 1.2 for new builds or 1.3 for existing facilities by 2027

Are designed to minimise land usage by **building several stories high**, unlike those in other markets

Are being constructed **farther from the centre of the city**, in neighbouring localities such as Mainz and Hanau

Source: DC Byte; CBRE; Data Center Frontier; DDA; Netbeheer Nederland; Worldstream; Telegeography; Data Centre Dynamics; DLA Piper; Company websites; EY-Parthenon research and analysis

The interconnection process is lengthy and complex, creating substantial barriers to adding new connections to the grid

4 C **Permitting:** Interconnection process

Grid operators must review large projects before they are connected to the network to ensure that the project will not compromise the safety, reliability, and resiliency of the electricity system

Interconnection process timeline[1]

- ~6 months — **1** Interconnection request (queue window)
- ~6 months — **2a** Feasibility study
- ~6 months — **2b** System impact study
- ~6 months — **2c** Facility study
- ~1-2 years — **3a** Interconnection agreement; **3b** Design, procurement, construction; **3c** Operation and maintenance

Timeline varies by market. The more power constrained a market is, the longer it takes to connect to the grid

Project developer submits a detailed application/deposit and partners with local utility for interconnection studies

A Letter of Authority (LOA) permits installation of a cross connect in a data centre. It includes details about facility.

Interconnection studies

- **2a Feasibility:** Determines whether the project would cause electrical problems within the grid, and if transmission upgrades are needed
- **2b System impact:** Requires detailed information from project developer to assess grid impacts in more detail
- **2c Facilities:** Estimates cost of equipment, engineering, and construction of facilities needed to connect to the grid

Studies may need to be re-done if a generator higher in the queue decides to cancel their project/drop out of the queue, as the impact and upgrade needs for other projects may change. **This can cause delays in the review and approval process**

Studies culminate in an interconnection agreement between the utility and project developer followed by implementation of the project

- **3a Interconnection agreement:** Grants rights to interconnect and generate, defines project milestones, system upgrades and costs, and metering requirements
- **3b Design, procurement, construction:** Required construction on the distribution system, substation, or transmission system, if applicable (design and procurement occur before construction)
- **3c Operation:** Utility and developer are responsible for maintaining interconnection standards

Supply chain disruptions, permitting delays, and equipment shortages may cause additional delays after the interconnection agreement is reached

Contracts can be called power purchase agreements (PPAs), and typically include a customer disclosure statement, variable pricing, early termination provision (energy suppliers will seek to collect termination fees), requirements to notify of power need changes and general terms and conditions

Rule 9 agreements specify how connections are designed, constructed, inspected, paid for, owned and taxed

[1] Partially based on PJM's standard interconnection process, which is representative of other ISO/RTO processes
Source: PJM; LBNL; NV Energy; Data Centre Ops; Mission Critical Magazine; EY-Parthenon research and analysis

Appendix

Contacts

Dwayne Gibson-White

+44 7889 725 886

dwayne.gibson-white@parthenon.ey.com

Office: Ernst & Young LLP, London, United Kingdom

Rob Harris-Deans

+44 7555 143 874

rob.harris-deans1@parthenon.ey.com

Office: Ernst & Young LLP, London, United Kingdom

James Langlands

+161 7425 7375

james.langlands@parthenon.ey.com

Office: Ernst & Young LLP, United States

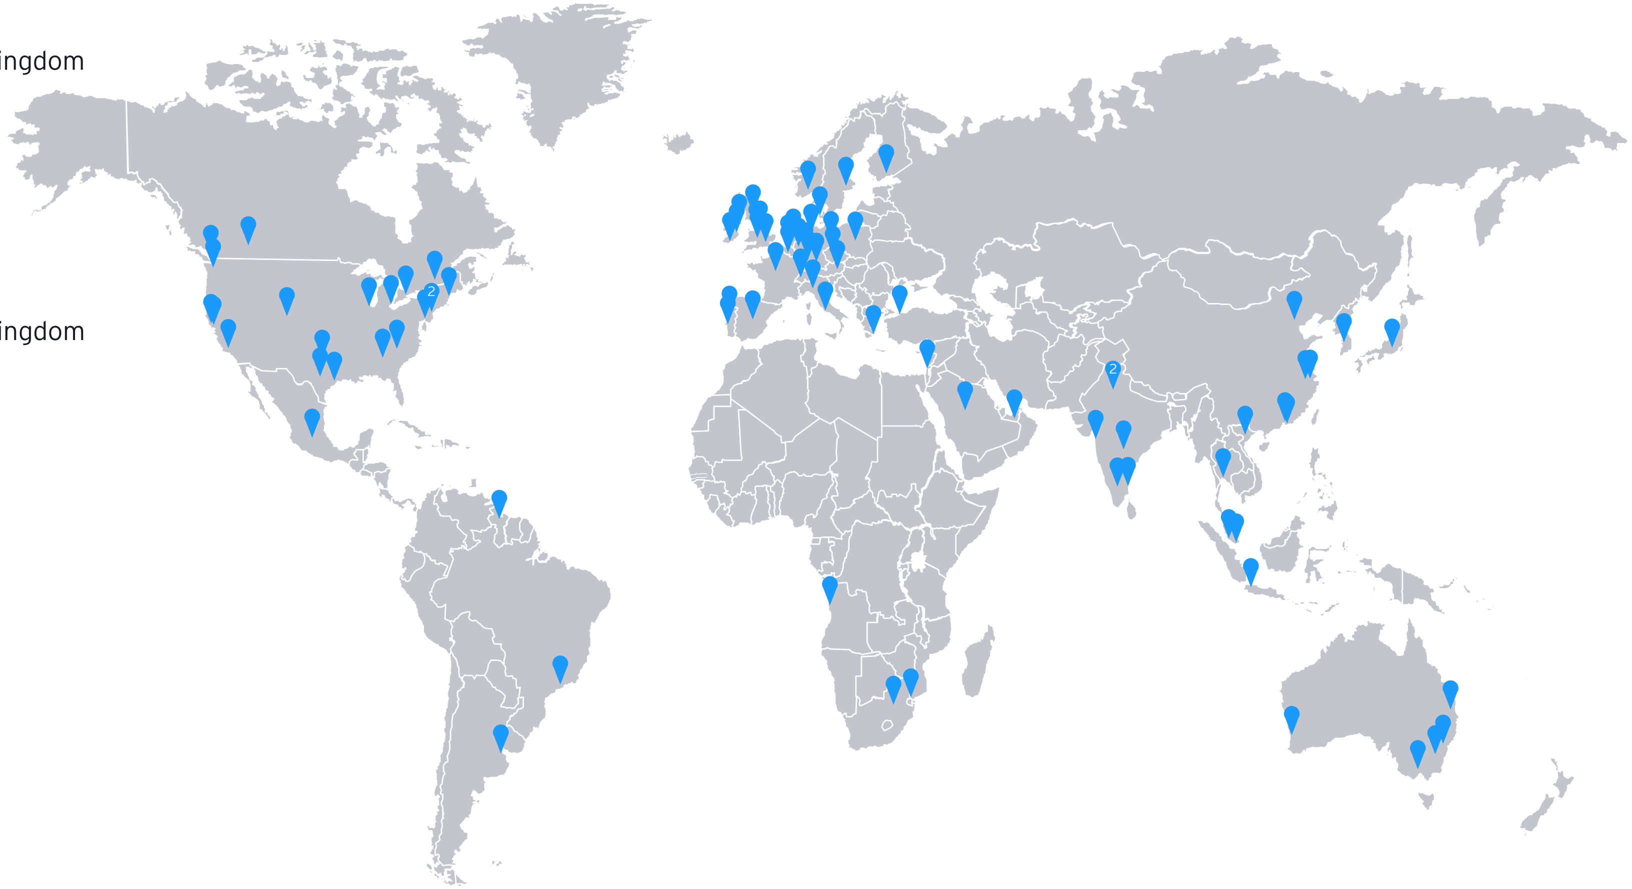

Glossary

List of abbreviations

Abbreviation	Terminology
AI	Artificial Intelligence
APAC	Asia-Pacific Region
AZ	Availability Zone
CAGR	Compound Annual Growth Rate
CDN	Content Delivery Networks
CEO	Chief Executive Officer
CSP	Cloud Service Provider
DC	Data Centre
EMEA	Europe, Middle East, and Africa
EPS	Earnings Per Share
ERP	Enterprise Resource Planning System
EU	European Union
FCF	Free Cash Flows
FY	Fiscal Year
GDP	Gross Domestic Product
GDPR	General Data Protection Regulation
GW	Gigawatt
IoT	Internet of Things
IP	Internet Protocol

Abbreviation	Terminology
ISO	International Organization for Standardization
IT	Information Technology
KSA	Kingdom of Saudi Arabia
LATAM	Latin America
LOA	Letter of Authority
LSEG	London Stock Exchange Group
MENA	Middle East and North Africa
ML	Machine Learning
MW	Megawatt
OECD	Organisation for Economic Co-operation and Development
P/E	Price-Earnings Ratio
PoC	Proof-of-Concept
PUE	Power Usage Effectiveness
ROI	Return on Investment
RTO	Recovery Time Objective
SaaS	Software as a Service
SLA	Service Level Agreement
TCO	Tota Cost of Ownership
YTD	Year to Date

EY Parthenon

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multidisciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

About EY-Parthenon
EY-Parthenon teams work with clients to navigate complexity by helping them to reimagine their ecosystems, reshape their portfolios and reinvent themselves for a better future. With global connectivity and scale, EY-Parthenon teams focus on Strategy Realized – helping CEOs design and deliver strategies to better manage challenges while maximizing opportunities as they look to transform their businesses. From idea to implementation, EY-Parthenon teams help organizations to build a better working world by fostering long-term value. EY-Parthenon is a brand under which a number of EY member firms across the globe provide strategy consulting services. For more information, please visit ey.com/parthenon.

UKC-040820 (UK) 02/26. Creative UK.
ED none
EYSCORE 007926-25-UK

ey.com