Data Centres

The Asset Class Powering the Digital Economy
Primer 1: An Introduction to Data Centres

July 2026

smartkarma

Garreth E. Elston
garreth@goldensection.capital

GOLDEN SECTION
REAL Asset Investing. Refined.

Executive Summary

Data centres have evolved from a niche property sub-sector into critical global infrastructure, yet most investment frameworks used to analyse them were designed for conventional real estate. That mismatch increasingly produces flawed conclusions. This series of two primers provides the conceptual foundation required to understand how the sector functions, what drives value creation, and why data centres should be analysed differently from traditional property assets.

At a physical level, a data centre is straightforward: a facility designed to house, power, cool, and connect computing infrastructure. Economically, however, the sector behaves less like conventional real estate and more like utility-linked digital infrastructure. Asset value is determined not primarily by the building itself, but by access to reliable power, network density, connectivity ecosystems, and the ability to deliver capacity at scale. A facility without secured power is not yet an operating asset, while a smaller, highly connected site in a constrained market may command materially higher valuations than a larger facility in an unconstrained but isolated location.

The sector is also far from homogeneous. Hyperscale, wholesale colocation, retail colocation, and edge facilities operate under different economic models, customer relationships, lease structures, and risks.

Revenue generation in the sector follows directly from this infrastructure-based model. Operators monetise colocation space, interconnection ecosystems, managed services, and power capacity (typically measured in kilowatts and megawatts rather than square metres). While revenues are generally contracted and supported by long lease durations with embedded escalations, the durability and quality of income depend heavily on tenant concentration, network density, operational reliability, and the operator's ability to re-lease capacity at attractive economics.

Capital intensity across the sector remains exceptionally high and continues to increase. Construction costs are rising globally due to power infrastructure requirements, supply chain constraints, and the transition toward AI-optimised facilities capable of supporting significantly higher rack densities. Advanced cooling systems, liquid cooling infrastructure, and enhanced power redundancy requirements are pushing development costs materially higher. In most major markets, power availability rather than customer demand has become the principal constraint on new supply.

Environmental regulation has also shifted from a disclosure exercise into a direct supply-side variable. Increasingly stringent efficiency standards, renewable energy requirements, and reporting obligations are shaping where and how new capacity can be developed. In constrained markets, these frameworks may strengthen the competitive positioning of incumbent operators, while simultaneously raising barriers to entry and extending development timelines.

Despite these constraints, the long-term demand outlook remains underpinned by powerful structural drivers. Cloud adoption, artificial intelligence workloads, enterprise digital transformation, sovereign digital infrastructure strategies, and rising data consumption continue to drive substantial global demand for computing capacity. The transition from AI model training toward inference workloads may further expand the geographic distribution of infrastructure requirements, as low-latency delivery increasingly requires capacity closer to end users. Asia Pacific remains structurally underpenetrated relative to its economic scale and digital intensity, supporting expectations for sustained long-term investment across the region.

Looking forward, the IEA projects that global data centre electricity consumption could reach approximately 945 TWh by 2030. IDC forecasts that Asia Pacific (excluding Japan) data centre capacity will reach 142,600 MW by 2029 at a 22% compound annual growth rate, and JLL estimates that the sector requires up to US$3 trillion of investment to meet that forecast demand.

This first primer aims to establish the analytical vocabulary required to evaluate the data centre sector rigorously. It provides the foundation necessary to understand how the asset class functions, where value accrues, and which factors are most likely to determine long-term investment performance. The second primer volume applies these concepts directly to the listed investment universe, with a particular focus on Singapore-listed REITs and the broader SGX data centre market.

Microsoft, Fairwater Class DC, Atlanta, Georgia, USA

Cover Image: Google Data Centre, 2 Jurong West, Singapore

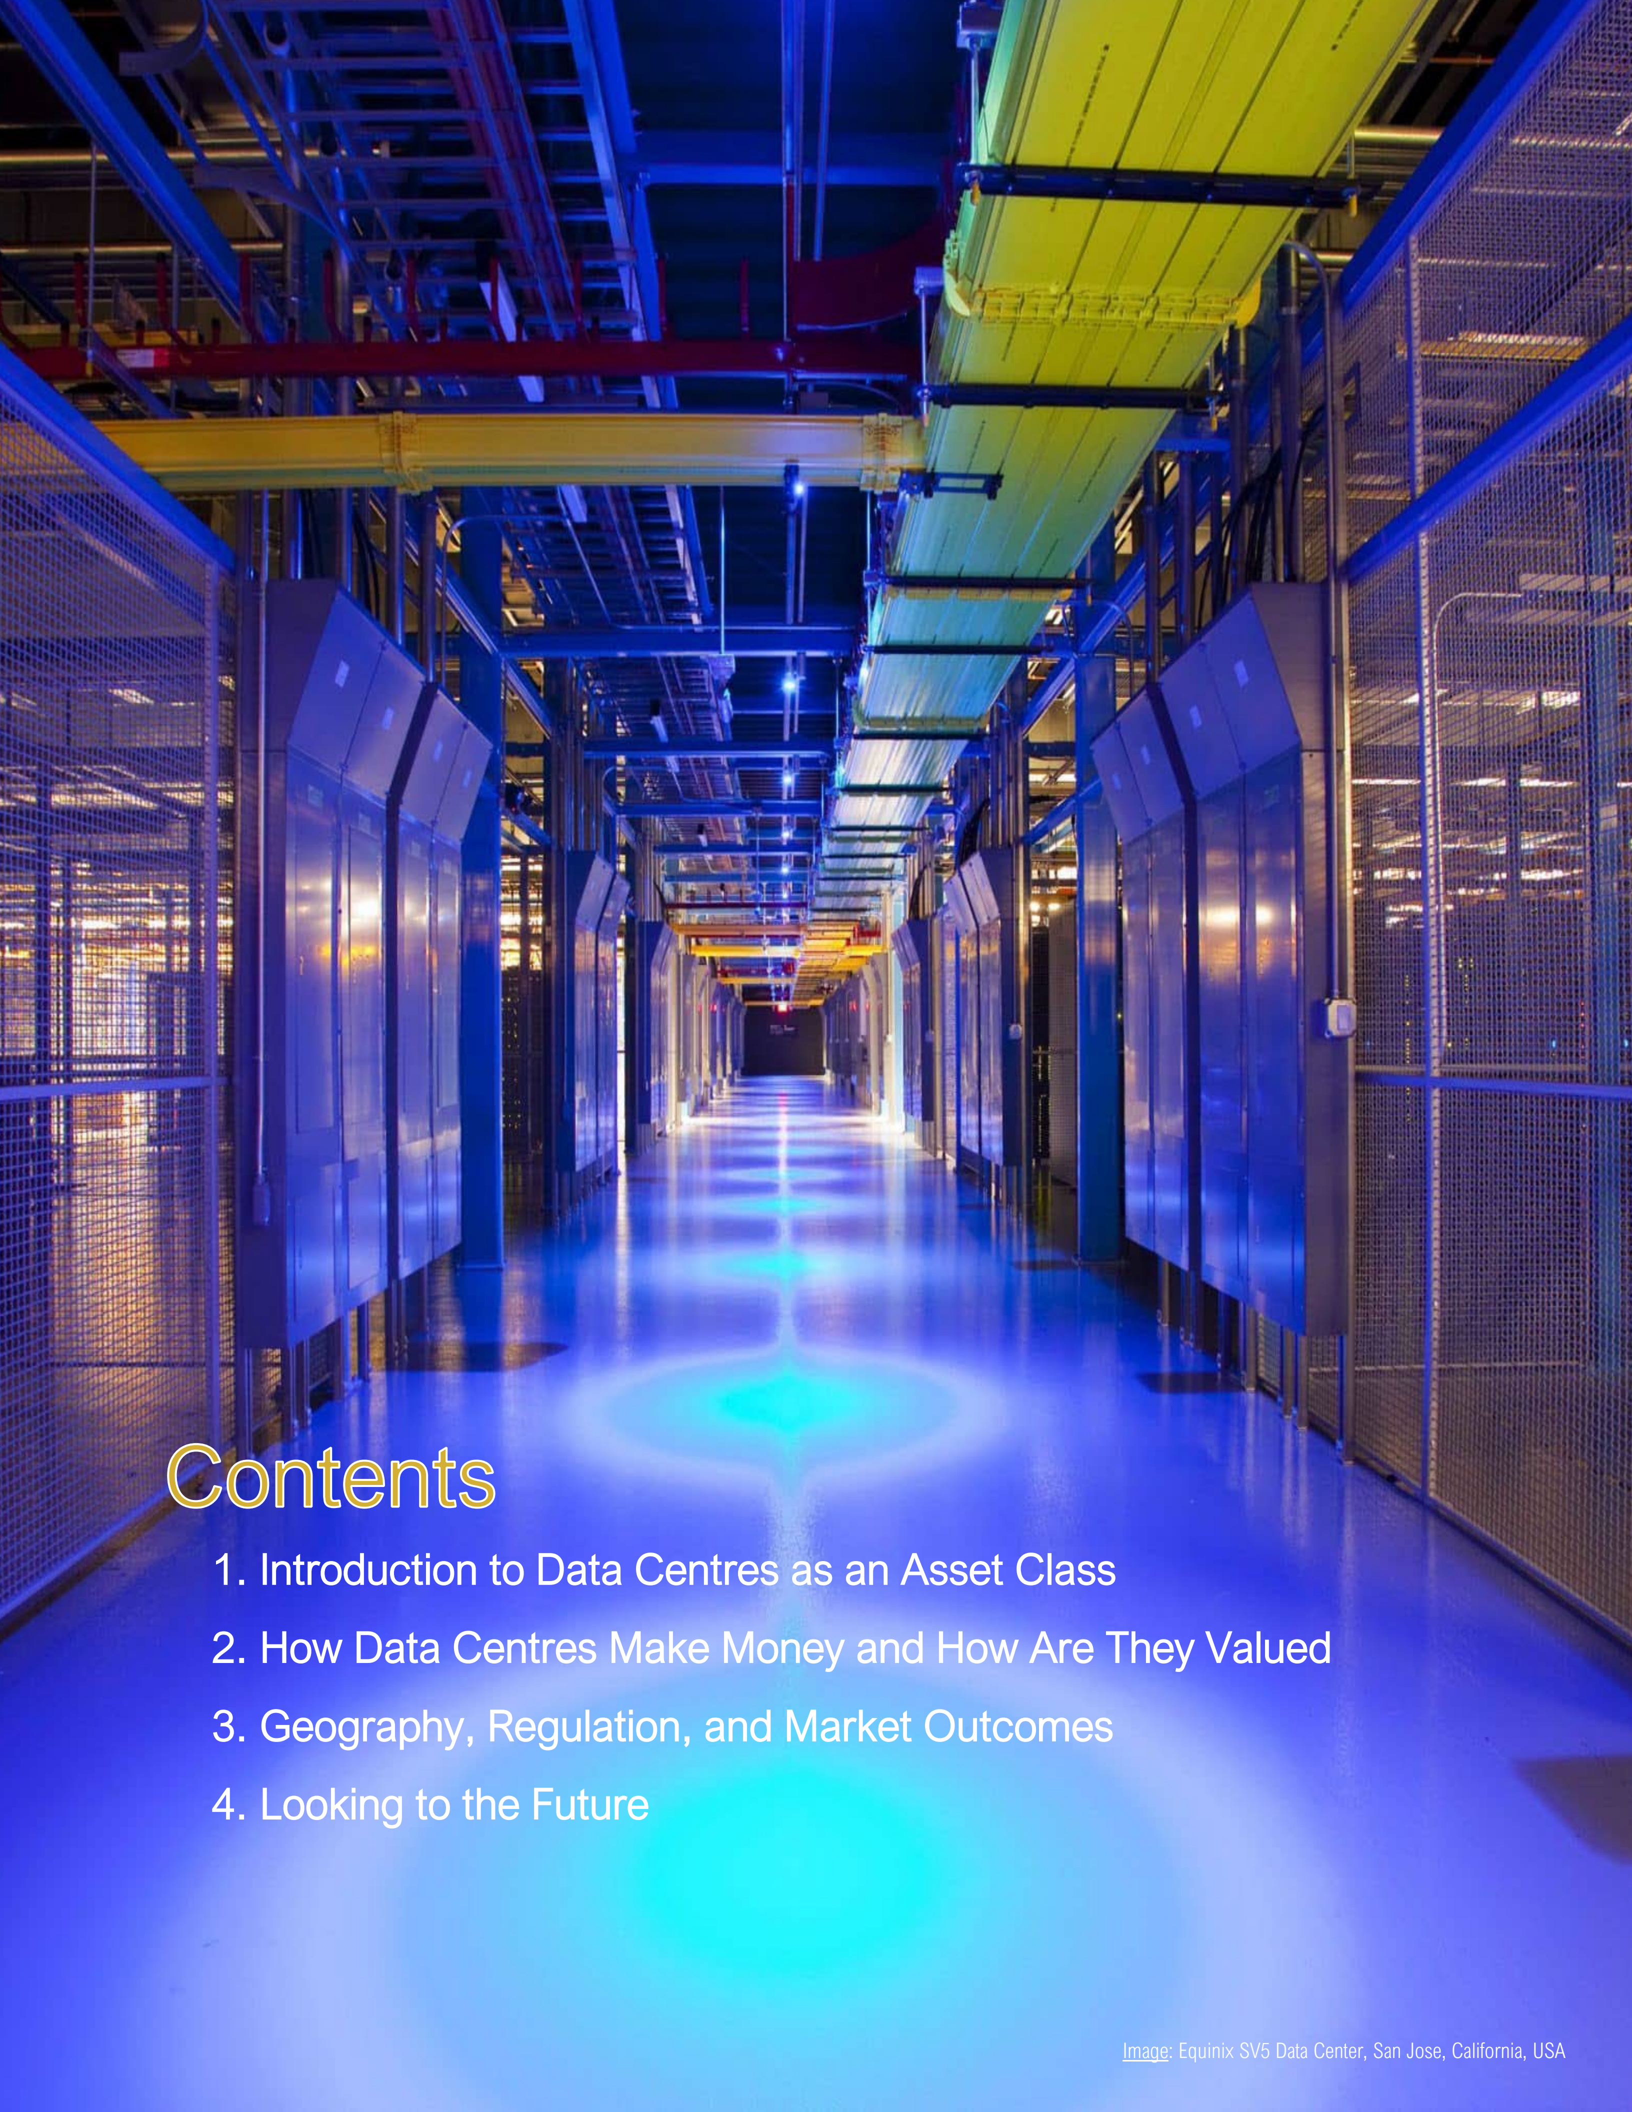

Contents

Image: Equinix SV5 Data Center, San Jose, California, USA

1. Introduction to Data Centres as an Asset Class

1.1 What is a Data Centre?

A data centre is a purpose-built facility that houses IT equipment (servers, storage, networking) and the mechanical/electrical systems that keep those workloads running reliably. For investors, the “product” is delivered power, uptime, and connectivity rather than floor area in the ordinary property sense. Data centre operators bill their tenants for price per kilowatt (kW) rather than the more traditional square metres (m^2) of traditional floor space used.

Table 1: Data Centre Types and Primary Attributes

Type	Size	Characteristics	Key Users	Key Benefits	Key Investor Aspects
Hyperscale Owner-built	Very large 100 MW+ 50,000 m²+	Built, funded, and occupied by the same company. No external tenant or lease. Entirely customised to the operator's workloads.	Amazon (AWS), Microsoft (Azure), Google, Meta.	Full control over design, power, and security. Optimised for proprietary cloud and artificial intelligence workloads. Operates as a long-term competitive moat.	Not directly investable as a standalone asset, exposure is through owning shares in the hyperscaler itself. There is no rent or income yield to value; the facility exists to defend and grow market share.
Hyperscale Leased / wholesale	Large 10-100 MW 10,000 - 50,000 m²	A specialist developer builds the facility and leases it to a hyperscaler or large artificial intelligence company. Contracts are measured in MW. The tenant typically dictates the technical specification.	Amazon, Microsoft, Google, large artificial intelligence companies, leasing from specialist operators such as Equinix, Digital Realty, and Keppel DC REIT.	Long lease terms (10-20 years) provide strong income visibility. High-quality tenants with investment-grade credit. Rents often linked to inflation.	Tenant concentration is the central risk, and losing one tenant can devastate income. More importantly, the largest tenants are actively building their own facilities, so lease expiry may mean the tenant leaves entirely rather than renegotiates (hyperscalers are both the customer and the competitor). Stress-test what the asset is worth vacant, not just on a renewal assumption.
Powered shell Build-to-suit	Large 20-100 MW 20,000 - 100,000 m²	The developer provides the building structure and power connection only. The tenant designs, installs, and pays for everything inside.	A single hyperscale tenant who requires a bespoke facility but prefers to control the interior specification and cost.	Lower developer capital outlay (no fit-out cost) Long lease term with a strong tenant covenant. Faster delivery than a fully fitted development.	When the tenant owns the fit-out, they effectively control the asset. If they leave, the building interior is bespoke and can have little value to an alternative occupier.

Hyperscale Leased, CloudHQ, Virginia, USA

Hyperscale Owned, Google, Singapore

Powered Shell, H5, Cleveland, USA

Type	Size	Characteristics	Key Users	Key Benefits	Key Investor Aspects
Retail colocation	Medium 1-10 MW 1,000 - 10,000 m²	A multi-tenant facility where businesses each rent a small amount of space, typically a rack, cage, or private suite. The facility also acts as a connectivity hub, allowing tenants to connect directly to other networks and cloud providers on-site.	Banks and financial services firms, media and content companies, network operators, mid-sized enterprises.	Revenue spread across many tenants reduces single-tenant risk. Direct network interconnection is high-margin and hard to replicate elsewhere. Network effects make well-established facilities sticky.	Headline occupancy is a lagging indicator; it is the last number to fall as the tenant base deteriorates. Many traditional enterprise tenants are gradually moving workloads to cloud providers. Analysis should look at net revenue retention within each tenant cohort by the year they joined, not total occupancy.
Edge / distributed	Small <1 MW <500 m²	Many small facilities located close to end users rather than centralised in one campus. Driven by the need to process data locally, either to reduce latency, meet legal requirements about where data must be stored, or reduce the cost of moving data between networks.	Telecommunications companies, Content Delivery Networks (CDNs) companies that deliver internet content quickly by storing copies close to users. Enterprises with real-time applications.	Serves latency-sensitive applications that a central data centre cannot. Supports legal data residency requirements across different jurisdictions. Growing demand driven by 5G networks and connected devices.	Revenue per site is modest and the economics of a micro site bear no resemblance to a large campus. The hidden risk is operational complexity (maintaining, staffing, and securing many geographically dispersed locations) generates cost that are routinely underestimated.
Enterprise On-premises	Small-medium <1-5 MW 500 - 5,000 m²	A private facility built and operated by a single company for its own internal use. Not a commercial transaction, funded through the company's internal technology budget and subject to its own capital expenditure review cycle.	Large corporations, government departments, regulated industries such as defence, healthcare, and financial services.	Complete control over data access and physical security. Meets stringent regulatory and data sovereignty requirements. Predictable internal cost base.	Not directly investable, exposure is through owning shares in the company that runs it. The key analytical error is comparing the cost of on-premises against cloud or colocation alternatives using only the upfront build cost. The true total cost of ownership (TCO) meaning the full lifetime cost including power, cooling, hardware replacement, and the staff required to run it.

Edge, LeadingEdge DC, Tamworth, Australia

Colocation, Equinix, Dallas, USA

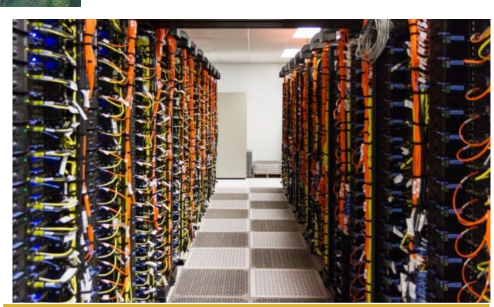

On-premises, NASA Goddard, Greenbelt, USA

Figure 1: Global Data Centre Locations

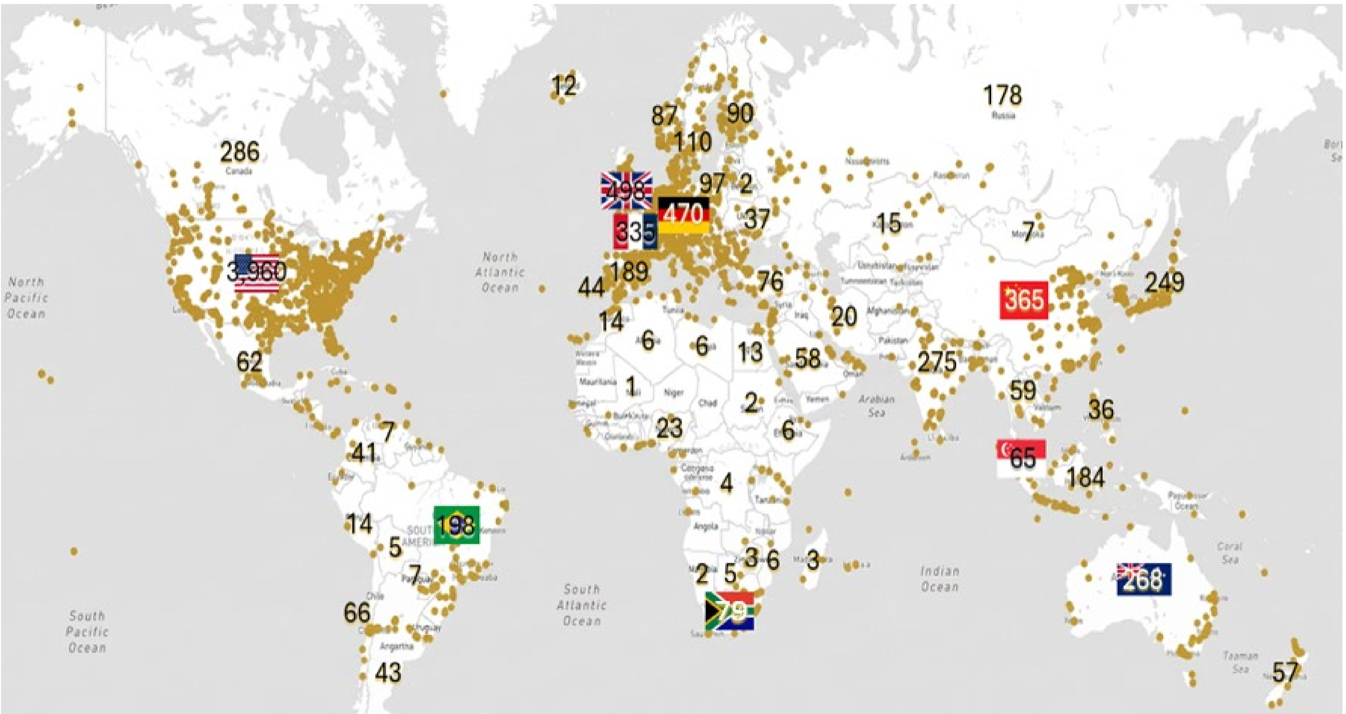

Source: Data Center Map, Golden Section Capital

Figure 2: Example of a Typical Data Centre

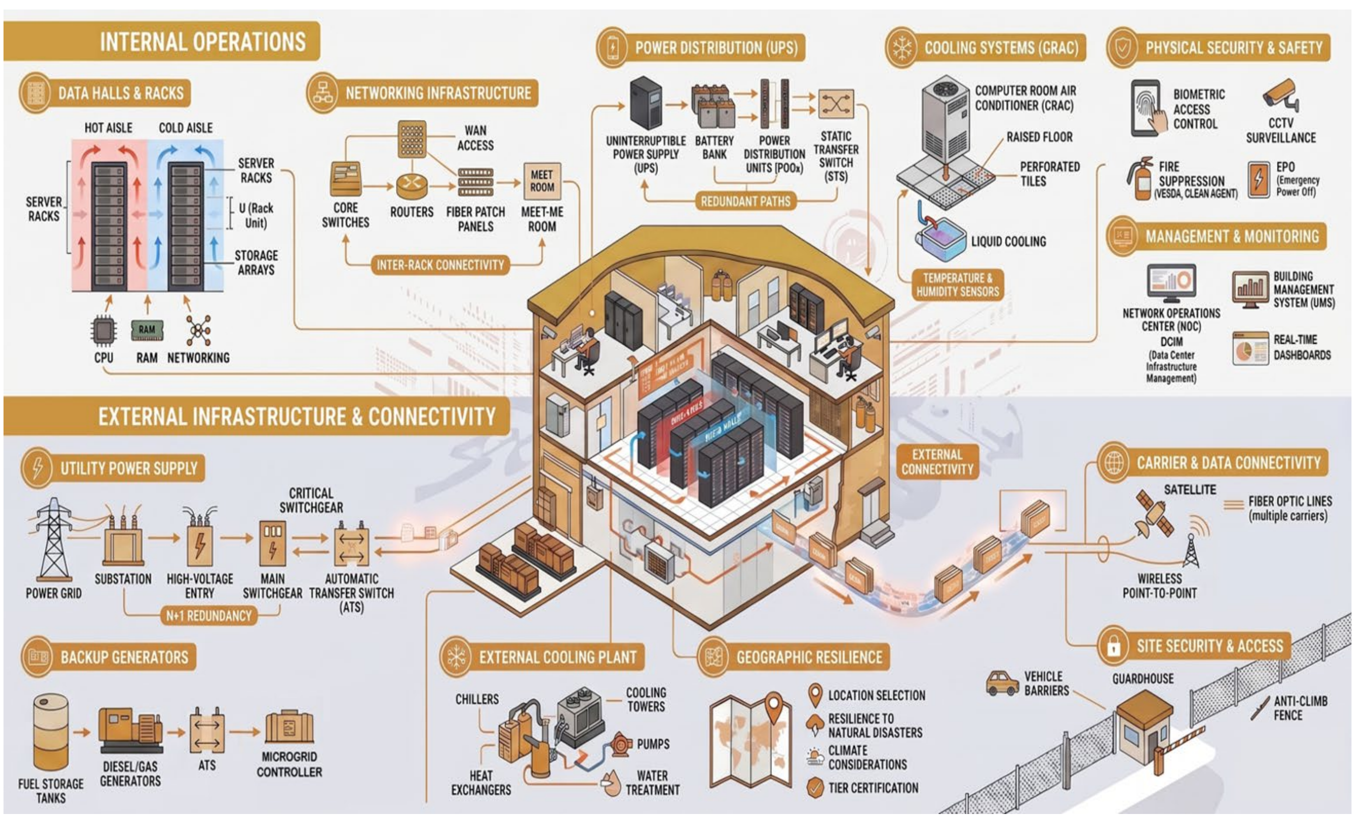

Source: Golden Section Capital

Figure 3: Primary Data Centre Type Usage for IT Workloads

The Uptime Institute's survey of 291 enterprises showed that the majority of enterprises still run the bulk of their IT workloads through their own data centres. The percentage has though steadily been dropping from 2020, with public cloud increasing and colocation remaining stable.

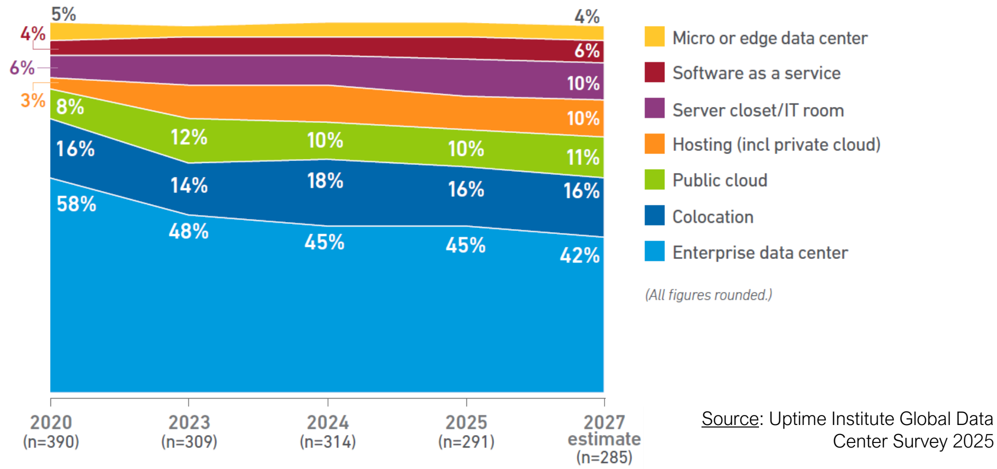

Source: Uptime Institute Global Data Center Survey 2025

1.2 How Data Centres Differ from Traditional Real Estate and Infrastructure

Data centres appear to match traditional property in legal form (leases, land, buildings) but behave like infrastructure in economic substance because power delivery, redundancy design, and uptime obligations dominate outcomes. The result is a hybrid risk profile: property valuation heuristics are effective only when power, capex, and operational performance are modelled explicitly.

Three structural differences matter most:

- **The binding constraint is often power, not space:** The IEA projects data centre electricity demand could reach ~945 TWh by 2030 in its Base Case, driven heavily by AI, and this scaling intersects directly with grid limitations. It is vital to rank secured power as not all secured power is equal, and there is a difference between: contracted vs delivered, firm vs interruptible, renewable vs grid mix. There is also a hierarchy for power from: Contracted, to Permitted, to Connected, to Energised.
- **Pricing is power-based:** CBRE's North America dataset reports asking rates in US$/kW-month and highlights faster growth in larger deployment tiers (3-10 MW) as AI and hyperscaler competition intensified.
- **Obsolescence is engineered in.** The Uptime Institute's 2025 survey notes average Power Usage Effectiveness (PUE)[1] has been relatively flat (average of 1.56, see Figure 6), partly because a large base of older facilities dilutes improvements, underscoring why reinvestment and retrofit cycles are a permanent feature.

The International Energy Agency (IEA) frames data centres as a material and fast-growing electricity load: ~415 TWh in 2024 (~1.5% of global electricity), with the United States holding the largest share of consumption (45%), followed by China (25%) and Europe (15%).

Figure 4: Electricity consumption in data centres by server type, and global and regional electricity demand 2024-2030

Note:
* Space and water heating in buildings.
TWh = terawatt-hour
AI = artificial intelligence
EIT = energy-intensive industry
US = United States
AE= advanced economies
EMDE = emerging market and developing economies.

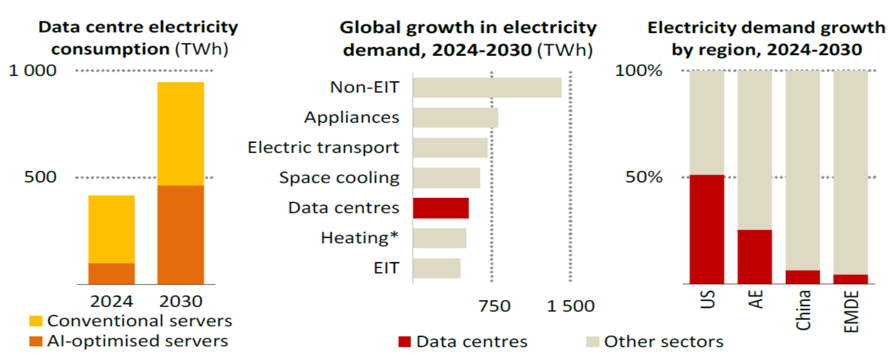

Source: IEA World Energy Outlook 2025

[1] Power Usage Effectiveness (PUE) is a key industry metric measuring data centre energy efficiency, calculated as the ratio of total facility power to IT equipment power. A PUE of 1.0 is ideal, meaning all power goes to computing; a lower ratio (typically 1.2-1.5) indicates better efficiency and lower cooling/overhead costs.

The reclassification of data centres from real estate to infrastructure is no longer merely an academic debate, as it has practical consequences for how capital is allocated, how assets are valued, and who owns them. Traditional real estate investors underwrite to occupancy, rental reversion, and capitalisation rate. While infrastructure investors underwrite to contracted cash flow duration, essential service characteristics, and the defensibility of the asset's market position.

Data centres increasingly exhibit the latter profile: long-term take-or-pay leases with investment-grade counterparties, revenue streams that are largely insensitive to economic cycles, and an asset base that is genuinely difficult to replicate due to power and permitting constraints. The growth of infrastructure fund ownership, through vehicles managed by Blackstone, KKR, EQT, Brookfield, and DigitalBridge among others, reflects this reclassification in practice. For listed investors, the implication is that the appropriate peer group for valuation and benchmarking is no longer purely other REITs, but increasingly includes listed infrastructure platforms, where different discount rates, hold periods, and income expectations apply.

1.3 Operators vs REITs: Business Model Differences and Key Operating Metrics

Operators monetise a service stack (colocation, interconnection, managed services) while REITs monetise contracted property cash flows and distribution capacity. This difference drives what "good performance" looks like, which metrics matter, and what valuation frameworks are appropriate.

Table 2: Key Operating Metrics - Data Centre Operators vs REITs

Metric	Operator Focus	REIT Focus
Occupancy / Utilisation	Billable power as % of total IT load capacity (MW utilised / MW available).	Leased area or contracted MW as % of total leasable capacity.
Revenue per MW	Colocation revenue per MW per month; blended across retail, wholesale, and hyperscale contracts	Annualised base rent per MW or per ft^2; typically fixed for lease duration.
EBITDA Margin	Typically 40-55% for scaled retail colocation operators; wholesale-heavy platforms typically run tighter margins; management teams focus on expansion through operational leverage.	Less directly comparable; distributable income margin more relevant.
Power Usage Effectiveness (PUE)	Critical efficiency metric: total facility power ÷ IT load power; best-in-class hyperscale approaches 1.20x; industry average is 1.54x (Uptime Institute, 2025).	Relevant but secondary; PUE affects operating cost, which flows into NPI margin.
Net Property Income (NPI) Margin	Not applicable.	NPI as % of gross revenue; reflects operating cost efficiency including utilities, property tax, and maintenance.
Contract Tenure / Churn	Average initial contract term (typically 1-3 years retail colocation, 7-15 years wholesale/hyperscale); renewal rate and churn rate are key forward indicators.	Not applicable as a standalone metric; subsumed into WALE.
Weighted Average Lease Expiry (WALE)	Not applicable; operators track contract duration and renewal rates by customer cohort but do not aggregate or report these as a standardised metric.	Key credit metric; DC REITs typically target WALE of 5-9 years, reflecting the long-dated nature of hyperscale and wholesale leases; longer WALE reduces re-leasing risk.
Distribution Per Unit (DPU)	Not applicable; operators do not distribute; they reinvest.	Core investor metric; expressed in cents per unit; year-on-year DPU growth is a primary performance signal.
Distribution Yield	Not applicable.	DPU ÷ unit price; benchmark for income investors; DC REITs in Singapore typically trade at 4-7% yield.
Debt/EBITDA	Primary leverage metric for operators; typically 4-6x for investment-grade rated operators.	Secondary; MAS aggregate leverage limit of 50% total borrowings to total assets is the binding constraint.
Interest Coverage Ratio (ICR)	Covenant metric in loan agreements; typically >3x for investment-grade operators.	Regulatory floor: MAS mandates minimum 1.5x ICR; rising to 2.5x for S-REITs with leverage above 45%.
Price/Earnings or EV/EBITDA	Primary valuation multiples for listed operators (Equinix, Digital Realty); EV/EBITDA typically 18-25x for scaled platforms, though the range compressed significantly post-2022 as interest rates rose.	Less commonly used; replaced by Price/NAV and dividend yield as primary valuation anchors.
Price-to-NAV (P/NAV)	Not applicable.	Core REIT valuation metric; compares unit price to appraised net asset value per unit; premium indicates market confidence in growth pipeline.

Source: Golden Section Capital

Metric	Operator Focus	REIT Focus
Funds From Operations (FFO) / Adjusted FFO	Used by US-listed operators (notably Equinix) as a REIT-equivalent metric; adds back depreciation to net income.	Equivalent to distributable income in the Singapore framework; basis for DPU calculation.
Contracted MW Pipeline	Forward indicator of revenue growth; signed but not yet commissioned capacity; critical for hyperscale-focused operators.	Relevant where REIT has development assets; reflects embedded NAV uplift potential.

Source: Golden Section Capital

Singapore's REIT framework adds a crucial constraint layer: the Monetary Authority of Singapore (MAS) rationalised leverage requirements such that all REITs are subject to a single aggregate leverage limit of 50% and a minimum interest coverage ratio (ICR) of 1.5x (rising to 2.5x for leverage above 45%). That matters because data centre growth is capex-hungry; and growth is therefore often constrained by equity valuation and debt capacity rather than "availability of deals".

Two additional structural differences follow from the table above. Firstly, operators have discretion over capital allocation in that they can retain earnings, acquire assets, or fund development from operating cash flow. Whereas S-REITs are required to distribute at least 90% of taxable income to maintain tax-transparent status, which structurally limits internal funding capacity.

Secondly, the depreciation treatment differs materially: operators carry data centre assets on their balance sheets and depreciate them, which suppresses reported net income relative to cash generation. This is why FFO and EBITDA are the standard analytical currency for operators, while distributable income and DPU serve the equivalent function for REITs. Investors who compare operator P/E multiples with REIT distribution yields without adjusting for this difference will systematically misread relative value.

1.4 The Data Centre Value Chain: Where Listed Companies Sit in the Ecosystem

The data centre value chain is a stacked system: power procurement and grid access → land and permits → shell/core build → fit-out → operations → connectivity ecosystem → workloads. Listed exposure usually captures only a slice of that chain, so "data centre exposure" is not one thing.

Figure 5: Data Centre Value Chain and Risk Allocation

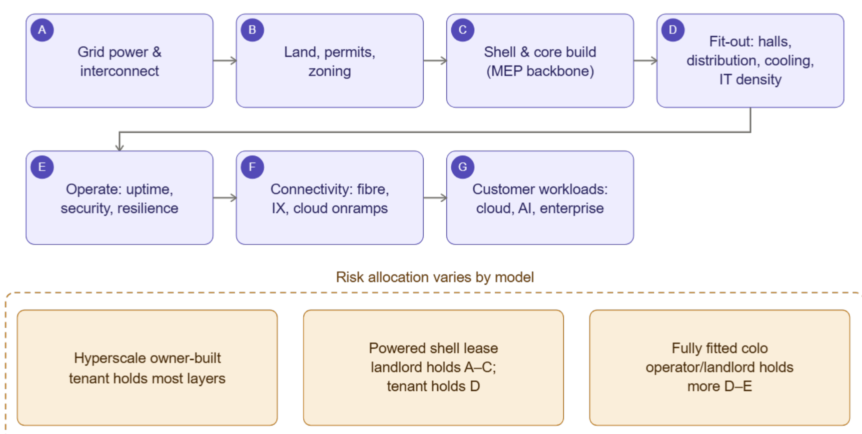

Source: Golden Section Capital

Where Singapore-linked vehicles often sit:

- **S-REIT data centre owners** (Keppel DC REIT, NTT DC REIT, Digital Core REIT): are primarily in the "own, lease, finance" layer with varying exposure to operations and fit-out decisions.
- **Platforms/operators (private or corporate)** (e.g., STT GDC via Singtel/KKR control): more direct exposure to build/operate/grow dynamics and power pipeline.

2. How Data Centres Make Money and How They Are Valued

2.1 Revenue Generation: Pricing by the Kilowatt, Lease Structures, and Power Cost Allocation

Data centre revenues are built around power-backed capacity (kW/MW) and, depending on the model, layers of connectivity and services. Lease structures often include escalators and pass-throughs, but the exact allocation of power cost risk is contract-specific and frequently under-disclosed in public summaries. Power is the defining economic variable in data centres, and how it is contractually allocated determines whether an asset behaves like stable infrastructure, a partially exposed operating business, or an implicit energy position.

In full pass-through models, tenants bear power costs, resulting in stable, infrastructure-like margins but limited upside. More commonly, partial pass-through structures introduce hidden volatility, where cost recovery is incomplete or delayed. At the extreme, fixed pricing entirely shifts the full power risk to the operator.

CBRE's North America market data illustrates why "power is the unit": the average monthly asking rate for a 250-500 kW requirement rose to US$195.94 per kW/month in H2 2025 (up 6.5% from H2 2024's US$184.06), while pricing for larger 3-10 MW requirements rose 12.5% YoY, reflecting competition for contiguous, scalable power. CBRE's global data highlights that the same dynamic is visible across APAC, where supply constraints are driving an even wider pricing dispersion. Singapore commands US$310-$470 per kW/month, Tokyo ranged between US$190-$355 per kW/month in Q1 2025 (a slight softening from US$200-$370 a year earlier as new suburban supply was absorbed), while Sydney declined modestly to US$145-$215 per kW/month.

Table 3: Data Centre Investment Structures - Operating Company versus REIT

Dimension	Data Centre Operating Company Private or Publicly Listed Not a REIT Structure	Data Centre Real Estate Investment Trust Listed Vehicle That Owns and Leases Data Centre Assets
Structure How the entity is organised and regulated	A standard company privately owned or publicly listed. Subject to normal corporate tax. No legal obligation to distribute earnings; can retain and reinvest profits freely. Growth is self-funded from operating cash flow or debt.	A regulated investment vehicle that qualifies for favourable tax treatment in exchange for distributing the majority of taxable income to shareholders (typically 90% or more). This constraint is structural and does not change in a downturn. Capital for growth must therefore come from issuing new debt or new equity, not retained earnings.
Revenue Unit What the customer is paying for	Varies by customer type and is critical to understanding margin: **Wholesale / hyperscale:** Priced per megawatt (MW) of committed power capacity. The tenant reserves power and pays for it whether or not they use all of it, similar to reserving a block of electricity, not a physical room. **Retail colocation:** Priced per cabinet or cage, per kilowatt (kW) of actual power consumed, plus fees for each network connection (cross-connect port) and any managed services. Operators with both segments have a blended revenue mix. Conflating MW pricing with cabinet pricing distorts any margin analysis.	Underlying revenue units are the same, MW capacity for wholesale tenants, cabinets and kW for retail colocation tenants, but the REIT passes this income through to investors rather than retaining it. The revenue mix matters because MW-priced wholesale leases carry very different power cost exposure and margin profiles compared to cabinet-based retail colocation. A REIT with a large wholesale book is more exposed to power price movements than one dominated by retail colocation.
Growth Drivers What moves the top line	New capacity additions and development pipeline Bookings growth and contracted backlog (MW pre-sold but not yet energised). Ecosystem density, the more networks and cloud providers interconnected in a facility, the harder it is for any single tenant to leave. Geographic expansion into new markets.	Acquisitions of stabilised or development-stage assets. Asset Enhancement Initiatives (AEIs), targeted capital investment to upgrade capacity or energy efficiency, within existing buildings. Rental reversions when leases are renewed at higher rents. Financing efficiency, because REITs rely heavily on debt to fund growth, the cost and availability of capital is itself a growth lever.

Source: Golden Section Capital

Dimension	Data Centre Operating Company Private or Publicly Listed Not a REIT Structure	Data Centre Real Estate Investment Trust Listed Vehicle That Owns and Leases Data Centre Assets
Primary KPIs Key performance indicators (The metrics that drive the model)	Recurring revenue growth (year on year). Churn rate, the proportion of revenue lost when tenants leave or downsize. Installed MW and utilisation rate. Interconnection revenue and density, the number of cross-connect ports per facility. Bookings and backlog, contracted future revenue not yet reflected in earnings.	Occupancy and power utilisation rate. Weighted Average Lease Expiry (WALE), the average time remaining on all leases, weighted by income. A longer WALE means more income certainty. Distributable income per unit, the cash available to pay to shareholders after expenses. Loan-to-Value (LTV) ratio, debt as a proportion of asset value. A measure of balance sheet risk. Interest Coverage Ratio (ICR), how many times operating earnings cover interest payments. A low ICR typically signals financial stress.
Capex Profile Capital expenditure spent on building and maintaining assets	Three distinct capex categories that must be modelled separately: **Growth capex:** New builds and capacity expansions. High and lumpy, large upfront spend before any revenue is earned. **Fit-out capex:** Equipping a new or vacant space for a specific tenant. Often funded by the tenant in wholesale; operator-funded in retail colocation. **Refresh capex:** Replacing ageing equipment (cooling systems, power infrastructure, servers) to maintain performance and avoid outages. Frequently underestimated in financial models.	REITs have the same three capex categories but a structural tension: the distribution requirement means there is less retained cash to fund them. In practice: Maintenance and refresh capex is often the first item cut to protect distributions, creating deferred risk that shows up later as outages or tenant loss. Growth is funded primarily through acquisitions and debt rather than organic development. AEIs are the main tool for organic value creation within the existing portfolio.
Income & Distribution How earnings are treated	No obligation to pay dividends. The company decides how much to distribute and when, most high-growth operators pay little or nothing, preferring to reinvest. Total shareholder return comes from capital appreciation, not income.	Mandatory high distribution, the defining feature of the REIT structure. Investors receive regular income, which is the primary reason many own REITs. However, this means the REIT cannot self-fund large capital projects. Every major acquisition or development effectively requires raising new equity or debt, which dilutes existing shareholders or increases financial risk.
Valuation Approach How analysts and investors typically price the asset	**EV/EBITDA** (Enterprise Value divided by earnings before interest, tax, depreciation, and amortisation) the standard metric for technology and infrastructure businesses. **Discounted Cash Flow (DCF)** particularly relevant for businesses with large, contracted backlogs of future revenue. Cap rates are not appropriate for operating businesses with significant reinvestment requirements.	**Net Asset Value (NAV)** the estimated market value of all assets minus liabilities. The primary anchor for REIT pricing. **Capitalisation rate (cap rate)** net operating income divided by asset value. Widely used but can be misleading (see below). **Price-to-Funds From Operations (P/FFO)** FFO is a REIT-specific earnings measure that adds back depreciation, which is a large non-cash charge in real estate. Price-to-FFO is the REIT equivalent of a price-to-earnings ratio.
The Valuation Trap The most common investor error	Applying a cap rate to an operating company treats it as a passive property investment when it is not. The business requires continuous reinvestment in equipment, power infrastructure, and capacity. A cap rate ignores all of this. Use EV/EBITDA and model the full capex cycle, including refresh spend that management teams routinely understate in their own guidance.	Cap rates applied to headline net operating income make a data centre REIT look attractively priced if the capex reserve is ignored. Maintenance and refresh capex can amount to 15-25% of revenue in an ageing portfolio. Understating that this flatters distributable income, inflates the apparent yield, and obscures the real cost of keeping the assets competitive. Always deduct a realistic capex reserve before applying any income multiple.

Source: Golden Section Capital

Table 4: Revenue Models and Lease Structures: Pricing Units, Escalations, and Pass-Throughs

Model	Typical Pricing Unit	Typical Term Range	Escalation Mechanisms	Power Cost Treatment	Key Investor Questions
Retail colocation	$/cabinet, $/kW-month + services.	Often shorter, renewal-heavy.	Fixed step-ups, CPI-linked, repricing on renewals.	Often metered; but billing mechanics vary.	What % of revenue is power pass-through vs margin revenue?
Wholesale colocation	$/kW-month or $/MW-month.	Often longer than retail.	Negotiated step-ups, indexation.	Often contractual pass-through + demand charge allocation.	Are there caps, lags, or floors that leak power volatility into NOI?
Hyperscale BTS lease	MW blocks, sometimes bespoke.	Often long.	Negotiated.	Typically structured to allocate utilities to tenant.	If tenant leaves, how re-tenantable is the asset without major capex?
Powered shell lease	$/m^2 plus power rights or $/kW.	Often long; net-lease common.	Negotiated.	Tenant often carries more opex.	What happens if power is not delivered on schedule?

Source: Golden Section Capital

2.2 Capital Intensity: Development Capex, Maintenance Capex, and Upgrade Cycles

Data centres are capital-intensive by design. Shell and core construction is expensive, and ongoing reinvestment is structural rather than discretionary. Development economics are increasingly shaped by labour constraints, lead times, and electrical equipment availability, pushing costs higher and elongating timelines.

JLL estimates average global construction cost increased from US$7.7m/MW (2020) to US$10.7m/MW (2025) and forecasts US$11.3m/MW for 2026. Turner & Townsend positions its Data Centre Construction Cost Index as a sector-specific cost benchmark.

Table 5: Capex Anatomy: Development, Maintenance, Upgrade Cycles, and What Breaks First.

Capex Bucket	What It Buys	Common Timing	What Tends to Go Wrong	Investor Modelling Implication
Development Capex	Land improvements, shell, MEP backbone, grid works.	Pre-Capacity on Demand.	Power delivery delays, equipment lead times.	Add scenario for timeline slippage, not just cost overrun.
Fit-out Capex	Data halls, distribution, cooling architecture.	Pre-lease or per-tenant.	Stranded design if density changes.	Separate “tenant-funded” vs “landlord-funded” economics.
Maintenance Capex	UPS/battery refresh, generators, switchgear, cooling.	Cyclical.	Underbudgeting increases outage risk.	Treat as recurring, not “one-off”.
Upgrade Capex	Liquid cooling enablement, efficiency retrofits.	Demand/regulation driven.	Spend becomes mandatory to remain competitive.	Include step-changes rather than smooth percentages.

Source: Golden Section Capital

2.3 Power as the Critical Input: PUE, Power Costs, On-Site Generation and Water

Power is both the essential input cost and the binding capacity constraint; efficiency and carbon pathways now shape both licensing and growth optionality. PUE remains the headline efficiency metric, but it is not a full sustainability metric and should be interpreted alongside disclosure quality and retrofit requirements.

The Uptime Institute reports an industry average PUE of 1.54 in its 2025 survey, noting that this average masks dispersion between modern and legacy facilities. This average has been relatively flat since 2021 (1.56 average over 5 years) implying that marginal efficiency gains increasingly come from capex-heavy redesigns rather than cheap operational tweaks.

Figure 6: Data Centre Average Annual PUE

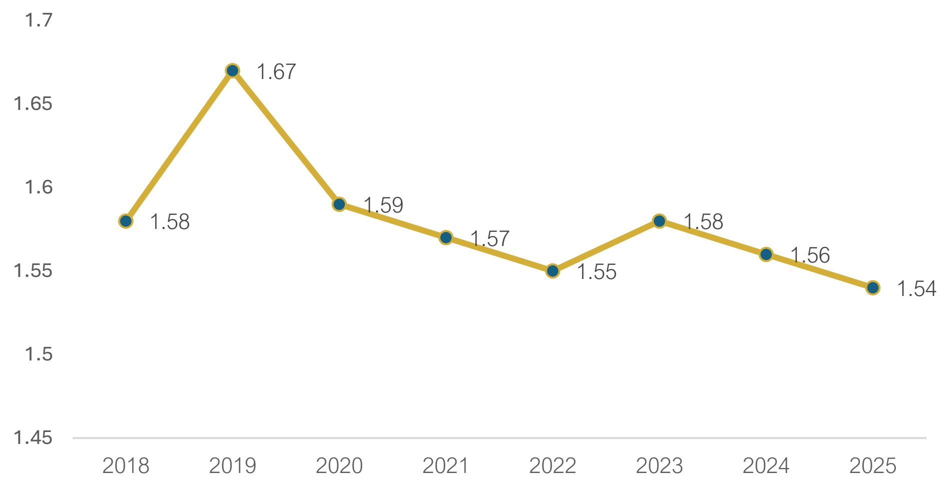

Source: Uptime Institute Global Data Center Survey 2025

Singapore's DC-CFA2[2] makes the ESG dimension explicit and investable. The programme requires, among other items: (i) BCA-IMDA Green Mark for Data Centres 2024 Platinum certification, (ii) PUE of 1.25 at 100% IT load or better, and (iii) at least 50% power from eligible green pathways.

Grid connection is no longer the only route to energised capacity. The combination of long interconnection queues, grid upgrade timelines measured in years rather than months, and the carbon commitments of hyperscale tenants has accelerated a structural shift toward behind-the-meter and on-site generation. Fuel cells (particularly solid oxide and molten carbonate variants) are being deployed at scale by operators in the United States, offering high efficiency, low emissions, and independence from grid congestion. Natural gas turbines and combined heat and power systems provide dispatchable baseload where grid reliability is uncertain. Solar and wind installations, typically paired with battery storage, are increasingly structured as long-term power purchase agreements directly with generation assets rather than through utility tariffs. Small modular reactors are at an earlier stage of commercial deployment but are attracting serious interest from hyperscalers seeking reliable, zero-carbon baseload at the scale required for large AI campuses.

For investors, on-site generation capability is becoming a valuation-relevant differentiator. An asset with secured, contracted on-site power is less exposed to grid queue risk, utility tariff volatility, and regulatory uncertainty around green energy procurement. Conversely, platforms that depend entirely on grid connection for all phases of a development pipeline carry an execution risk that should be explicitly modelled. Disclosure on generation mix, PPA terms, and the proportion of capacity that is grid-independent is therefore a material analytical input.

Water is the other vital component and Water Usage Effectiveness (WUE) measured as litres of water consumed per kilowatt-hour of IT load, is emerging alongside PUE as a core operational and regulatory metric. Cooling systems account for the majority of data centre water consumption, with evaporative cooling towers in particular drawing significant volumes in warm or water-stressed locations. In markets such as the US Southwest, parts of Southeast Asia, and the Middle East, water scarcity is becoming a genuine site selection and permitting constraint, analogous to the role that power availability plays in European and Singaporean markets. Singapore's DC-CFA2 framework already embeds sustainability thresholds that extend beyond PUE, and broader water-related disclosure requirements are emerging in the EU and in several US states. Investors should treat WUE alongside PUE as part of the standard ESG operating metric framework for data centre assets, and should assess whether management teams disclose site-level water consumption data or aggregate only at portfolio level, where individual high-consumption assets can be obscured.

[2] Singapore's Data Centre Call for Applications (DC-CFA2), administered by the Infocomm Media Development Authority, allocates data centre capacity directly and ties approvals to quantified efficiency and sustainability thresholds, including PUE ≤ 1.25 at full IT load, Green Mark Platinum, and at least 50% power from eligible green pathways.

2.4 Understanding Cash Flows: Contracted Revenue vs Earnings Volatility

Data centre cash flows can be contract-backed yet still volatile because timing, capex, and pass-through mechanics matter as much as headline occupancy. Investors should underwrite to "distributable cash after maintenance and upgrade capex", not accounting earnings or headline rent.

Singapore-listed REIT disclosures show how "contracted" still requires judgement, for example: Keppel DC REIT explicitly computes distribution based on distributable income after setting aside capex reserves and other reserves, indicating that cash distributions are structurally linked to reinvestment policy rather than purely rental receipts.

In Keppel DC REIT's own portfolio update, valuation methodologies differ by region: the presentation notes that for certain Europe/UK assets valuers adopted a discounted cash flow model rather than cap rates, implying cap rates are not always the natural valuation language for data centres. In February 2026 Mapletree Industrial Trust shows that 74.6% of its data centre portfolio (via its JV interest) is on triple net leases, which shifts many outgoings to tenants but does not eliminate refinancing and vacancy risk, as shown by its discussion of managing upcoming North American expiries.

A simple investor heuristic that usually improves models: build a "power-to-cash" bridge and force every revenue line to declare whether it is (a) margin revenue, (b) pass-through, or (c) capex recovery in disguise. If management cannot explain it, it is probably because it is complicated, but not immaterial.

2.5 Valuation Frameworks: DCF, Yield-on-Cost, and Replacement Cost

Data centre valuation works best as triangulation:

1. Using DCF to capture contractual ramps and reinvestment,
2. Yield-on-cost to test development economics, and
3. Replacement cost as a sanity check in scarcity markets.

Each framework fails in predictable ways, particularly when power delivery or obsolescence is mis-modelled. JLL's cost anchors (US$11.3m/MW average global construction cost forecast for 2026) provide a replacement-cost benchmark, but investors must be explicit about what is included (land, grid works, redundancy tier, fit-out scope). Meanwhile, pricing anchors such as CBRE's US$195.94/kW-month asking rate help frame revenue, but contract specifics drive the real economics.

Where a discounted cash flow output diverges materially from replacement cost, the divergence should be decomposed into development margin and revenue-anchor error before it is treated as value.

Table 6: Valuation Frameworks and Failure Modes

Framework	Best For	Biggest Failure Mode	What to do About it
Discounted Cash Flow (DCF)	Assets with ramps, complex pass-throughs, varied lease terms.	Modelling NOI without capex reality; ignoring power delay optionality.	Separate maintenance vs upgrade capex; includes power delivery scenarios.
Yield-on-Cost (YoC)	Development and forward funding underwriting.	Treating "stabilised yield" as guaranteed without ramp risk.	Add ramp, delay, and exit valuation assumptions; compute IRR not just YoC.
Replacement Cost	Boundary check in scarcity markets.	Assuming "replacement is possible" when power permits are the real moat.	Treat as upper bound; adjust for non-replicable power rights.

Source: Golden Section Capital

2.6 Why Cap Rates Can Be Misleading in Data Centre Valuation

Cap rates can mislead because data centre "NOI" is not standardised and can be distorted by (i) pass-through accounting, (ii) service revenue inclusion/exclusion, and (iii) near-term capex needs that are value-preserving rather than optional. Put differently, cap rates are fine as outputs, but dangerous as a sole valuation input.

Two concrete examples illustrate how cap-rate applications can distort underlying economic realities:

- **NOI definition drift:** if one portfolio's revenue includes large power pass-through and another excludes it, the same "cap rate" can reflect different economic margins even if both are reported as "property income." Limited public standardisation exists across issuers.
- **Capex masking:** a low apparent capex reserve may inflate distributable income in the short term but can create future outages or expensive retrofits.

Investor conclusion: where cap rates are used, normalise NOI and explicitly deduct a maintenance-and-upgrade capex reserve that matches the asset's redundancy and cooling profile.

2.7 Valuation by Data Centre Type: How Metrics Differ Across REITs, Operators, Hyperscalers, and Colocation

Valuation multiples differ because growth engines and risk layers differ across data centre types. REITs are usually valued as yield plus financed growth; operators and platforms are valued more like growth infrastructure with optionality embedded in expansion pipelines and ecosystem effects.

A practical investor mapping:

- **Data centre REITs (e.g., Keppel DC REIT):** sensitivity is often to financing costs and speed of acquisitions. Keppel DC REIT reports aggregate leverage, cost of debt, and ICR sensitivity to rate changes, illustrating that financing is a first-order variable in equity outcomes.
- **Pure-play S-REIT data centre owners with global portfolios (e.g., Digital Core REIT, NTT DC REIT):** market pricing often reflects a blend of US/Europe/Japan fundamentals plus SGD investor preferences; disclosed occupancy/WALE and leverage give a window into renewal and refinancing risk.
- **Platforms/operators (e.g., ST Telemedia Global Data Centres - STT GDC):** private valuations can reflect expansion pipelines. STT GDC's stated design capacity (2.3 GW) and multi-market footprint were central to its S$13.8bn enterprise value in the 2026 control transaction.
- **Hyperscalers:** valuation is not tied to property cash flows; as for hyperscalers data centre capital expenditure does not generate a discrete income stream that can be valued independently as it is a cost of running the underlying technology business.

2.8 The Public-to-Private Valuation Gap: What It Means for Listed Investors

Private capital is visibly paying for platform scale, secured power pipelines, and expansion optionality, sometimes at valuation levels public-market investors struggle to justify during higher interest rate regimes.

For listed investors, the key is not to "follow the private money"; it is to identify whether listed vehicles have the same scarcity assets (power, permits, expandable campuses) and funding capacity to monetise them without dilution.

Two recent deal datapoints illustrate the private playbook:

- **atNorth acquisition (US$4bn EV):** In February 2026, Equinix and CPP agreed to acquire atNorth for US$4bn; atNorth reportedly has 1 GW of secured power and an ~800 MW pipeline planned over five years, explicitly framing "secured power" as value.
- **STT GDC control transaction (S$13.8bn EV):** February also saw KKR and Singtel announce that they will acquire full control of ST Telemedia Global Data Centres at an enterprise value of S$13.8bn, with STT GDC operating across 12 markets and holding 2.3 GW of design capacity.

Investor implication is that public-to-private gaps can persist if listed vehicles are constrained by leverage rules (e.g., 50% cap) and equity issuance appetite, even when underlying data centre fundamentals are strong.

atNorth, ICE03 Data Centre, Akureyri, Iceland

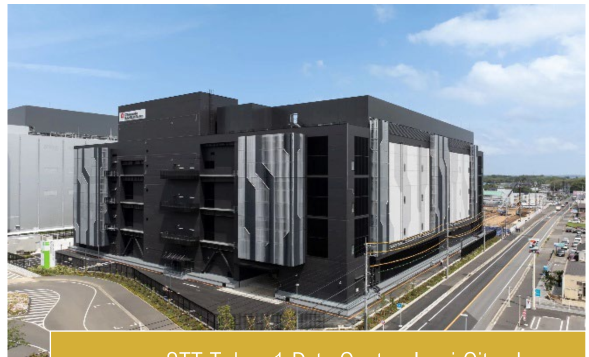

STT Tokyo 1 Data Centre, Inzai City, Japan

2.9 Pricing Power: Who Has It, When, and Why?

Pricing power is increasingly structural in power-constrained metros, and increasingly scarce in markets where capacity is replicable and grid access is easy. The winners tend to be owners/operators who can deliver MW on time, in the right location, with credible efficiency and compliance credentials.

CBRE's H2 2025 North America report shows pricing pressure persists even as the rate of growth slows: 250-500 kW asking rates reached a record US$195.94/kW-month, and larger 3-10 MW requirements saw 12.5% YoY increases, consistent with intensified competition for contiguous power. Singapore's regulatory gating (PUE ≤ 1.25; minimum 50% green energy) implies that compliant capacity is structurally scarcer than nominal MW counts suggest, which can support pricing power for "approved supply".

Examples

The following examples are illustrative, not investment recommendations. They deliberately separate (i) capacity pricing, (ii) opex, and (iii) recurring maintenance capex to show how results can swing under reasonable assumptions that are often hidden in headline cap rates. Pricing and cost anchors are sourced from CBRE and JLL; and model structure is generic.

Meta Data Centre, Tanjong Kling, Singapore

Table 7: Example - Unlevered DCF for a Wholesale Data Centre (Illustrative)

Item	Assumption	Rationale / Anchor
Capacity	20 MW (20,000 kW)	Illustrative "large facility" scale.
Starting capacity charge	US$195.94 per kW-month	CBRE H2 2025 average asking rate for 250-500 kW requirements (used as anchor; large-block pricing may differ).
Utilisation ramp	60% (Y1), 85% (Y2), 95% (Y3+)	Illustrative lease-up curve.
Annual escalation	2.5%	Simplified proxy for CPI-linked escalators.
Opex ratio	55% of revenue (NOI margin 45%)	Conservative normalisation to avoid NNN "fantasy" NOI.
Maintenance capex	2% of revenue	Illustrative recurring reinvestment.
Discount rate	8.5%	Illustrative unlevered rate.
Terminal cap rate	6.5%	Illustrative exit yield.
Result: Present Value	~US$304m	Implied value ~US$15.2m/MW (model output).

Note: This DCF uses a single asking-rate anchor; actual underwriting should use metro- and tier-specific pricing, contract terms, and power pass-through disclosures. The explicit forecast period is 10 years, with terminal value applied thereafter. The terminal value capitalises Year 11 NOI before maintenance capex, on the basis that a purchaser's own reserve is reflected in the exit yield; deducting the 2% reserve in the terminal year as well would reduce the present value to ~US$298m.

Reconciliation to replacement cost: The model output of ~US$304m implies ~US$15.2m/MW, against a replacement cost anchor of ~US$12.4m/MW (JLL's US$11.3m/MW 2026 forecast plus 10% soft costs). The ~22% premium is the development margin, and it arises solely because the implied stabilised yield-on-cost of ~8.5% exceeds the 6.5% terminal cap rate by roughly 200 bps. Capitalising stabilised NOI directly at 6.5% would produce ~US$325m, or ~31% above cost; the discounted model returns less because the lease-up ramp and the maintenance capex reserve absorb roughly nine percentage points of that uplift.

Two consequences: Firstly, the margin accrues to the developer, not to an acquirer. A purchaser paying US$15.2m/MW for the stabilised asset earns approximately the 6.5% exit yield, not the 8.5% yield-on-cost. Secondly, the spread is a direct function of the revenue anchor. The model applies a 250-500 kW retail-tier asking rate to a 20 MW wholesale block, and wholesale pricing per kW is materially lower. Table 9 shows what happens when that assumption is relaxed: at an effective capacity charge of US$170/kW-month the stabilised yield-on-cost falls to 7.0%, and at US$140/kW-month it falls to 5.8% (at or below the exit cap), at which point the premium to replacement cost compresses to zero or inverts. Replacement cost is the constraint that disciplines the revenue assumption.

Table 8: Example - DCF sensitivity - Discount Rate vs Terminal Cap Rate (Illustrative)

	6.0% terminal cap	6.5% terminal cap	7.0% terminal cap
7.5% discount rate	343.2	327.4	313.8
8.5% discount rate	318.4	304.0	291.6
9.5% discount rate	295.8	282.6	271.3

Interpretation: a 100 bps move in either discount rate or terminal cap rate can shift value by ~7-9% in this stylised model, before even touching utilisation or capex.

Table 9: Example - Yield-on-cost and project IRR scenarios for a 30 MW development (Illustrative)

Assumptions: cost per MW US$11.3m (2026 forecast) + 10% soft costs; opex ratio 55%; stabilised occupancy 95%; escalation 2.5%; terminal cap 6.5%; full development cost incurred at Year 0; 10-year operating hold; exit at 6.5% on forward NOI; no maintenance capex deducted.

Effective capacity charge (US$/kW-month)	Stabilised YoC (NOI/Cost)	Project IRR (2-year dev)	Project IRR (3-year dev, delay)
140	5.8%	6.3%	5.8%
170	7.0%	8.7%	8.0%
200	8.3%	10.9%	10.0%

Key point: A one-year delay reduces IRR by ~0.5 to 0.9 percentage points in these scenarios even when stabilised economics are unchanged. In constrained markets, time-to-power becomes a valuation risk.

3. Geography, Regulation, and Market Outcomes

3.1 Why Location Drives Data Centre Economics

Location determines latency, interconnection value, land feasibility, and, critically, deliverable power and regulatory permission to use it. As demand scales, it becomes vital to ensure users can get MW online, and this becomes the first-order determinant of rent growth, preleasing behaviour, and terminal values. This is clearly indicated by the world's premier data centre market, Northern Virginia, in the USA.

Data centres remain heavily concentrated geographically. The United States, China and Europe collectively account for roughly 82% of global installed capacity, and more than 85% of planned capacity additions over the next few years are expected to occur in these three regions.

Their impact on electricity demand, however, differs materially by market. In China and the European Union, data centres are projected to contribute around 6-10% of total electricity demand growth through 2030. In the United States, as the world's largest data centre market, the sector is expected to account for roughly half of the country's incremental electricity demand over the same period.

An IEA review of the global data centre development pipeline shows that more than half of projects under construction, or announced, are located in, or near, cities with populations exceeding one million, where electricity grids already serve large demand centres.

Figure 7: Characteristics of Data Centre Capacity Pipeline

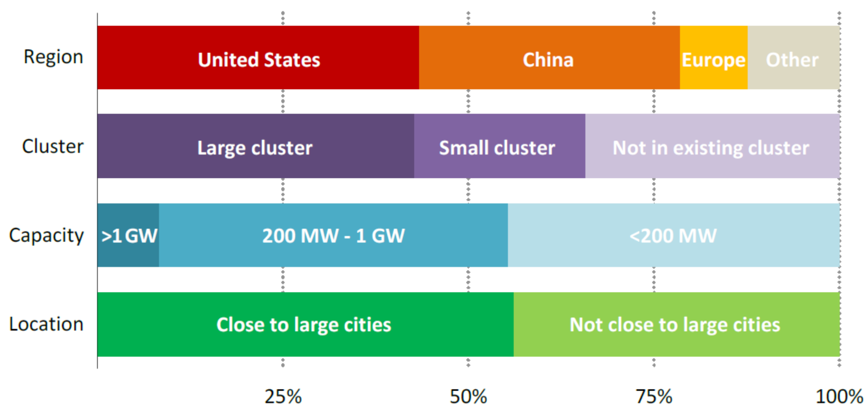

Source: IEA World Energy Outlook 2025

The scale of new facilities is also notable. Around 55% of pipeline data centres exceed 200 MW of capacity, with each facility expected to consume electricity equivalent to roughly 200,000 households once fully operational. In addition, close to two thirds of these developments are concentrated within established data centre clusters rather than new geographic markets.

Table 10: USA H2 2025 Wholesale Primary Market Fundamentals

Market	Inventory (MW)	YoY Change (MW)	Available MW	Vacancy Rate	YoY Change* (bps)	2025 Net Absorption (MW)	YoY Change (MW)	Rental Rates (US$kW/mo)** 250-500kW
Northern Virginia	4,039.6	1,109.5	21.5	0.5%	-5	1,102.0	650.3	$190-$235
Atlanta	1,459.2	458.8	28.5	2.0%	-60	456.0	-249.8	$165-$185
Dallas-Ft. Worth	1,067.3	476.3	26	2.4%	-113	470.8	424	$140-$175
Chicago	904.6	263.7	10.7	1.9%	-130	264.2	191.8	$190-$220
Phoenix	807.3	204.5	10.5	1.3%	-60	205.5	-40	$170-$210
Silicon Valley	489.2	20.8	23.2	4.7%	-76	-1.4	-46.6	$180-$275
Hillsboro, OR	475.4	0.0	1.0	0.2%	1	1.9	-240.2	$175-$225
New York Tri-State	190.0	0.0	13.6	7.2%	74	-1.4	-1.4	$180-$225

* Vacancy year-over-year change is between H2 2025 and H2 2024.
** Rental rates are quoted asking rates for 250+ kW at N+1/Tier III requirements.

Source: CBRE Research
CBRE North America Data Center Trends H2 2025.

Table 11: APAC H2 2025 Primary Market Fundamentals

Country	APAC Primary Markets	Operational Capacity	Under Construction	Planned	Colocation Vacancy	Rental Rates (US$kW/mo)* 250-500kW
Japan	Greater Tokyo	1,179MW	228MW	1,705MW	5.6%	$190-$355
Singapore	Singapore	1,043MW	20MW	237MW	4.9%	$310-$470
Australia	Sydney	786MW	189MW	1,102MW	3.0%	$145-$215
Greater China	Hong Kong	581MW	161MW	510MW	19.1%	$190-$350
Malaysia	Johor	897MW	315MW	2,099MW	0.7%	$90-$130

Source: Cushman & Wakefield Asia Pacific Data Centre H2 2025, * Monthly pricing from CBRE Research

Looking at 1-20 MW data centres, JLL research forecasts the strongest rental growth in the Americas through to 2030, followed by EMEA and then APAC.

Figure 8: Average 1-20 MW Lease Rates by Region ($/kW/month)

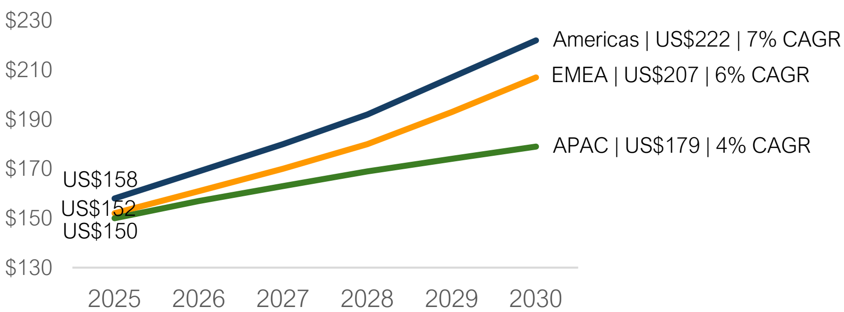

Source: JLL Global Data Center Outlook 2026

3.2 Constraints

Regulation is now a primary supply variable, as some jurisdictions ration capacity directly, while others impose disclosure and efficiency regimes that raise compliance costs and accelerate obsolescence in older stock. In both cases, policy is no longer peripheral.

Outcomes diverge because the binding constraint differs by location. In some markets it is policy, in others it is power availability and grid connection timelines, and increasingly it is both. The constraint set is local, often at metro or substation level. That is where pricing power, delivery risk, and returns are set. National averages mask interconnection queues, zoning friction, and transmission limits. Underwriting therefore must be bottom-up, anchored in specific grid nodes, permitting pathways, and utility counterparties.

In the USA, demand is not the issue. The US Energy Information Administration is forecasting record electricity consumption in 2026 (+1.9%) and 2027 (+2.5%), with incremental load from AI and crypto data centres. Infrastructure was not built for dense, latency-sensitive loads at this scale. The result is longer interconnection queues, higher upgrade costs, and a growing gap between contracted capacity and actual energisation.

Pricing reflects that constraint. CBRE reports North American colocation pricing of about US$195.94 per kW-month for 250-500 kW deployments, with 12.5% year-on-year uplift for 3-10 MW requirements. Smaller deployments pay for immediacy; larger blocks face time risk and limited availability. The common denominator is scarcity in tier-one hubs such as Northern Virginia, Dallas, and Silicon Valley, where both land and power are constrained.

By contrast, Singapore is explicitly policy constrained by the DC-CFA2 and capacity is effectively rationed. Approved megawatts can remain scarce even if developers are willing to build, supporting sustained pricing power for compliant assets and disadvantaging sub-scale or non-compliant operators.

Europe sits closer to a compliance-led constraint. In Ireland, operators face mandatory reporting of energy use and KPIs under the Energy Efficiency Directive and Delegated Regulation 2024/1364. This does not cap capacity directly, but it raises operating complexity and tightens scrutiny on energy performance, which over time constrains inefficient stock.

The synthesis is straightforward. Grid-constrained markets carry queue risk, capex volatility, and delivery uncertainty. Policy-constrained markets impose volume caps, higher barriers to entry, and more predictable but rationed growth. Where both apply, structural undersupply emerges, but only for those with secured and energised capacity. Broad statements about "strong US demand" or "Asian supply constraints" are not investable. The relevant questions are specific: which metro, which substation, what queue position, and under which regulatory regime. The period of unconstrained data centre development is over. Governments are now active allocators of capacity, either directly through permitting and quotas or indirectly through energy policy and reporting regimes. Investors who treat regulation as a secondary variable will misprice both risk and return.

3.3 Regional Comparison: ASEAN, Broader APAC, and US/Europe

Regional outcomes are diverging based on two primary constraints: policy gating and grid queues. While governments are committing billions to AI as "strategic infrastructure," the physical ability to energise these sites remains the ultimate arbiter of value. Investors must underwrite at the metro level; country-level averages are a blunt instrument that masks localised bottlenecks.

Key Performance Indicators and Regional Signals

- **US Grid Pressure:** The US Energy Information Administration (EIA) forecasts record electricity consumption through 2027, catalysed by AI and crypto. This macro-grid constraint is driving price escalation in prime hubs. CBRE's H2 2025 North America average asking rate for 250-500 kW requirements reached US$195.94/kW-month, up 6.5% year-on-year, while larger 3-10 MW requirements saw a steeper 12.5% uplift as competition for contiguous, scalable power intensified.
- **Singapore's Rationed Model:** The DC-CFA2 framework represents the ultimate "policy gate," where the state explicitly allocates MW based on quantified requirements. Here, policy is not just an incentive; it is the direct regulator of supply.

Governments are treating AI and semiconductor capacity as sovereign assets, deploying massive capital to de-risk domestic supply chains. However, this creates a "subsidy-demand paradox" where capital is plentiful, but power is scarce.

Table 12: Global Policies and Strategic Focus

Region	Primary Policy Vehicle	Capital Commitment	Strategic Focus
United States	CHIPS and Science Act / AI Action Plan.	~US$30bn	Onshoring fabrication (TSMC/Samsung) and deregulation of AI ecosystems.
European Union	EU Chips Act / InvestAI.	~€80bn	Expanding "Gigafactories" (e.g., Infineon in Dresden) and EuroHPC supercomputing sites.
Middle East	KSA PIF / Abu Dhabi Digital Strategy.	~US$68.50bn	Transitioning to regional AI hubs; includes a US$25bn energy infrastructure play by ADQ and Energy Capital Partners (ECP).
Japan	AI Supercomputing Subsidy.	¥114.60bn	Robotics and domestic supercomputing capacity.
South Korea	National AI Computing Center.	~US$1.52bn	Targeted 15x increase in GPU capacity by 2030.
United Kingdom	AI Opportunities Action Plan.	"Growth Zones"	Aiming for a 20x increase in national compute capacity by 2030.

Source: Golden Section Capital

The global data centre development pipeline is concentrating in a small number of established hubs while simultaneously generating a second tier of emerging markets whose growth is directly linked to constraints in those hubs. Understanding the relationship between the two tiers is essential for investors evaluating APAC exposure.

Among established hubs, Northern Virginia remains the world's largest single market by installed capacity, with vacancy rates below 1% and asking rents that have risen sharply as power constraints have intensified. Singapore occupies an equivalent position in Southeast Asia: dominant, expensive, and supply-constrained by design. Both markets function as pricing anchors and ecosystem centres, but neither offers meaningful new development capacity in the near term. Frankfurt, London, Amsterdam, and Paris (the FLAP cluster) perform a similar role in Europe, with regulatory and grid constraints increasingly limiting new supply across all four.

The emergence of Johor Bahru, Malaysia as a major data centre market is the most direct expression of the overspill dynamic in APAC. Located immediately across the causeway from Singapore, Johor offers land availability, lower power costs, and access to Singapore's connectivity ecosystem without the supply restrictions of DC-CFA2. GDS Holdings' Nusajaya Tech Park, the Johor-Singapore Special Economic Zone, and a series of hyperscaler commitments have converted Johor from a speculative market into an active development pipeline. Cushman & Wakefield's Asia Pacific Data Centre Update reports that Johor more than doubled its operational capacity through 2025, growing from 401 MW to 897 MW, with a further 315 MW under construction and 2,099 MW planned. Tenaga Nasional Berhad, the transmission and distribution utility for Peninsular Malaysia, guides to national electricity demand growth of 4.0%-4.5% in 2026, with data centre load the principal incremental driver. Investors should note that utility-reported data centre pipelines are typically stated as contracted maximum demand under signed electricity supply agreements rather than energised load, and the gap between the two (and the rate at which this closes) is the relevant variable in this market.

Investors should be aware that Johor's growth is partly a function of Singapore's constraints rather than entirely independent demand, and that Malaysian regulatory developments (including tightened permit requirements introduced in late 2025) introduced execution risk that did not exist when many projects were announced.

The broader SIJORI growth triangle of Singapore, Johor in Malaysia, and the Riau Islands in Indonesia, is increasingly functioning as a single integrated digital infrastructure zone, with capacity, connectivity, and cooling resources distributed across three jurisdictions. Batam in the Riau Islands is attracting early-stage development capital, offering low land costs and proximity to Singapore's subsea cable ecosystem. Bangkok, Jakarta, and Kuala Lumpur are establishing independent demand bases driven by domestic cloud adoption, e-commerce, and government digitalisation programmes, though grid infrastructure and cooling water availability vary significantly across these markets.

Beyond Southeast Asia, India represents the most significant structural growth opportunity in APAC. Hyperscaler commitments to Mumbai, Chennai, and Hyderabad have accelerated sharply, and domestic regulatory reform has improved the financing environment for data centre development. Japan, and in particular the Tokyo and Osaka markets combine mature demand with severe land scarcity, making it one of the higher-cost development environments in the region but also one of the most defensible once capacity is secured.

The Middle East, led by Abu Dhabi and Riyadh, is rapidly scaling through sovereign-backed programmes that treat AI infrastructure as strategic national investment, though grid carbon intensity and cooling water requirements in desert environments impose sustainability constraints that complicate green energy compliance. The proliferation of these national programmes, and others like Canada’s US$2bn data centre expansion and Taiwan’s NT$15tn 2040 vision reinforces the status of data centres as the new base layer of the global economy.

For investors assessing hub selection, the relevant questions are whether the market's growth is demand-driven or overspill-driven; whether power and permitting constraints are structural or temporary; and whether the connectivity ecosystem is mature enough to support the tenant types being targeted. Markets that score well on all three are genuinely scarce.

Policy competition is inflating demand for compute while simultaneously putting immense pressure on energy transmission. Investors should view government grants as a demand signal but must verify the right-to-power at the specific substation level before assuming these projects are viable.

3.4 AI Consortia and the Implications for Independent Data Centre Owners

The formation of large-scale AI infrastructure consortia, examples which include: Stargate (OpenAI, SoftBank, Oracle, MGX and others), the AI Infrastructure Partnership (BlackRock, MGX, Microsoft, Nvidia), and various sovereign-backed vehicles in the Middle East and Asia, represent a structural development whose implications for independent data centre owners and listed REITs are not straightforward.

The bullish case is that these consortia are accelerating demand at a scale that independent platforms and REITs cannot match organically. Stargate's stated ambition to deploy US$500 billion in AI infrastructure over four years, and AIP's initial US$30 billion equity target, represent capital formation that will need to be deployed through, or alongside, existing operators. REITs with hyperscale-ready capacity, long WALEs, and established relationships with cloud providers are natural beneficiaries of this demand wave. Keppel DC REIT's acquisition of Singapore hyperscale assets in FY2025, and NTT DC REIT's California and Virginia portfolio, are examples of exactly the type of powered, compliant capacity that AI consortium demand requires.

The negative case is that the same consortia are vertically integrating in ways that could over time reduce the addressable market for independent landlords. When Microsoft, Nvidia, and sovereign wealth funds form a dedicated vehicle to own and develop AI infrastructure directly, they are in effect bypassing the traditional landlord-tenant relationship that underpins REIT income streams. Hyperscalers that build and own their own campuses (such as AWS, Google, and Meta) increasingly do so at scale, and represent both the largest source of demand and the most credible competitive threat to third-party colocation and wholesale platforms. If hyperscalers and sovereign wealth funds successfully build closed-loop infrastructure ecosystems, independent REITs face structural subordination: they will either be relegated to fighting over lower-margin tier-two enterprise tenants, or forced to accept severely compressed yields as passive, powered-shell landlords stripped of all pricing power.

The relevant investor question is whether demand will flow through independent data centre owners or be captured within vertically integrated structures. The answer is likely to be both, but the allocation between the two will shift over time, and assets with undifferentiated power and location characteristics are most exposed to disintermediation.

OpenAI, Stargate Data Center Campus, Abilene, Texas, USA

3.5 Access to Capital

Markets with enforceable legal frameworks, transparent permitting, and institutional-grade development ecosystems consistently attract deeper, cheaper pools of capital. Jurisdictions where approvals are opaque, or grid access is discretionary, trade at a discount, regardless of how strong headline demand looks.

The current cycle is better understood as an infrastructure build-out than a traditional real estate expansion. Talk of a global investment supercycle is directionally correct but risks papering over the dispersion in outcomes. Aggregate capacity may approach 200 GW by 2030, but the distribution of that capacity, and therefore of returns, will be highly uneven. Capital is concentrating in markets where delivery risk is visible, manageable, and contractually bounded. This is now the US, select European hubs, and parts of North Asia. Southeast Asia and emerging markets are attracting a more structured, cautious approach.

Transaction evidence reinforces this shift. Deals like the atNorth acquisition show that investors are increasingly pricing not just stabilised income streams but the embedded optionality in secured power and development pipelines. The asset being acquired is no longer just the building; it is the right to convert power into revenue over time. Platforms with demonstrable access to power, rather than just land or customer demand, command structural premiums as a result. Land bank “plus” concept strategies are being repriced downward, particularly where grid timelines remain uncertain.

There is a subtler implication here too, and this is that capital is bifurcating. Core capital is targeting de-risked, powered, partially leased assets with visible cash flow ramps. Opportunistic capital is moving upstream, but only where it can secure asymmetric exposure to power allocation or policy-driven supply constraints. The middle, which is speculative development without firm power visibility is increasingly difficult to finance on attractive terms, and at low risk.

3.6 Structural Constraints: Power Scarcity, Policy Intervention, and Valuation Implications

Power availability, grid timelines, and policy allocation mechanisms determine whether a project ever transitions from concept to cash flow. In that sense, development looks increasingly like an options portfolio. Each stage, from grid application to energisation, represents a conditional probability.

Value leakage tends to cluster at predictable points in the development timeline. Site selection looks low risk at first, but errors here cascade through everything that follows. Grid applications and interconnection studies introduce the first real uncertainty, especially in constrained markets. Permitting phases are increasingly embedding sustainability and efficiency requirements that alter both capex and schedules. Long-lead equipment, such as transformers, switchgear, and cooling systems has become a bottleneck. Construction risk is now secondary to power risk, though still relevant in tight labour or materials markets. Utility energisation remains the single most binary event in the entire sequence; delays here have a disproportionate impact on returns. Even after energisation, how quickly fit-out, customer onboarding, and stabilisation are achieved determines whether the asset transitions smoothly from development yield to operating yield.

Singapore offers a useful counterpoint. Rather than allowing these risks to play out across the development timeline, policy effectively front-loads them into the approval stage. Under DC-CFA2, capacity is only allocated to projects that meet stringent efficiency and sustainability criteria. If the criteria are not met, the project does not proceed. For investors this removes a layer of execution risk post-approval but replaces it with a binary gate upfront. The megawatt either exists, or it does not. There is limited scope for speculative underwriting making Singapore a more predictable market for investors and developers.

The practical implication for valuation means that investors cannot just look at cap rates, they must fully understand time-to-power, probability-weighted lease-up schedules, and contingency capex needs as explicit investment consideration inputs.

Cooling system piping, Google Douglas County DC, Georgia, USA

Compu Dynamics power distribution unit whip cable assemblies for powered shell DC, USA

4. Looking to the Future

The next phase of data centre expansion will be shaped primarily by artificial intelligence infrastructure, rising cloud adoption, and the growing power intensity of advanced computing. Industry forecasts from IDC estimate that the global AI infrastructure market could reach approximately US$758 billion by 2029, implying growth above 40% per annum from current levels. Hyperscale cloud providers are expected to account for more than 85% of this investment. The spending mix is heavily skewed toward hardware, which represents roughly 95% of the total as demand concentrates on GPUs, AI accelerators, networking equipment, and high-performance servers rather than traditional facilities.

A structural shift is also underway in AI focused workloads. Historically, most computing demand was linked to model training. Over the remainder of the decade, inference workloads are expected to expand rapidly as AI models move into production deployment across enterprises and consumer platforms. Current estimates suggest inference could rise from roughly one quarter of total AI compute spending today to nearly half by 2029. This transition materially expands the scale of required infrastructure because inference workloads must be delivered continuously, geographically distributed, and close to end users.

Power demand from AI infrastructure is therefore accelerating well beyond traditional data centre growth rates. The IEA sees AI-driven accelerated computing as the dominant force reshaping data centre electricity demand. Accelerated server consumption is projected to grow at approximately 30% per year in the base case, against 9% annually for conventional servers, and accelerated servers alone account for nearly half of the net increase in global data centre electricity consumption through 2030. Conventional servers contribute around 20% of the net increase, with other IT equipment and facility infrastructure (principally cooling) accounting for roughly 10% and 20% respectively. The growth is broad-based across facility types: enterprise, colocation, and hyperscale data centres all contribute meaningfully to the aggregate demand trajectory, reflecting the fact that AI workload adoption is not confined to any single operating model.

These technological shifts are significantly increasing rack power densities. Introl sees GPU racks that typically operated between 130 kW and 250 kW in 2024-2025 are expected to rise to between 250 kW and 900 kW by 2026-2027, with densities approaching 1 MW per rack later in the decade. This requires parallel innovation in electrical infrastructure, cooling systems and facility design. Liquid cooling, advanced thermal management and redesigned power distribution systems are becoming essential for next generation AI facilities.

At the same time, artificial intelligence is driving a new phase of cloud revenue expansion. Major cloud providers are currently reporting annual growth rates between 20% and 40% as enterprises scale AI applications across software development, analytics, automation and consumer platforms. Quarterly cloud revenues among the largest providers now range from roughly US$8 billion to US$50 billion. Backlogs and long-term contracts have expanded sharply, with some providers reporting contracted revenue pipelines exceeding US$155 billion and in certain cases more than US$500 billion. These multi-year commitments are increasingly linked to dedicated AI compute capacity.

The scale of capital expenditure required to support this growth is unprecedented. Hyperscale capital spending is expected to rise from roughly US$500 billion in 2025 to approximately US$660 to US$690 billion in 2026 as per capex guidance provided by Amazon, Alphabet, Meta, Microsoft, and Oracle. Spending composition is shifting toward shorter-lived compute assets such as GPUs and accelerators, which now account for an estimated 50-60% of capex at leading cloud platforms. As a result, financing structures are evolving. Joint ventures, equipment financing, operating leases and multi-party funding vehicles are becoming more common, with some initiatives mobilising more than US$100 billion in capital from sovereign wealth funds and private investors.

Despite aggressive investment, supply constraints remain acute. AI demand is currently exceeding available capacity across several critical inputs including GPUs, data centre space and grid power. Hyperscalers report that many GPU clusters are operating at near full utilisation, including older hardware generations. Power availability has become the most significant constraint on expansion, prompting long term renewable procurement agreements, colocation with generation facilities and strategic partnerships with utilities.

Expansion timelines are increasingly shaped by execution constraints rather than pure capital availability. Operators are responding by pre-ordering long-lead electrical equipment, adopting phased construction strategies and securing land and grid capacity years in advance. The United States currently delivers new data centre capacity at a materially faster pace than Europe or Asia due to more flexible land availability, permitting frameworks and energy market structures.

The utility sector is therefore becoming a central participant in the data centre growth story. However, transmission capacity and interconnection timelines remain the primary bottlenecks, with some large projects unlikely to be fully energised until the early 2030s.

Energy mix considerations are also shifting. While renewable generation and battery storage continue to expand, utilities increasingly emphasise the role of dispatchable power sources in supporting 24-hour AI workloads. Natural gas, existing nuclear plants and potentially small modular reactors are being positioned as key sources of reliable baseload supply. Nuclear power is receiving renewed attention as a long-duration, zero-carbon energy source capable of supporting high density data centre clusters.

Several broader structural themes are also reinforcing long term growth. Governments are accelerating investment in domestic AI infrastructure through industrial policy initiatives. The United States has allocated roughly US$30 billion under the CHIPS and Science Act to support semiconductor manufacturing and AI capability. The EU launched the €43 billion EU Chips Act alongside additional AI investment programmes. Interestingly, as seen in Table 12 the EU has seen more than €80 billion in chip-related investments. Similar initiatives are emerging globally, including large scale funding commitments in Saudi Arabia, Japan and China as countries seek technological leadership in artificial intelligence.

New service models are also expanding demand. Neo-cloud providers offering GPU-as-a-service are emerging alongside traditional hyperscalers, targeting enterprises that require AI compute without building dedicated infrastructure. These businesses typically rely on multi-year capacity reservation agreements, often structured as take-or-pay contracts, to finance large GPU clusters. In many cases the primary constraint is not electricity supply but the ability to deliver AI-ready facilities on schedule due to permitting, labour availability and equipment lead times.

Looking further ahead, emerging technologies such as quantum computing may eventually introduce a new layer of demand for specialised data centre infrastructure. Quantum and classical computing systems are expected to operate in parallel, with data centres offering hybrid access to both architectures. Although widespread commercial deployment remains some distance away, sectors such as pharmaceuticals, chemicals and advanced encryption are likely to be early adopters.

Taken together, these dynamics suggest that data centre demand over the next decade will be driven by a combination of AI compute growth, escalating power density, and strategic national investment in digital infrastructure. The sector's growth trajectory will therefore increasingly be determined by energy availability, grid expansion, and the ability of developers to deliver large scale facilities at speed.

Primer 1 has established the analytical foundation of what data centres are, the next question for investors is how to access this asset class, how to evaluate the vehicles through which exposure is available, and how to apply a rigorous investment framework to a sector that does not behave like conventional property.

These questions are the subject of the second primer in this series, and Primer 2 will examine the structural differences between operators and REITs, expand on the valuation frameworks appropriate to each business model, and apply them directly to the Singapore-listed data centre universe.

Huawei Cloud Data Centre, Guiyang City, PRC

Appendix A: Uptime Institute's Tier Classification System

The Uptime Institute's data centre classifications are divided into four Tiers that match a particular business function and define criteria for maintenance, power, cooling and fault capabilities. The Tiers are progressive, so each Tier incorporates the requirements of the lower Tiers. This progression does not automatically mean that a Tier IV data centre is better than a Tier II, it means that these levels fit differing business operations.

Operational sustainability is the second essential component of the Tier classification. It refers to the behaviour and risks apart from infrastructure design that determine the ability of the data centre to meet long-term business goals.

The data centre Tier levels are:

Tier I

A Tier I data centre is the basic capacity level with infrastructure to support information technology for an office setting and beyond. It has an expected uptime of 99.671% (28.8 hours of downtime annually).

The requirements for a Tier I facility include:

- An uninterruptible power supply (UPS) for power sags, outages, and spikes.
- An area for IT systems.
- Dedicated cooling equipment that runs outside office hours.
- An engine generator for power outages.
- Tier I protects against disruptions from human error, but not unexpected failure or outage. Redundant equipment includes chillers, pumps, UPS modules, and engine generators. The facility will have to shut down completely for preventive maintenance and repairs, and failure to do so increases the risk of unplanned disruptions and severe consequences from system failure.

Tier II

Tier II facilities cover redundant capacity components for power and cooling that provide better maintenance opportunities and safety against disruptions. The distribution path of Tier II serves a critical environment, and the components can be removed without shutting the data centre down. Like a Tier I facility, unexpected shutdown of a Tier II data centre will affect the system. It has an expected uptime of 99.741% (22 hours of downtime annually).

These components include:

- Engine generators.
- Energy storage.
- Chillers.
- Cooling units.
- UPS modules.
- Pumps.
- Heat rejection equipment.
- Fuel tanks.
- Fuel cells.

Tier III

A Tier III data centre is concurrently maintainable with redundant components as a key differentiator, with redundant distribution paths to serve the critical environment. Unlike Tier I and Tier II, these facilities require no shutdowns when equipment needs maintenance or replacement. The components of Tier III are added to Tier II components so that any part can be shut down without impacting IT operation. It has an expected uptime of 99.982% (1.6 hours of downtime annually).

Tier IV

A Tier IV data centre has several independent and physically isolated systems that act as redundant capacity components and distribution paths. Separation is necessary to prevent an event from compromising both systems. The environment will not be affected by a disruption from planned and unplanned events. However, if the redundant components or distribution paths are shut down for maintenance, the environment may experience a higher risk of disruption if a failure occurs.

Tier IV facilities add fault tolerance to the Tier III topology. When a piece of equipment fails, or there is an interruption in the distribution path, IT operations will not be affected. All of the IT equipment must have a fault-tolerant power design to be compatible. Tier IV data centres also require continuous cooling to make the environment stable. It has an expected uptime of 99.995% (26.3 minutes of downtime annually).

Source: Uptime Institute

References

Al Jazeera. (2026, March 24). Amazon says AWS Bahrain region 'disrupted' following drone activity.

AWS. (Accessed 2026. April 7). Security, Identity, and Compliance: Taiwan Data Privacy.

CapitaLand Ascendas REIT. (2026, January). Investor presentation

CapitaLand Ascendas REIT. (2026, April 27). 1Q 2026 Business Updates.

CBRE. (2025, June 24). Global Data Center Trends 2025. CBRE Research / Data Center Solutions.

CBRE. (2026, February 25). North America Data Center Trends H2 2025. CBRE Research / Data Center Solutions.

Challani, A. (2025, February 28). Data Localization Laws in India

Cushman & Wakefield. (2026, March 27). APAC Data Centre H2 2025 Update.

DCD. (2026, April 1). Fire at AWS data center in Bahrain after Iranian attack.

Digital Core REIT. (2026, January). Investor Presentation.

Digital Core REIT. (2026, April 22). 1Q26 Business and Operational Update.

DigiTimes. (2026, March 20). Analyzing NVIDIA's GTC 2026 Potential Strategies: How Will Jensen Huang Block ASIC Competitors?

Electrical Trader. (2026, February 2). Transformer Shortages: Supply Chain Impact on Pricing.

Equinix. (2026, February 27). CPP Investments and Equinix to Acquire atNorth for US$4 Billion (press release).

Government of Canada. (2024, December 5). Canada to drive billions in investments to build domestic AI compute capacity at home.

Government of Ireland. (2025, February 24). Data Centre Energy and Sustainability Performance Reporting Obligations.

Government of Taiwan. (2025, September). Promoting the Ten AI Initiatives - Implementing the policy vision of making Taiwan a smart technology island nation.

Hogan Lovells. (2025, September 15). The key characteristics of data centres in the US and Europe - overview for those involved in financings.

IDC. (2025, October 25). Artificial Intelligence Infrastructure Spending to Reach $758Bn USD Mark by 2029, according to IDC.

IDC. (2025, December 2). Asia/Pacific Datacenter Market to Grow 3x by 2029, Driven by Cloud, AI, and Hyperscale Growth.

Amazon Project Rainier 2.2 GW DC, New Carlisle, Indiana, USA

Indonesia Government. (Accessed 2026. April 7). Regulation of the Government of the Republic of Indonesia, Number 71 Of 2019 on The Organization of Electronic Systems and Transactions.

Infocomm Media Development Authority. (2025, December 1). Launch of second data centre - Call for application.

International Energy Agency. (2025, April). Energy and AI (executive summary and data centre demand sections).

Introl. (2026, February 20). Hyperscaler CapEx Hits $690B in 2026: Inside Microsoft's $80B Azure Power Bottleneck.

Introl. (2026, February 20). High-Density Racks: 100kW+ Designs for AI Data Center Infrastructure.

Jones Lang LaSalle. (2026, January 5). 2026 Global Data Center Outlook.

Keppel DC REIT. (2026, January 30). FY2025 Results Presentation and portfolio updates.

Keppel DC REIT. (2026, April 16). 1Q 2026 Operational Updates.

Kong, P. (2026, January 2). Renewable Energy. Research by Kenanga Investment Bank Berhad.

Mapletree Industrial Trust. (2026, February 17). Investor Presentation Slides.

Mapletree Industrial Trust. (2026, April 28). 4Q & FY25/26 Financial Results.

McKinsey and Company. (2025, August). The data center balance: How US states can navigate the opportunities and challenges.

Monetary Authority of Singapore. (2024, November 28). MAS rationalises leverage requirements and introduces additional disclosures for REITs. MAS media release.

NTT DC REIT. (2025, November 12). 1H FY25/26 Financial Update.

NTT DC REIT. (2026, May 12). FY25/26 Financial Results.

Picus Labs. (2026, March 12). Iranian Threat Actors: What Defenders Need to Know.

Reuters. (2026, February 3). Power grid delays challenge Amazon's data center expansion in Europe.

Reuters. (2026, February 4). KKR and Singtel pay $5.2 billion for full control of data centre operator STT GDC.

Reuters. (2026, February 27). CPP, Equinix to buy Nordic data center operator atNorth in $4 billion deal.

Reuters. (2026, April 7). US power use to beat record highs in 2026 and 2027 as AI use surges, EIA says.

Rust, M. and Parker, W. (2026, April 7). These Cities and States Are Taking Aim at Data Centers. Wall Street Journal.

Sightline Climate (2026, February). Q1 Data Center Outlook.

Singtel. (2026, February 4). KKR-led consortium with Singtel Group to fully acquire ST Telemedia Global Data Centres at S$13.8 billion enterprise value (news release).

Science|Business. (2026, February 19). Chips Act spurs semiconductor investments in Europe.

Tenaga Nasional Berhad. (2026, February 27). Analyst Briefing 4Q FY2025.

Turner & Townsend. (2025, November 5). Data centre construction cost index 2025-2026.

Uptime Institute. (2025). Global Data Center Survey Results 2025.

US Energy Information Administration. (2026, February). Fossil generation could rise with faster-than-expected growth in data center power demand.

Wood Mackenzie. (2026, March 16). Newly added US data center capacity slows down considerably in Q4 2025, as market struggles to keep up with explosive demand.

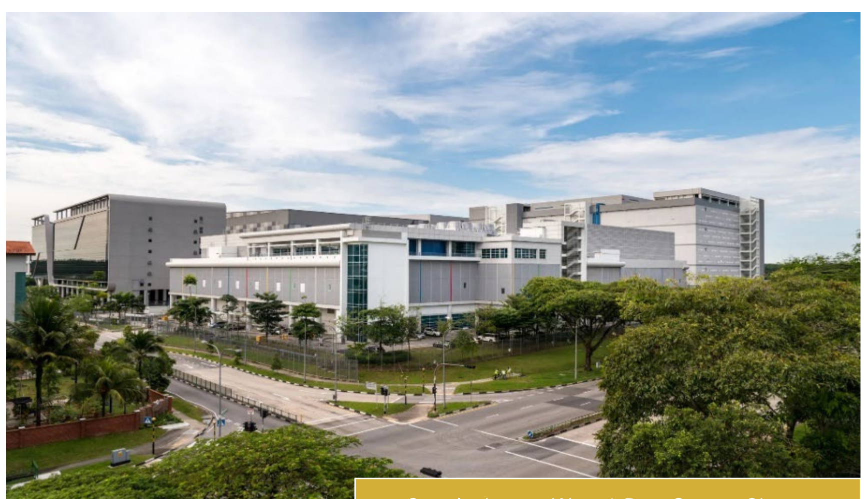

Google Jurong West 1 Data Centre, Singapore

Disclaimer and Copyright

The information contained in this document is being provided to you for informational purposes only and is not, and may not, be relied on in any manner as legal, tax, or investment advice. A Recommendation (Buy/Sell/Hold) is not provided. Forecasts and any valuation are the independent view of Golden Section Capital. Company management, or representatives, have no editorial input. The views expressed in this report represent those of Golden Section Capital at the time of publication and may be subject to change without notice.

Except where otherwise indicated herein, information provided herein is current as of 21 July 2026, and there is no obligation to update or otherwise revise such information after such date. Certain assumptions have been made in calculating the return information and preparing the other information set forth in this document. While made in good faith, there can be no assurance that such assumptions will prove correct. Any prices or quotations provided are for informational purposes only and should not be used for valuation or any other purpose. Golden Section Capital has not independently verified all information used in the preparation of this report. You should be aware that there may be errors or omissions in the information presented. This report uses information sources believed to be reliable, but their accuracy cannot be guaranteed.

This report is not a solicitation to buy or sell any product. The companies mentioned in this report may not be suitable for all investors, and certain legal jurisdictions may have restrictions on their sale. You are solely responsible for determining whether any particular security is appropriate for you in light of your investment goals, financial situation, and risk tolerance.

Certain information contained herein may constitute forward-looking statements, including, but not limited to, the key themes, outlooks and key strategic priorities and statements regarding potential liquidity events. Forward-looking statements may be identified by the use of terms such as “may,” “will,” “should,” “expect,” “anticipate,” “project,” “estimate,” “intend,” “continue,” or “believe,” (or the negatives thereof) or any other variations thereof.

Due to various risks and uncertainties, actual events or results or actual performance may differ materially from those reflected or contemplated in such forward-looking statements, and there can be no assurance that unrealised investments used to calculate the return information set forth herein will be realised for their assumed values. As a result, investors should not rely on such forward-looking statements. Any projection of the performance of a fund or an individual investment or targets provided by Golden Section Capital herein, or in any related discussion is highly speculative, and represents Golden Section Capital’s opinion, which may change. Any modelling, scenario analysis, or past performance included in this report is not indicative of future results.

To the fullest extent permitted by law, neither Golden Section Capital, nor the author(s) accept responsibility or liability for any loss or damage arising from the use of or reliance on any information contained in this report, even if due to negligence or errors in the information provided. The information is presented in good faith and is based on sources believed to be dependable, but independent verification may not have been conducted on all aspects of information, and/or data.

Golden Section Capital does not conduct any investment business and does not hold positions in the companies mentioned in this report.

This report is intended for institutional and professional investors who meet the experience requirements defined by applicable national laws and regulations. It is not intended for retail or non-qualified investors and may not be distributed in any jurisdiction where the information is prohibited.

This communication is for the intended recipient only. If you have received this report in error, please destroy it immediately and notify Golden Section Capital.

These disclaimers and exclusions shall be governed by and construed in accordance with South African law. If any provision of these disclaimers is deemed unlawful, void, or unenforceable, such provision shall be struck, and the remaining provisions shall remain valid and enforceable.