



INFRASTRUCTURE INVESTMENT

The Third Order

Capital is competing harder for yield than at any time since 2008. The technology underneath these assets is now changing faster than the models financing them. The gap between those two facts is where the return sits. This paper discusses whether the investments team are staffed for it.

SCG Partners · October 2026 · 11 minute read



Marina Bay Sands, Singapore, where Data Centre World Asia occupied three floors this week, against one and a half last year.

Data Centre World Asia occupied one and a half floors of Marina Bay Sands last year. This week it occupied three. Tech Week Singapore as a whole drew 38,000 attendees and 720 exhibitors across eleven country pavilions, up from nine, and the organisers chose The Infrastructure Era as their theme. Senior Minister of State Tan Kiat How opened it. The first Open Compute Project

Southeast Asia Tech Day ran alongside, built around how inference workloads will shape facility design through 2030.

Doubling your floor space in twelve months is a reasonable proxy for a market in a hurry. But the scale was not the interesting part. Our own observation from walking the floor, and the one we keep returning to, was the sheer technical complexity now sitting between the grid and the chip. That complexity is not a detail for engineers to resolve quietly. It is about to reprice a large amount of infrastructure and it is going to do so faster, in our view, than capital models are built to absorb.

This paper follows the one we published before the show, which argued that rising long dated yields are forcing a reckoning on business models built for cheap capital. This is the second half of that argument. If money has become expensive and impatient, and the asset it is funding has started to change shape every eighteen months, then the question for anyone deploying capital over fifteen or twenty years is no longer which data centre to buy. It is which layer of this industry still has a clock slow enough to finance.

The first order is crowded and priced

The direct play, owning the data centre and leasing it, has been thoroughly discovered. Blackstone with CPP Investments took AirTrunk at an enterprise value above twenty four billion Australian dollars. A KKR led consortium with Singtel committed 1.75 billion Singapore dollars to ST Telemedia Global Data Centres. A Stack Asia portfolio has drawn reported interest above thirty billion dollars. Roughly five hundred billion dollars of artificial intelligence related debt has been issued this year alone on Goldman Sachs estimates, with hyperscaler capital expenditure heading towards eight hundred and twenty billion dollars in 2026 and a trillion or more in 2027.

Nobody finds an information or competitive advantage in a market that crowded. Spreads have already begun to tell the story: a basket of high yield data centre bonds is trading wider than June 2022 levels, seventeen of twenty three data centre joint venture deals are trading wide to their originated yields, and S&P cut Oracle to BBB minus in July with its credit default swap spread near an eighteen year high.

The second order, briefly

We have argued at length that the more interesting return sits in what the data centre makes possible rather than in the data centre itself. A hyperscale operator is the first large, concentrated, creditworthy electricity buyer this region has ever had, and that changes the credit profile of everything downstream of it. Thailand turned a two gigawatt direct power purchase pilot designed for data centres into an open national market in eight months.

RMI's modelling of Peninsular Malaysia suggests that allowing data centre load to shift during under seven percent of the hours in a year defers up to eight hundred megawatts of new gas and draws in an extra 4.6 gigawatts of solar.

Grid, generation, storage, water and waste all become financeable on the back of that anchor. The reason the opportunity persists is mundane: a digital infrastructure fund cannot buy a waste plant and an energy transition fund will not underwrite a data centre, so the value sits in the gaps between mandates that firms wrote for themselves.

That argument survives this week intact. What this week added is a third layer, and it is the one that should worry people.

The third order: the asset has more than one clock

Here is what was actually on display at Marina Bay Sands.

Delta Electronics gave a keynote on the road from the grid to the chip at 800 volts direct current. Huawei presented on building grid interactive artificial intelligence data centres optimised for compute delivered per watt. Wärtsilä argued for securing power first and worrying about the grid connection second. Midea launched an architecture that fuses electrical power, energy storage, magnetic bearing chillers and liquid cooling distribution into a single integrated stack, with an industrial coolant distribution unit rated at 2.6 megawatts. Nidec showed an in rack coolant distribution unit delivering three hundred kilowatts of cooling in four rack units, and an in row unit at two megawatts.

Behind all of that sits one transition. NVIDIA has moved 800 volt direct current distribution from concept to production roadmap, with an MGX compatible power rack arriving in the second half of this year, row level power centres supporting two megawatts per row in 2027, and eventually facility scale conversion from 13.8 kilovolt grid supply straight to 800 volts direct current through 4.8 megawatt power blocks. ABB, Eaton, Schneider Electric, Vertiv and Delta are all building for it.

The reason is physics rather than fashion. A one megawatt rack running on the current 54 volt distribution would need up to sixty four rack units of power shelves, which is the entire cabinet, leaving no room for compute, and roughly two hundred kilogrammes of copper busbar per rack. At gigawatt campus scale that is two hundred tonnes of copper in rack busbars alone. The 800 volt architecture cuts copper requirements by something like forty five percent, improves end to end efficiency by around five percent and, NVIDIA claims, reduces maintenance cost substantially by removing power supply units from inside the rack.



Switchgear panels labelled at 13.8 kilovolts. The roadmap described here would convert from that supply level straight to 800 volts direct current.

Now hold that against where rack density has come from and where it is going. A conventional enterprise hall was designed around five to twenty kilowatts per rack. Today's high density artificial intelligence cabinets run at three hundred to five hundred kilowatts. The Kyber platform expected to ship in volume through 2027 is designed to draw up to 1.4 megawatts per rack.

That is a two order of magnitude change in the fundamental design parameter of the building, inside a decade.

Three consequences follow, and none of them are priced in most models we have seen.

The first is that colocation operators are expected to lag dedicated hyperscale campuses by two to three years on this transition, which creates a window in which a great deal of recently completed capacity is technically current and commercially second tier.

The second is that there is a genuine standards fight underway between the NVIDIA aligned 800 volt approach and Google's own rack voltage, which is expected to fragment the supply chain for switchgear and busways well into 2027. Anyone procuring electrical infrastructure this year is making a bet, whether they know it or not.

The third is the one that matters for capital. Air cooled facilities completed eighteen months ago cannot house the chips now shipping. The building is fine. The land is fine. The grid connection is excellent. The electrical and thermal architecture inside it is wrong, and fixing it is not a refurbishment.

What this does to the model

An institutional investor underwriting a data centre is in practice underwriting at least three assets with three different lifespans bundled into one structure.

The land, the building and the grid connection run on a twenty to twenty five year clock, sometimes longer. The electrical and cooling architecture now appears to run on something closer to a seven to ten year clock, and nobody knows precisely where it sits because the 800 volt transition is the first serious test. The silicon runs on a three to six year clock and is the subject of an unresolved accounting argument.

That argument is not academic. Microsoft extended the depreciable life of server and network equipment from four years to six. Oracle moved from four to five and now runs six. Amazon has shortened the assumed life of some hardware while Meta has extended others, which is two sophisticated operators reaching opposite conclusions about the same technology cycle. Michael Burry put the resulting understatement of depreciation across the sector at roughly one hundred and seventy six billion dollars between 2026 and 2028. The International Monetary Fund modelled a stylised three year useful life and found aggregate hyperscaler operating margin falling by more than nine percentage points with debt rising from around eight hundred billion dollars to over a trillion.

There is a serious counterargument and it deserves stating. Google reports running some eight year old tensor processing units at full utilisation, and the case for longer lives rests on workload cascading: frontier training moves to new silicon, older silicon picks up inference and less demanding work, and the asset keeps earning. Useful life on that view is set by utilisation rather than by accounting convention, and demand is broad enough to support it.

Both positions cannot be right, and the honest answer is that nobody knows yet. Which is precisely the point. An investment committee approving a fifteen year financing structure against an asset whose core economic assumption is actively contested by its own operators is not underwriting a risk. It is expressing a hope.

The building is fine. The land is fine. The grid connection is excellent. The electrical and thermal architecture inside it is wrong, and fixing it is not a refurbishment.

Meanwhile the money has got more expensive and less forgiving. The United States thirty year Treasury sits above five percent, the United Kingdom sold a thirty year gilt at the highest yield since comparable records began in 1998, the German ten year is at its highest since 2011 and the Japanese ten year crossed three percent for the first time in around three decades. Technology issuers are competing directly with sovereigns for the same capital and are largely insensitive to price. Data centre securitisations are expected to rise by close to half to more than thirty billion dollars. More than ten percent of new

single asset single borrower commercial mortgage issuance is now data centre paper.

So the structure is this. Long dated, inflexible capital, raised at the highest cost in a generation, secured against an asset containing a component that may be obsolete before the first refinancing.

Where the third order return actually is

The conclusion is not that this is a bubble. It is that the value is migrating between layers, and the migration is predictable even if the timing is not.

Value accrues to the layers with the slowest clock. Land with a firm grid connection. Generation and transmission. Storage. Water. Cooling infrastructure designed to be re laid rather than replaced. And the operating capability that keeps all of it running regardless of what is installed in the halls above it.

Value erodes fastest in the layers tied to a specific compute generation. That includes the electrical and thermal architecture inside the building, which is exactly the layer most investors treat as a construction line item rather than a strategic bet.



Insulated pipework and gauges. The paper argues that value is migrating towards the layers that outlast any one generation of compute.

This strengthens rather than weakens the second order argument. If the compute layer is where obsolescence lives, then the anchor offtake thesis becomes more attractive, not less, because the power, grid, storage and waste infrastructure the anchor makes financeable runs on a twenty five year clock and does not care which generation of chip is drawing the load.

It also has a consequence for the host economies we have been writing about. As value migrates towards power, land, water and grid, the governments of Malaysia, Indonesia, Vietnam, Thailand and the Philippines hold more leverage than they have been exercising. Singapore has already worked this out, which is why Tan Kiat How has said the country will develop capacity in a disciplined and calibrated way rather than adding megawatts for their own sake, with proposed legislation to impose efficiency requirements on existing facilities as well as new ones, and why the government has stated plainly that leaving this entirely to commercial arrangements would be irresponsible given the spillover effects on the wider economy. That is a government pricing the second order effects deliberately.

The challenge to capital providers

Here is the uncomfortable part.

If you are deploying capital into this sector and nobody on your investment team can independently assess whether a given facility's electrical architecture will be commercially viable in 2030, you are not underwriting the asset. You are underwriting the sponsor's assertion about the asset, and taking a fee for doing so.

Most infrastructure and real estate funds we deal with are staffed with real estate professionals, project finance specialists, asset managers and sustainability leads. That was the right team for a market where the building was the asset. It is the wrong team for a market where the building is the slowest moving part of a system whose fastest moving part determines the revenue.

The capabilities that are missing, almost universally, are these.

-
- 01 Someone who can read a silicon and power roadmap and translate it into a facility specification, which is a different skill from reading a construction programme.
 - 02 Someone who can price retrofit optionality. The question is no longer whether a site is good. It is what it costs to convert it to 800 volt distribution and direct to chip liquid cooling, whether the structure and the floor loading permit it, and what that does to the exit multiple.
 - 03 Someone who genuinely understands the interaction between flexible load, behind the meter storage and local grid tariffs, because that is where a meaningful part of the operating margin now sits and it is specific to each market.
-

- 04 And an operating partner who has actually run a data hall, rather than an asset manager who has read reports about them. Availability record and operating cost per megawatt are valuation inputs in a fund owned asset, and a thin operating team shows up as a discount at exit at the worst possible moment.
-

None of these people are in the market in quantity. They are currently employed by operators, hyperscalers, specialist engineering firms and the equipment manufacturers who exhibited this week. Hiring them is difficult, which is why most funds have not done it, which is why the capability gap persists while the technical clock speed increases.

Planning a team that has to change every year

There is a second problem behind the first, and it is one we think is underappreciated.

Workforce planning in most investment firms is an annual exercise tied to a budget cycle. The technology cycle in this sector is now faster than that. A team designed around the capability requirements of 2025 is already behind the requirements of 2027, and will be redesigned at least twice before a current fund reaches the end of its investment period.

Value accrues to the layers with the slowest clock.

This is a legitimate use of artificial intelligence, and one of the few in talent work where the value is real rather than marketed. Skills can be inferred from actual work history rather than job titles. Existing capability can be mapped against a forward requirement and the gap quantified. Scenarios can be modelled across labour mix, location, cost and time to hire without needing a specialist planner for each one. Planning can become continuous rather than annual, which matches the clock speed of the underlying problem.

What it cannot do, and here we would urge caution, is tell you whether the person it has identified can actually be hired in Johor, on your timeline, at that price. That data does not exist in any public system. It exists in the heads of people who are placing those candidates and watching offers get accepted and rejected week by week, which is a deliberate argument for combining the two rather than trusting either alone.

There is a commercial consequence that ties this back to our last paper. If the technology clock accelerates while your revenue model is contractually fixed for fifteen or twenty years, your margin erodes by default. The teams that

keep those models profitable are the ones who can re price, renegotiate and restructure as the underlying asset changes. That is a commercial and technical capability sitting in the same person, and it is rare.

What SCG Partners do about it

We are a Singapore based executive recruitment and talent advisory firm working across infrastructure, energy and the energy transition in Asia Pacific. Most of the regional platforms building in Johor, Batam, Jakarta and Bangkok make their senior appointments here, which is where we are.

For capital providers, the work that matters usually starts before a search. A capability assessment of an investment team, measured against what this sector is about to require rather than what it required when the team was assembled, is a short piece of work that most firms have never commissioned. A workforce and delivery capacity assessment during diligence establishes whether a target asset can actually be staffed in its market on its timeline, which is a risk almost nobody diligences and which we can quantify because we recruit into it. A talent risk review across a portfolio identifies where a platform is exposed if two or three individuals leave.

Then there is the search itself, for the operating partners, technical leads and commercially structured people described above. That is our core business and it is genuinely difficult work, which is the reason to use somebody who knows where these people sit.

One thing we do not do is advise on capital raising or financing structures. That is regulated activity in Singapore and it belongs with licensed advisers.

The industry spent this week in Singapore demonstrating, at considerable expense across three floors, that it is capable of extraordinary technical advance. The open question is whether the capital funding it can change at the same speed. That will not be decided by better models. It will be decided by who is in the room when the decisions are made.

Talk to SCG Partners · satchfieldconsulting.com

Sources. Event detail from Tech Week Singapore and Data Centre World Asia, 29 and 30 September 2026, and from the published conference programme and exhibitor announcements. Technical material from NVIDIA, ABB, Eaton, Delta Electronics, Midea Building Technologies and Nidec. Market and financial data from Goldman Sachs Research, Morgan Stanley, UBP, Cambridge Associates, the International Monetary Fund, S&P Global Ratings, RMI, and published hyperscaler reporting. Policy statements from the Ministry of Digital Development and Information, Singapore. Conditions were current on 2 October 2026 and continue to evolve.

Images. Marina Bay Sands by Calvin Seng. Switchgear panels by Shameer Vayalakkad Hydrose. Insulated pipework by Vladimir Srajber. All via Pexels.