



Five things that could stop the data centre boom and why we're not calling it a bust

Key points

- ▶ The data centre boom in Australia and abroad faces five main risks: electricity and grid constraints, opaque debt financing, weaker-than-expected future revenues, community opposition, and concentrated supply chains.
- ▶ We do not see a bust as the base case. Improving model efficiency, broader demand for computing, staged development, investment in energy infrastructure, and interdependency between countries should help negate some of these risks.
- ▶ The AI capex cycle has further to run but is likely to plateau around 2028 as the technology matures and revenues become more reliable.
- ▶ Investors should diversify beyond hyperscalers and favour flexible, reusable data centre assets.
- ▶ Australia can benefit from data centres through grid investment, mining, engineering and data sovereignty but only with supportive regulation and community backing.


The basics of data centres

Data centres are nothing new: they are buildings that house computing infrastructure. Before the AI boom, they mainly handled cloud storage, website hosting, remote computing and networking. It perhaps sounds boring because it is! However, centralising computing has made IT infrastructure cheaper and easier to scale for most companies on earth. Overall, these traditional workloads still account for roughly 50% of global data centre capacity and will remain an important part of the mix.

What has changed since late 2022 is the boom in larger AI data centres, which comes with much investor hype around them. Unlike traditional facilities built around CPUs and relatively predictable workloads, AI training relies heavily on GPUs (aka NVIDIA's bread and butter), which use more electricity and generate more heat. McKinsey estimates that more than half of data centre workloads will be AI-related by 2030. These facilities could either be colocation sites, where clients rent space from operators such as NextDC, AirTrunk or Goodman; build-to-suit projects for a single customer; or centres owned outright by hyperscalers such as Amazon, Microsoft and Google.

Data centres need cheap power, land, capital, a supportive regulatory environment and, for cloud facilities, proximity to customers. Australia has had plenty of these advantages, which explained the five-fold rise in data centre buildings, server racks and processing equipment capex since 2020, driving much of the growth in business investment for the country. Meanwhile, hyperscaler capex plans abroad point to strong growth in data centre investment through to 2030. So what could slow the boom? Below are five risks to data centre development, as well as five reasons we think the capex cycle will plateau rather than collapse, without taking broader share and bond markets down with it.

Big Tech's capital expenditure




Source: Bloomberg, AMP

1. Power demand and grid bottlenecks

Once operational, power is one of the biggest operating costs for data centres, which account for around 40-60% of most data centre cost bases. Since 2020, electricity consumption by data centres has doubled and now accounts for close to 2% of world electricity demand. The International Energy Agency still projects it to double again by 2033. Meanwhile the AEMO estimates that data centres currently account for a similar 2% of Australia's electricity use, but their share could rise to around 10% over the next decade.

Electricity consumption by data centres, 2020-2035



Source: IEA, AMP

As a result, limited generation and transmission capacity can slow how quickly data centres are built and connected. Their large, reliable power needs mean grid-connection waits range from about three years in Dallas and Sydney to seven years in Paris, London and Frankfurt, and up to 10 years in Amsterdam and Tokyo, according to JLL.

If new data centres rely mainly on existing supply, they could compete with households and energy-intensive industries such as mining and

manufacturing, pushing up electricity prices and effectively limiting how quickly the sector can grow without weighing on the broader economy. Their round-the-clock demand (even when wind and solar output is low) could also keep coal- and gas-fired generation running for longer, delaying the shift to renewables.

Why we're not worried


Long-term forecasts may underestimate how quickly chips, models and data centre operations are becoming more efficient. According to the [IEA](#), energy use per AI task has fallen by around 90% a year in recent years; with recent rises in electricity consumption reflecting wider adoption and the training of more advanced models instead. A simple text query now uses less electricity than running a TV for the same period, while developers such as DeepSeek and Moonshot AI's Kimi have achieved comparable capabilities with less computing power and older chips. Total electricity demand will still rise as the economy grows and AI adoption expands, but improving efficiency should slow consumption growth as the technology matures.

In addition, large loads can improve grid economics. Research by the [Berkeley Lab](#) showed that if there is spare capacity in the grid (e.g. midday in Australia when too much solar energy means wholesale electricity prices drop below zero), having large, utilised loads means the large fixed costs are spread out over more customers.

Finally, capital flowing into data centres could help finance new renewable generation, R&D and grid connections. Long-term power purchase agreements provide predictable demand, making new wind, solar and battery projects easier to finance, much like an anchor tenant in a shopping centre. Some governments also require data centres to support new clean-energy capacity: For example, in Ireland, new facilities must provide generation or storage and meet 80% of demand with additional renewables; in Singapore, operators must source at least 50% of new capacity from green energy. In short, data centre investment could expand energy infrastructure and raise the power "ceiling" on how many facilities can operate at once.

2. Financing and refinancing risk

Data centres are increasingly financed through a mix of on-balance-sheet bonds and project finance via special purpose vehicles (SPVs), rather than equity or operating cash flow alone. This is because while free cash flow across the five hyperscalers has grown by around 26% a year over the past three years, capex has risen much faster, by 55%, 72% and 96% yoy between 2024 and 2026. As a result, debt markets are filling the shortfall.




Source: Bloomberg, AMP

Corporate bonds remain the main funding source, with issuance by the top five companies exceeding US\$100 billion in each of the past two years. Most debt is long-dated to match multi-year reconstruction timelines. On-balance-sheet debt has grown by 35% a year over the past two years for the largest hyperscalers. Yet that roughly US\$200bn increase covers only a quarter of this year's planned capex, and an even smaller share of next year's plans.


That creates a need for other funding. Some lower-rated or unrated firms such as OpenAI and Anthropic may rely on equity partnerships with investors including SoftBank, Macquarie or GIC. In addition, most firms also use off balance sheet arrangements built around long-term commitments.

These can include commitments to builders, electricity providers and chipmakers, or long-term leases with a Special Purpose Vehicle (SPV) that owns the data centre. Because these future payments provide predictable cash flows, the SPV can borrow against them. Under accounting rules, they are permitted to be "off balance sheet" unlike conventional debt, helping companies preserve cash, save on depreciation expenses, lower leverage ratios, preserve their credit ratings and access to bond markets, but also making their true downside risks harder to assess. Below is a diagram of the cash flows of a typical arrangement.



Source: AMP


These obligations (in blue) are generally larger and growing faster than on balance sheet debt (in red), as the chart below shows. Disclosure is also less standardised (I had to dig through the footnotes of 10-Ks to find these numbers).



Source: Bloomberg, WSJ, AMP

The risks to the boom in debt financing are threefold. First, opaque and interconnected contracts make it harder for investors to judge leverage, cash flows and concentration risk in their portfolio. Second, higher rates or wider credit spreads could increase refinancing pressure as they come due in about five years. For now, investment-grade corporate spreads over Treasuries have narrowed, keeping debt financing relatively affordable. Though, further Fed rate hikes could slow the economy and widen spreads, particularly as weaker issuers enter the market (for example, Oracle is rated BBB with a negative outlook). Third, this relatively new asset class lacks historical performance and has not yet been tested through a full economic cycle.

US Corporate BBB to Treasury 10 Year Spread



Source: Bloomberg, AMP

Why we're not worried


With any new asset class, investors need time and periods of stress to understand and price them properly. Some weak deals will inevitably emerge, but they should help expose possible risks and standardise issuance and contract terms. And if a major hyperscaler's cash flow falters, the SPV will likely be able to re-lease the space to a more successful player (although this protection weakens in a broader downturn). Finally, hyperscaler debt represents only around 3.5% of the outstanding US investment-grade market for now, so there is little evidence so far that wider hyperscaler spreads are spilling into the broader market.

3. Demand risk

Related to this, a risk for future debt repayment as well as the broader sharemarket is whether AI users consume enough computing power and pay enough for the capacity being built. Currently, share prices are being justified by lofty earnings growth forecasts: hyperscalers' revenues are projected to grow around 21%pa for each of the next three years before tapering around 10%pa by 2029, leaving 2029 revenue roughly double its 2025 level.

If revenue falls short, companies may struggle to meet long-term lease commitments and share prices need to be adjusted downward versus earnings.

Big Tech's revenue estimate



Source: Bloomberg, AMP

On top of this, data centre construction and fit-out can also take years, creating a risk that revenue arrives too slowly or flows to more efficient operators (such as Chinese model developers or new competitors, rather than the hyperscalers locked into the leases). At the same time, chips could become obsolete quickly, further reducing a facility's value or the computing power needed per AI task, leaving some capacity underused.


Why we're not worried

For now, strong balance sheets, positive operating cash flow, resilient cloud businesses and signed customer contracts should help the major hyperscalers meet their financing commitments in the next few years.

In the longer term though, while overcapacity could still hit tech valuations, the technology behind data centres is here to stay. Hyperscalers cannot indefinitely sustain the double-digit annual share price gains as we have seen since 2023, but we expect that gains could broaden out to other industries in the next decade and the AI boom will continue beyond any sharemarket pullbacks.

In particular, AI will mature further, and cheaper computing (see the chart below) will actually create more demand from end users as well as preserve the revenues for overall data centre services. It's the classic economics situation where lower prices increase quantity sold. Cheaper and more accessible computing services will lead to wider adoption and automation in more industries like health, agriculture, industrial, airlines, resources, and professional services.

Silicon Data LLM Token Expenditure Index



Source: Macrobond, AMP

And remember that data centre demand is broader than AI training, it also includes cloud, storage, inference, cybersecurity and enterprise workloads which are still growing quickly (given a growing economy) and could use lower end chips.

This is similar to the internet and fibre optic wave of the 2000s – where there was indeed some overbuilding. However, the downturn helped filter out weaker operators of the facilities, and the excess capacity continued to be used in addition to benefits for the broader economy despite the sharemarket pullbacks.

Finally, big data centre projects are also usually developed in stages, meaning operators can delay equipment installation rather than abandoning the entire site.

4. NIMBYism

NIMBYism (Not In My Back Yard) is a major pushback to data centres in democratic countries, and it's fast becoming a major constraint on data centre development. In 2Q 2026, US-based Data Center Watch reported that local opposition disrupted 45 projects worth US\$68bn. In the upcoming US midterms election this November, 34 data centre-related measures will appear on the ballots, while local governments in 40 states have active bans or moratoriums and further restrictions are being promised to voters on the campaign trail (S&P). In Sydney, Lane Cove residents successfully opposed Goodman Group's proposed 24-hour, 90MW facility.

In general, NIMBYs' concerns include noise, competition for scarce land, electricity, water, and environmental impacts, and at the same time, the "communities" would have to bear all these costs while financial benefits mostly flow to "big corporations". Some of the protests are also related to the broader fear that AI will displace jobs.

Growing resistance could slow approvals, raise compliance costs and reduce the supply of suitable sites, particularly when Australia's construction productivity is already weak.

Why we're not worried

Many concerns can be mitigated through better design and regulation (see point #1). In addition, we think there could be more to be done to educate the community on the real costs and benefits.

Sydney already has large facilities close to homes, including NextDC's 93MW Artarmon complex, which is larger than the rejected Lane Cove proposal and yet it has attracted relatively little opposition. Noise can be engineered and insulated, as demonstrated by Sydney data centres already operating directly across from apartment buildings. Water use is a genuine concern, as conventional estimates suggest a facility this size could use up to 6 million litres a day – almost 10% of Australia's total daily water use! However, that's only 0.013% of daily water use for Australia; and these figures often assume older, less efficient systems, open-loop cooling and continuous 24/7 peak usage on hot days. In reality, many hybrid designs could use no water under certain weather conditions. Reclaimed water and more efficient systems can further reduce demand with the right regulations; for example, Victoria requires new data centres to use recycled or non-drinking water, while Google uses reclaimed non-potable water at a quarter of its campuses.

The economic benefits also extend beyond construction. Data centres pay land and income taxes and support round-the-clock jobs in engineering, maintenance, security, cooling and electrical systems. The ongoing energy needs also mean demand for installation and maintenance staff for the grid. Over time, all of this infrastructure adds to Australia's capital stock, productivity and capacity, even for other industries who will eventually need more data centre space.


Related to this, the Balassa–Samuelson effect suggests that productivity gains in one part of the economy can lift wages elsewhere, as industries compete for the same workers. Australian wages, for example, are higher than those in emerging markets partly because Australia is more productive in industries such as mining, not because Australians are necessarily more productive in services such as hairdressing or teaching, where productivity differs less across countries.

In short, communities need to see the benefits, even if some take time to emerge. Australia has ample space, relatively cheap renewable energy, and a need to lift weak business investment and preserve data sovereignty. It would be a waste to let NIMBYism unnecessarily constrain the boom.

5. Supply chain and geopolitical risks

Supply chain disruption has become a defining risk since COVID, and AI data centres are particularly exposed because advanced chips depend on a small number of companies and countries.

The chart below shows NVIDIA supplies around 75% of GPUs used for AI training. Meanwhile TSMC manufactures NVIDIA chips as well as advanced designs from AMD, Intel and hyperscalers, giving it a market share above 90%. In the next step in the supply chain, ASML is the sole producer of the advanced lithography machines used to print the smallest chips.



Source: siliconanalysis.com, AMP

Geographic concentration further compounds the supply-chain risk. TSMC operates one 12-inch fabrication plant in Arizona, but all its advanced facilities producing chips at 6 nanometres and below are in Taiwan, a continuing source of tension between the US and China. ASML has a relatively diverse manufacturing footprint across Europe, the US and East Asia, but assembles its most advanced machines only at its Dutch

headquarters. Even their competitors are concentrated in South Korea, Japan, China and a handful of US sites.

This leaves the data centre rollout vulnerable to geopolitical tensions, export restrictions, natural disasters and even minor manufacturing delays like those seen during Covid. Australia faces even more risks because we are a net importer of both construction materials, including steel and refined copper, and fit-out equipment such as server racks.

Why we're not worried

Even a severe supply shock is more likely to delay than stop the boom because 1) most governments are prioritising high-tech manufacturing as a propeller of the economy (e.g. China has doubled its share of chips in the exports mix in the last 1.5 years as the government wants to pivot away from low-tech construction and manufacturing), and 2) the countries involved remain interdependent and have often carved chip exports out of broader trade barriers.

Even as a net importer of materials, Australia is also a beneficiary of the AI supply chain, as we are a net exporter of copper ore, iron ore, critical minerals (including lithium and rare earths), as well as a net energy exporter. This does not only increase national incomes but also gives us leverage in securing downstream products, as seen in the LNG-for-gasoline deal with Singapore during the April 2026 fuel supply crunch.


Implications for investors

Like any new investment, data centres carry plenty of risks, which is partly why investors are being compensated with higher bond yields. Overall, we think the capex cycle has further to run: total US tech equipment investment is now 2.6% of US GDP, below its 2.9% peak in 2000, and is likely to plateau around 2028 as the technology matures, adoption broadens and direct and indirect revenues become more reliable. Even if AI-training capacity is overbuilt by then, the data centre facilities can remain as industrial sites for corporate computing and cloud services.

For longer-term investors, the key is to diversify, given that the beneficiaries of the AI boom are not the hyperscalers but are likely to be the broader supply chain, energy providers, network developers and end users.

Corporate bond, private debt and ABS investors should favour data centre assets that can be repurposed and transferred quickly after a default over highly customised facilities.

For policymakers, including in Australia, the data centre boom provides an opportunity beyond leasing industrial sites or developing frontier AI models. It could also support grid upgrades, growth in mining and engineering, and having the leverage in one of the world's most important assets: data centres. Though, realising these benefits will require having the right tax settings and incentives, community education, sufficient labour and materials, and regulations that support grid investment.



Source: Macrobond, AMP