NOMURA

Global Markets Research
26 October 2026

Asia AI Semi & Server

EQUITY: TECHNOLOGY

AI data center build-up tracker

More upside from our prior update, but with rising funding costs: risk/reward turning more neutral

Our new data center build-up tracker suggests further upside; rising funding costs drive investor concerns, but might not slow down AI infra investments; AI infra build-up would surge into 2027F, but there is no clarity on 2028 yet

We initiated our global new data center build-up tracking in Dec-2025 to have a leading indicator beyond Asia tech semi and hardware supply chain where data points, in our view, are "not useful", and "only positive or super positive" (*Dec 2025 anchor report*). Our continued tracking showed an upside since then (*March 2026 update*; *June 2026 anchor report*) — which is one (not the only one) reason for our high convictions in accumulating tech stocks on any fall from current levels in the past one year. During the period, TAIEX also touched new highs, and reached nearly TWD50k in Sept-2026.

In this tracker update, we see further upside to our *prior update in June* (demand rising to 26/40/28 GW in 2026/27/28F, from 27/32/23 GW in prior update; *Fig. 11*). Interestingly, if we include the comments from SpaceX (SPCX US, Not rated) and Broadcom (AVGO US, Not rated), the upside would be meaningfully higher for both 2027 and 2028 (demand rising to 49/46 GW in 2027/28, which include some of our subjective judgements). See *What could be more?* section. We maintain our view that Asia AI tech semi and hardware supply earnings would trend higher into 2027F given the ongoing shortage (further upside for demand, while new capacities from greenfield investments from end-2025 are not yet to start contributing through 2027).

That said, in terms of the cycle (and on investment stance, accordingly), we now turn less aggressive as risk/reward appear turning increasingly neutral to us for following reasons: **1)** share prices of AI tech stocks (either in Taiwan or the US) have rallied hard into this year (following three super bull years of 2023, 2024 and 2025), before increased volatility from 3Q26 (*Fig. 1* - *Fig. 4*), driven by the positive developments in AI infra build-up and strong consensus earnings upward revisions YTD (*Fig. 5* - *Fig. 6*); **2)** however, compared to a year ago when we recommended accumulating into year-end declines driven by "AI bubble" concerns (*Dec 2025 anchor*), we are now a year closer to the time when new greenfield capacities will be coming out from 2028 (hence, to some extent, it is the later part of that period, in our view); **3)** while our new data center build-up tracker continues to suggest upside (and even more so by including SpaceX and BRCM, suggesting significant numbers from their oral comments during the earnings call, see below sections), whether these numbers are to cover up for the seemingly shrinking funding pool would be debatable; **4)** with our unchanged view from *June* that top US hyperscalers' FCF would be insufficient to support capex upside, our tracking of top US hypercalers' funding costs indeed (unsurprisingly) has also been rising (*Fig. 16*); **5)** as we flagged in our *June 2026 update*, we were hoping to have full 2028 visibility with a stronger 2028 (than 2027) data center build-up demand, but we are yet to reach any conclusion in this quarter's survey.

However, the equity market is efficient, too. The valuations of SOX, Nasdaq and TAIEX have started correcting again these days (*Fig. 1* - *Fig. 4*), on concerns of Anthropic's ARR trend (*Fig. 8*) along with potential lower IPO valuation (*news*), rising capex funding cost (*Fig. 16*) along with Coreweave's surging CDS (*news*), on top of macro factors such as the US 10-year bond yield breaking a record level at 5% (*chart*) and US midterm elections (on 3 Nov), but not driven by consensus earnings cuts (according to *Fig. 5* - *Fig. 6*) – so that tech valuation overall is not demanding, in our view, which suggests that share prices could also easily rise, triggered by further upward earnings forecast revisions into 2027. Also, **we don't think the AI infra build cycle will be over in 2027F.** Semi/hardware upcycles usually do not end when there is still shortage, and those

Research Analysts

Asia Technology

Aaron Jeng, CFA - NITB
aaron.jeng@nomura.com
+886(2) 21769962

Anne Lee, CFA - NITB
anne.lee@nomura.com
+886(2) 21769966

Vivian Yang - NITB
vivian.yang@nomura.com
+886(2) 21769970

Eric Chen, CFA - NITB
eric.chen@nomura.com
+886(2) 21769965

Donnie Teng - NIHK
donnie.teng@nomura.com
+852 2252 1439

Production Complete: 2026-10-27 00:21 UTC

See Appendix A-1 for analyst certification, important disclosures and the status of non-US analysts.

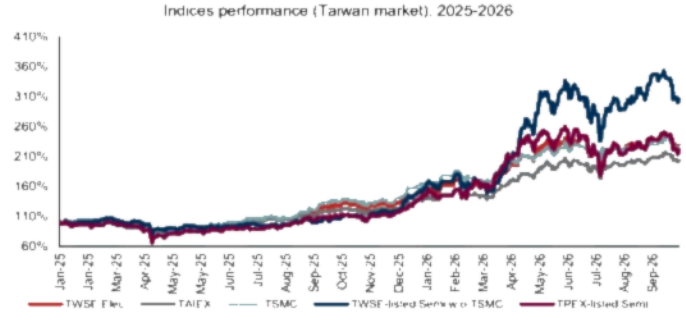
greenfield capacities still won't come online before 2028, in our view.

What's more, the rising funding costs to build AI infra is a reality, but compared with USD30bn ARR per GW (assuming sustainable or even upside) earned by those top token providers/front tier model suppliers (Anthropic, Open AI, Google, xAI, etc), even funding costs at around 10% (vs. their current funding costs at 5-10% level now) might not stop them from borrowing/investing given superior ROIC — which suggests to us that US hyperscalers will likely continue to be positive on their future AI investments (not to mention their consistent messages in the past few quarters that "under-investing AI is riskier than over-investing"). xAI's recent bullish plan to build 20GW new capacity over 2027-28 is one good example, despite credit investors' negative view on Coreweave. The AI infra build demand is still on the rise, but some vendors with stronger execution and credibility might outrun those weaker players, e.g. neoclouds.

Net net, the AI trade will likely re-emerge once recent credit market concerns stabilize and/or OpenAI/Anthropic ARR uptrend re-accelerates; but, we won't expect the return to be as substantial/sustainable as they were over the past two-three years (SOX 2024: 19%; 2025: 42%; 2026 YTD: 52% / TAIEX: 2024: 28%; 2025: 26%; 2026 YTD 59%)**.** The upside catalysts would be: 1) further ARR upside surprise of Anthropic/Open AI after their significant growth over the past three years (*Fig. 8*). These two companies might explain 1/3 incremental AI compute demand over the next two years, on our estimates, on announced incremental GW deployments. They are not the only buyers in AI compute but their ARR trend would have a significant impact on the Asia semi and hardware supply chain (as CSPs and sovereign AI have not disclosed details); 2) a positive ARR trend could also help ease AI infra build-up funding costs, or if the yield curve of US government bond fall when inflation concerns ease; 3) hyperscalers' comments on AI ROIC and capex outlook in 2027-28; the top 4 US hyperscalers may account for another 1/3 of incremental AI compute demand (*Fig. 13*); and 4) further upward revisions to consensus earnings estimates for tech supply chain companies would still be a positive factor once the dust settles. We believe the ARR trend (thus ROIC) might be the most important factor among all – which is subject to AI monetization capability and if more killer applications are launched – as it would affect credit providers' confidence, hyperscalers' aggressiveness in building AI infra, and the demand outlook of Asia's tech semi and hardware supply chain.

Fig. 1: Indices' performance (Taiwan market): 2025-2026 YTD

TAIEX broke high in Sep 2026 before correction in Oct



Note: Shadowed area indicates earnings period
Source: Bloomberg Finance L.P., Nomura research

Fig. 2: Indices' performance (US market): 2025-2026 YTD

SOX peaked in July 2026 with Sep rebound weaker than TAIEX



Note: Shadowed area indicates earnings period
Source: Bloomberg Finance L.P., Nomura research

Fig. 3: Taiwan large-cap semi P/E (excluding TSMC)

Taiwan large cap semi, excluding TSMC, trades at 14x now, vs. 10 year range of roughly 10-20x



Source: Bloomberg Finance L.P., Nomura research

Fig. 4: SOX P/E

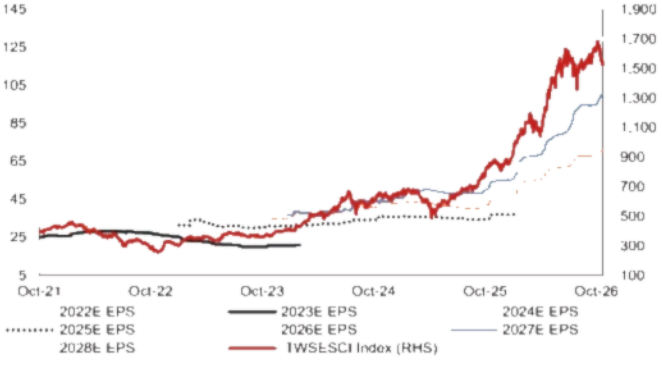
SOX trading at 18x now, vs. 10 years range of roughly 13-27x



Source: Bloomberg Finance L.P., Nomura research

Fig. 5: TWSESCI index vs Bloomberg consensus EPS revisions

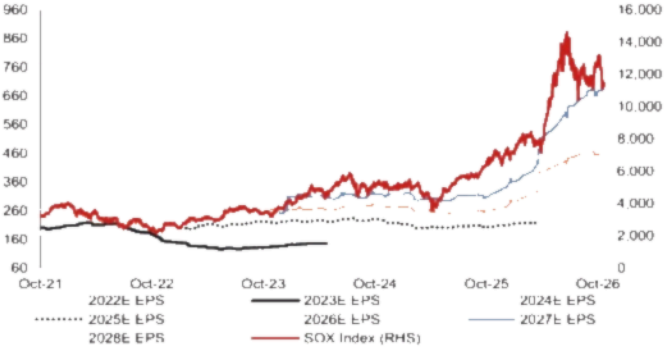
Earnings revision still on the rise



Source: Bloomberg Finance L.P., Nomura research

Fig. 6: SOX index vs Bloomberg consensus EPS revisions

Earnings revision still on the rise



Source: Bloomberg Finance L.P., Nomura research

Fig. 7: Cisco share price vs. internet capex over 1996-2002

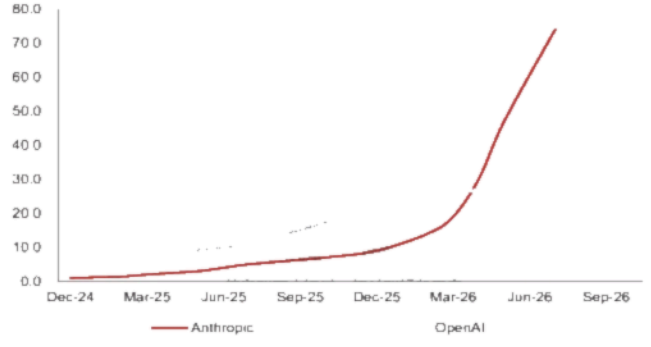
Cisco share price peaked in March 2020, along with peak capex in the same year



Source: Bloomberg Finance LP, US Census Bureau, Nomura research

Fig. 8: ARR trend of Open AI and Anthropic

Their ARRs have recorded growth in the past three years. Their future ARR trend would affect Asia semi/hardware stocks meaningfully, in our view (USDbn)



Note: Anthropic reports gross ARR while OpenAI reports net ARR.
Source: Nomura research

We see growing GW-deployment numbers

We still observe GW-deploy numbers going up, with increased visibility

Since *June*, our tracker on global data center project buildouts continues to see an upward trajectory, despite many positive and negative developments around them. Key observations and dynamics are below:

1) Updates on global data center project tracker and implied CoWoS demand

Based on our latest tracking, incorporating project delays, cancellations and additions, we now have 51 GW-scale data center projects in our tracking base. Our latest calculation indicates 26GW/40GW/28GW incremental deployment in 2026F/2027F/2028F, from 27/32/23 GW in *June update*. We do see more project roll-outs with extended visibility (naturally), and SpaceX/xAI's's aggressive expansion also provides solid support to our numbers (see below and *report*). We estimate 26GW/40GW/28GW could translate into the demand for 5-8mn AI chips.

By geography, the US has captured nearly half of the global pipeline at 24 projects. Beyond US, five projects are structured as multi-region (e.g., worldwide) rollouts, anchored by OpenAI's multi-vendor silicon partnerships and CoreWeave's 5 GW master plan with nVidia. Outside North America, buildouts are heavily clustered in nations aggressively pursuing domestic AI capacity, led by South Korea, France, India, and the Middle East.

By project builders/owners/initiatives and traditional hyperscalers (top-4 CSPs) account for 31% of the list, led by Meta and Google, followed by AWS and Microsoft. Agile Neoclouds and crypto-miners-turned-AI-operators represent the second pillar at 26%, with Fluidstack, IREN, Hut 8, and CoreWeave (1 massive 5 GW global deal) driving expansion. The remaining projects are spearheaded by regional telecoms, conglomerates, and state entities, frontier AI labs like OpenAI and xAI, and pure-play infrastructure developers with tenants.

From another angle, Sovereign AI initiatives account for 30% (15 projects) within the 51 samples. Rather than operating in isolation, sovereign mandates frequently co-opt commercial operators.

Beyond the 51 GW-scale projects, our collection of 50+ smaller projects (less than 1GW) could contribute additional 16GW. Besides, simply applying 100MW each for other projects within our project collection but without compute capacity disclosure (110+ projects), we estimate more than 10GW potentially comes from these projects without details. Combined, around 25-30GW hardware demand could be added, depending on their deployment period, equivalent to 5mn+ Rubin GPUs, or ~580k CoWoS demand.

2) CSP's activities over the past several months

Over the past three months, we do see CSP continues investing around the globe; key announcements include – Meta *announced* in Jul-26 that it is expanding Hyperion in Louisiana to 5 GW and more than USD50 billion. Also in Jul-26, Meta *announced* a 1GW data center in Sturgeon County, Alberta, its first in Canada. Google *announced* on 9 Sept that it will invest EUR13bn in Finland, building new data centers with initial construction phase across 2027 and 2028.

In our last update in June (*report*), we mentioned that Crusoe (unlisted) *announced* on 10 June that it would cease the expansion of Project Jade on a client request, a 1.8GW (up to 10GW) project, and we removed the project from our calculation. However, a Laramie County memo on 8 Jul identified Google, through Jupiter Star Holdings, as the company behind the renamed Project Tembo, and we thus add it back (*news*).

It was *reported* in Oct-26 that Microsoft ceased one computing deal with a neocloudin the UK's datacenter, but no official announcement has been made so far. According to our understanding, this is a small-scale project (less than 100MW). However, it was also reported that Microsoft withdrew from the Monarch Compute Campus, while no official confirmation has been made on that as well (*news*).

3) Uncertainties of OpenAI/Stargate — Our previous worries realized

Bloomberg reported in Mar-26 that Oracle and OpenAI scrapped an expansion of their flagship Texas data center after talks dragged over financing and OpenAI's changing needs. Microsoft has since leased the planned expansion capacity (*news*). Another news article reported that OpenAI paused Stargate UK on 9 April, citing high energy costs and

regulatory challenges (*news*). OpenAI withdrew from both Stargate Norway and Stargate UK.

We also witness some developments around Oracle's project build outs: Regulatory and power-supply setbacks have put Oracle's delivery targets for Project Jupiter and Project Lighthouse at risk. These two projects are also OpenAI/Stargate-relevant.

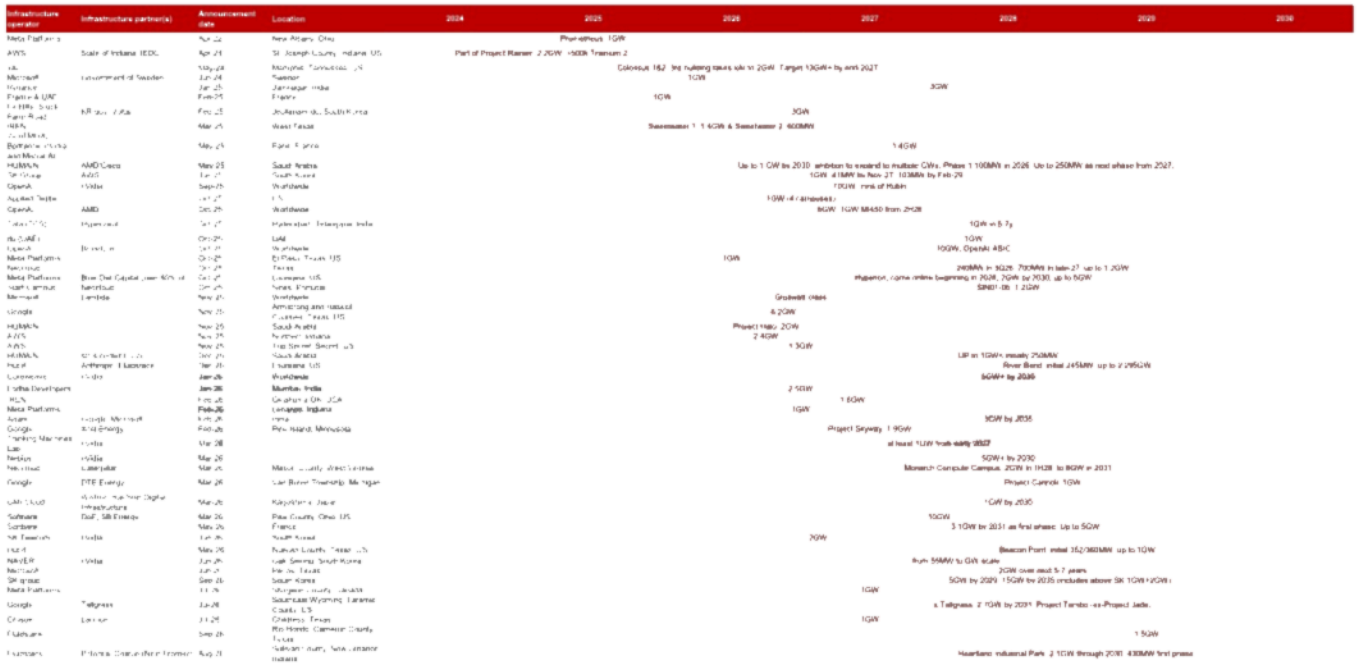
When we revealed our data center tracker in *Dec-25 for the first time*, we noted that one major risk could be OpenAI's execution. We believe these pullbacks somewhat support our concerns. However, note that in our calculation, we set OpenAI's 26GW chip agreement as their GW deployment celling, and will not change the numbers based on the project dynamics underneath (before we can assume that the company will not fulfill the 26GW commitment).

Meanwhile, OpenAI *announced* on 2 July a 3.2 GW campus, Project Camellia, in Effingham County, Georgia, with investment of at least USD20bn. We believe the project is either a power-centric deal, and may somewhat overlap with the 26GW chip agreement, and we thus do not factor in this project into our calculations.

4) Other dynamics

- Firmus *withdrew* its ASX listing on 9 Oct and said it will pursue private funding. In Jun-26 the company *announced* a partnership with nVidia through 2034, anchored by a 360 MW campus in Batam, Indonesia, which DayOne is building, and Firmus will rent. In Sep-26 it *signed* a multi-year OpenAI deal for two Malaysian sites, and Meta *agreed* to take GPU capacity at its Southeast Asia sites. Separately, CDC ended the Project Southgate partnership with about 42 MW deployed (planned 1.6GW originally).
- SpaceX/xAI: We mentioned in *June* that SpaceX could be a new source of computing power. Over the past several months, we see aggressive procurements and deployments from the company (see July *report*). We do not 100% factor in the company's announcement; see details below. For the past two to three years, the AI infrastructure cycle has been driven by nVidia and Broadcom on the silicon side, the hyperscalers as the largest buyers, and OpenAI and Anthropic, together with the neoclouds orbiting them, as the main sources of incremental demand. SpaceX's IPO may be shifting that balance. The company is also placing less like a pure model lab and more like a hybrid buyer and compute supplier, renting Colossus capacity to Anthropic and to Google, which makes its direct competitors into customers. Our view is that SpaceX will build as aggressively as it can while the financing window remains open. We can see their leveraging activities from reporting numbers (*Fig. 21*), as well as buildout plans. From our observations from official announcements, its relationships with the other labs will be a mix of competition and mutual dependence.

Fig. 9: Major data center infrastructure build-out



Source: Company data, Nomura research

Fig. 10: Our back-of-the-envelope calculation on GW deployment trend

We see growing GW deployments into 2026-28F

	2025F	2026F	2027F	2028F	2029F	2030F
Incremental capacity deployment (GW)	4.93	26.25	40.43	28.29	20.10	10.08
- OpenAI	-	3.50	7.50	8.50	6.50	-
- OpenAI (%)	*0%*	*13%*	*19%*	*30%*	*32%*	*0%*
- Top 4 CSPs	2.35	9.00	8.35	2.69	1.80	1.80
Top 4 CSPs (%)	*48%*	*34%*	*21%*	*9%*	*9%*	*18%*
- SpaceX	0.50	1.95	7.30	3.00	-	-
SpaceX (%)	*10%*	*7%*	*18%*	*11%*	*0%*	*0%*
- Others	2.08	11.80	17.28	14.10	11.80	8.28
- Others (%)	*42%*	*45%*	*43%*	*50%*	*59%*	*82%*

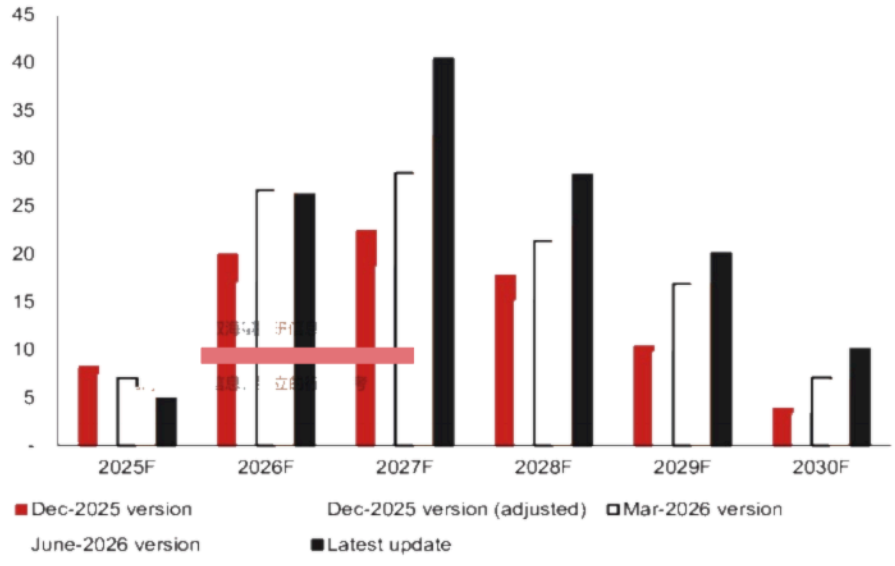
From GW to Chips			
Computing power % as an infra	70%		
Chip (W)	TDP		
GB300	1,400		
VR200	2,300		
Rack power incl. redundancy (kW)		Racks # per GW	Chips # per GW (k)
GB300 NVL72	180	3,895	280
VR NVL72	266	2,630	189

	2025F	2026F	2027F	2028F	2029F	2030F
Incremental capacity deployment (GW)	4.93	26.25	40.43	28.29	20.10	10.08
Assume all are GB300 for 2024-26F:						
Implied chip demand (k)	1,383	7,363				
Implied CoWoS demand (k/16)	86	460				
Assume all are VR for 2027-30F:						
Implied chip demand (k)			7,658	5,357	3,807	1,910
Implied CoWoS demand (k/9)			851	595	423	212

Source: Nomura estimates

Fig. 11: Incremental capacity deployment (GW)

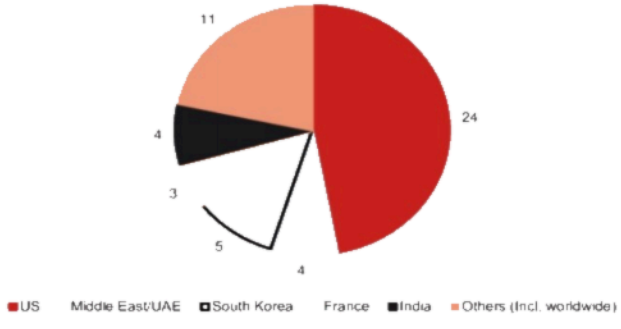
We see more project buildout over past several months



Note: Adjusted: We remove same infrastructure projects for better comparison.
Source: Nomura estimates

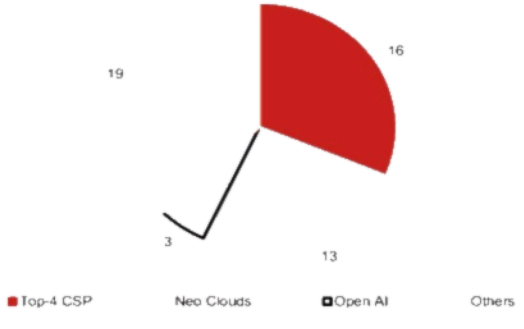
Fig. 12: GW-scale projects by geography

US accounts for nearly half



Source: Nomura research

Fig. 13: GW-scale projects by owners/operators



Source: Nomura research

What could be more?

We highlight few special cases here: SpaceX and Anthropic (Broadcom/AMD). These are potential data center deployment as announced by the companies, while we do not fully factor in them into our tracking in the above section.

For SpaceX, Elon Musk updated on the company's data center expansion plan on 21 Oct 2026 (*transcript*):

- *"We're tracking to be able to bring roughly 200,000 GB300s online per month for every month for the remainder of this year and continuing that pace next year. With next year being primarily Rubin, Vera Rubin compute. At this point, I would say* ***I'm fairly confident of being able to bring an incremental 10 gigawatts of compute online in 2027. That would be in addition to the roughly 3 gigawatts, maybe 2.5 to 3 gigawatts that we'll have online by the end of this year.*** *So I put it this way. I would be surprised if we brought on less than 13 gigawatts of compute by the end of next year."*
- *"And as I mentioned earlier, power cooling and electrical equipment infrastructure is somewhat of a limiting factor. So to make sure that it is not a constraint for our build-out,* ***we're actually going to aim for an incremental 20 gigawatts of power cooling and infrastructure next year.*** *So that the amount of power and cooling that*

we have is far in excess of what is needed for -- to turn the GPUs on. So there's a non-zero possibility. Maybe it's a small possibility at this point, but it's not zero, of having -- of ***SpaceX having over 20 gigawatts of Super Intelligence compute online next year.*** *A small possibility, but it used to be zero, it's not zero anymore."*

Thus, the plan is to deploy incremental 10GW of compute by end-2027, from 2.5-3GW by end-2026 (accumulated 13GW by end-27). In our tracking above, we tentatively assume accumulated compute capacity to reach 10GW by end-2027 (incremental 7+GW in 2027), and additional 3GW in 2028 (accumulated 13GW by end-2028). There could be a 2-3GW gap between the official announcement and our calculation in 2027, out of conservatism. We also don't factor in any potential to reach incremental 20GW in 2027 (10GW more).

Broadcom rolled out customers' plans of chip deployment during its earnings call on 2 Sept 2026 (*transcript*):

- *"Under this agreement, we are planning to deliver multi-tens of billions of dollars of TPUs annually over the next several years. We expect this growing demand in '28 and '29 to be fulfilled through successive generations of the increasingly complex TPUs we are developing today with Google."*
- *"Anthropic, starting with the 1 gigawatt of Ironwood we are deploying in 2026. We expect Anthropic to deploy another 5 gigawatts of TPU version 8i in 2027. And in 2028, we have clear line of sight to deliver another -- an incremental 10 gigawatts"*
- *"For OpenAI, Jalapeno is on track for the planned deployment of 1.3 gigawatts in 2027. Together with OpenAI, we are deep in the development of the next-generation XPU beyond Jalapeno, which is approaching tape-out. In 2028, we have line of sight for OpenAI to deploy over 5 gigawatts of Jalapeno and its successor generation of XPU, which would make OpenAI our second largest XPU customer.*
- *"Between now and the end of 2027, we will be delivering three generations of MTIA accelerators to Meta. Across these three generations, we have line of sight to deploy 3 gigawatts through '28."*

These numbers appear much larger than what we have factored in our tracking in the above section. However, we acknowledge the fact that companies do not necessarily announce projects with GW-disclosure, and within their datacenters, there could still be other sources of accelerators (nVidia/AMD GPUs), other than Broadcom's ASIC.

In our tracking, we still incorporate OpenAI's 10GW agreement (*press*) from 2H26 to 2029, and leave Google's and Meta's projects as what we did before. The one could be more eye-catching or worth-discussed could be Anthropic. We've discussed some agreements between Anthropic and external parties *in our June update*.

According to our observations on public announcements, official press releases, and collaboration announcements in the media, Anthropic's compute sources can be divided into the categories below:

1) Self-build: Through leases powered data center space directly from landlords (L0/L1) and Anthropic arranges its own hardware. TeraWulf's (L0/L1) Kentucky *Justified Data campus* is one example. Anthropic does not construct these buildings; it is the anchor tenant.

2) Self-own?: Through operators such as Fluidstack (L2.5), which sign the lease and builds and runs the site, such as *Hut 8 (L0/L1)'s River Bend (with Google backstop)*. Fluidstack is the tenant of L0/L1 campuses, not Anthropic, but GPU/hardware could potentially be purchased and deployed by Anthropic (*announcement*).

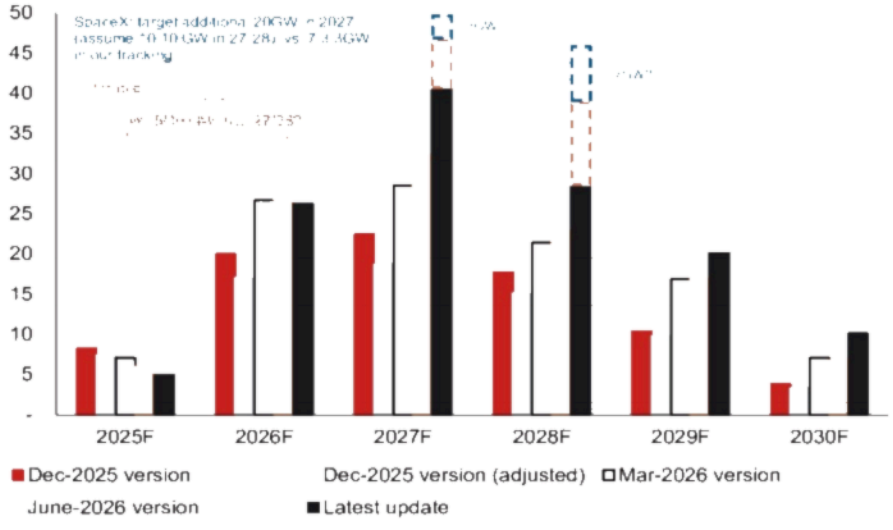
3) Rent external existing capacity as a service from hyperscalers: directly uses AWS, Google Cloud, or Microsoft Azure and pays service fees. The cloud provider owns the hardware.

4) Buys compute from neoclouds and other GPU owners (neoclouds): signs computing capacity leasing deals externally, such as SpaceX/xAI's Colossus, and other companies such as Lambda. The providers own the GPUs and Anthropic pays for usage.

We don't have details about Broadcom's comment on Anthropic's chip deployment, and Anthropic's and AMD's agreement announcement of 2GW MI450 from 2H27 (*press*), because inclusion/exclusion of self-owned or end-usage can change calculations significantly. However, if we only consider large-scale projects (our 51 samples), and only

take into account: **1) self-build** and **2) self-own?**, likely less than 1GW is to be deployed by Anthropic each year in our tracking period. Of course, there could be some small projects or undisclosed projects ongoing, but still implies a significant gap from our current tracking (we projected less than 1GW in 2027-30F vs Broadcom's comment about 1GW/5GW/10GW+ in 2026-2028E, plus AMD's 2GW without a timeline). Considering double-count risks and a lack of evidence and documentation, not to mention execution risks, we don't include these in our calculation base. Of course, these risks can be applied to other projects as well, but not at such a large scale.

Fig. 14: Incremental capacity deployment (GW) – gap between our tracking and some official announcements



Note: Adjusted: We remove same infrastructure projects for better comparison.
Source: Nomura research

Be mindful of capital shortfall

CSP debt expansion has entered a rapid phase

CSPs have entered a phase of rapid debt expansion in 4Q25. As debt financing has scaled to a massive level and credit spreads have widened, CSPs' equity financing and backstop arrangements have boomed starting in 2Q26.

The bond market remains the largest financing channel. CSPs generate strong operating cash flow, carry low leverage, and hold investment-grade credit ratings, enabling them to secure large-scale, long-term, and relatively low-cost funding in the global bond market. In 1H26, Net addition of long-term debt by the five major CSPs has reached USD158bn, a record high.

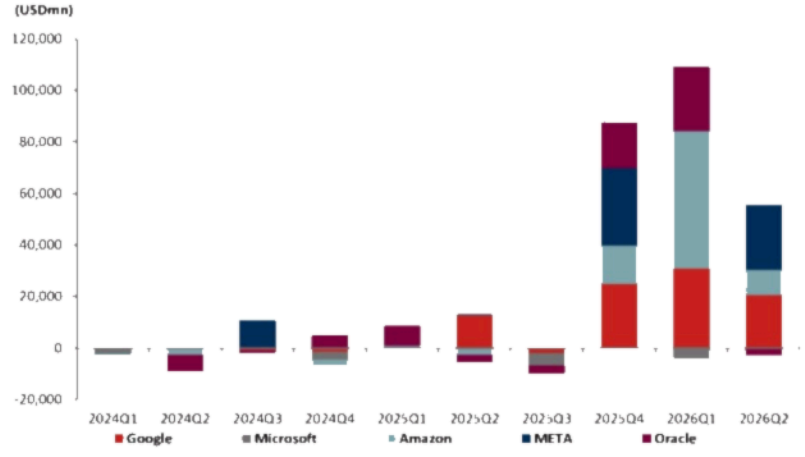
Meanwhile, heavy bond issuance has pushed up interest rates and credit spreads, driving AI infrastructure financing costs higher.

In October 2026, SpaceX is finalizing a debt financing of USD40bn. The riskiest tranche of the new October debt is expected to yield approximately 8.3%. This new October issuance has risen sharply compared to its June bond offering of 6%.

As Big Tech has flooded the bond market and investor confidence in AI monetization has been shaken, CoreWeave's secondary-market bonds is under heavy selling pressure. The high-yield bonds it issued in June with a coupon rate of 9.6% have seen its secondary-market prices plunge, with yields spiking to ~15%.

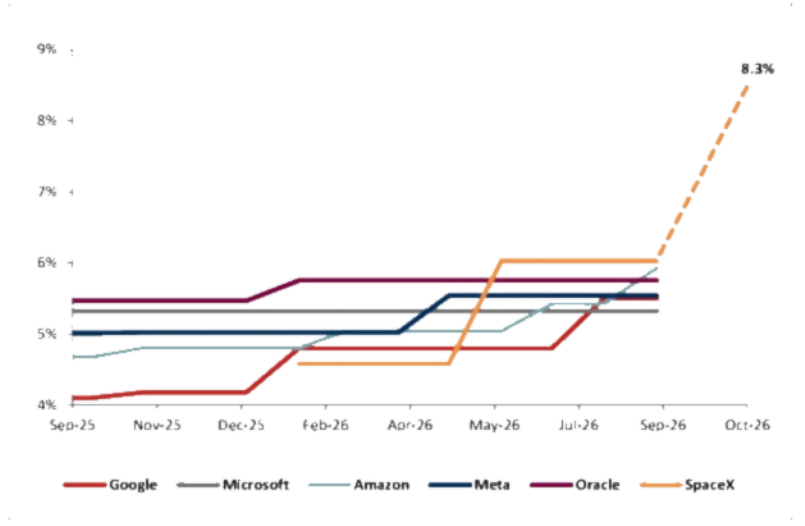
We observe that the currency denomination of debt issuance by CSPs such as Google and Amazon has expanded beyond US dollars into euros, Canadian dollars, Japanese yen, and other currencies – a diversification that helps mitigate part of the issuance pressure.

Fig. 15: Long-term debt financing – net



Source: Company data, Nomura research

Fig. 16: Heavy bond issuance has pushed up interest rates

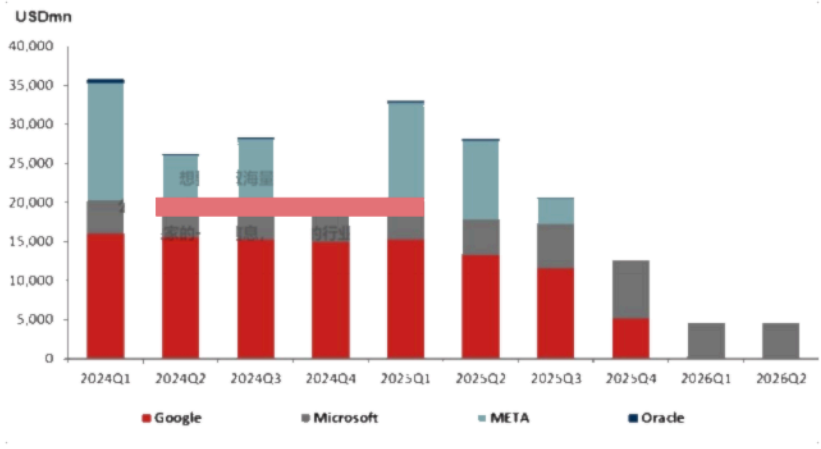


Source: Company data, Nomura research

From share buybacks to equity Issuance

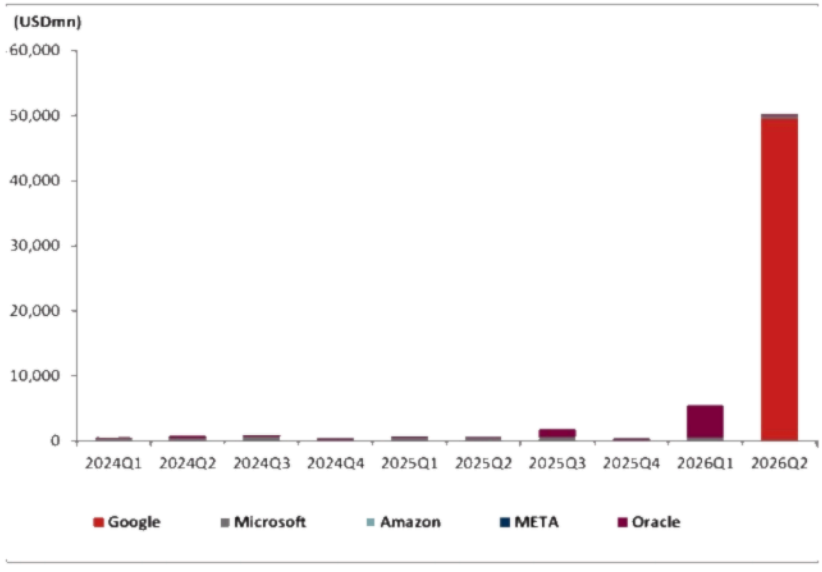
In the past years, CSPs were more inclined to repurchase shares to support EPS. That trend has now reversed: CSPs have paused buybacks and started to rely on equity issuance to raise funds again.

Fig. 17: Equity repurchases



Source: Company data, Nomura research

Fig. 18: Equity financing



Source: Company data, Nomura research

In the latest AI debt shockwave of late 2026, backstop financing has emerged as the core instrument supporting AI-related capital raising.

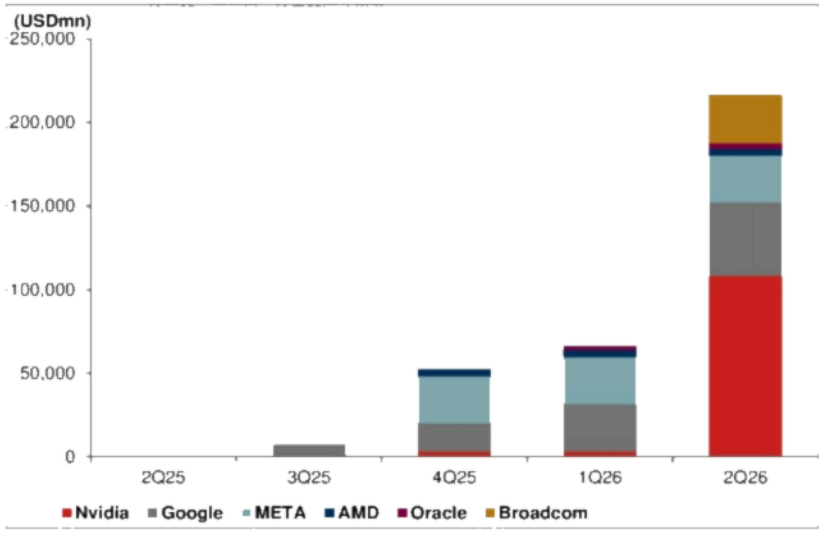
As of 2Q26, the total exposure of key AI companies has reached USD216.7bn, with nVidia alone accounting for USD108.5bn (~50%).

Fig. 19: Disclosed backstop exposure in companies' financial reports

USDmn	2Q25	3Q25	4Q25	1Q26	2Q26
Nvidia		860	3,500	3,500	108,500
Google		6,529	16,940	28,436	43,785
META			28,000	28,000	28,000
AMD			4,100	4,100	4,100
Oracle				2,200	3,300
Broadcom					29,000
Total		**7,389**	**52,540**	**66,236**	**216,685**

Source: Company data, Nomura research

Fig. 20: Backstop exposure



Source: Company data, Nomura research

The rationale behind the massive backstops provided by CSPs and chipmakers, along with their operating mechanics, can be summarized in two key points:

1. Rescuing the overheated debt market. As CSPs and neoclouds issued debt aggressively to fund GPU purchases and data center construction, bond yields and spreads spiked. To reassure Wall Street, chipmakers and CSPs have begun leveraging their own corporate credit to guarantee downstream borrowers. Without big companies' backstops, institutional lenders (e.g., BlackRock, Apollo) would no longer extend credit to lower-rated neoclouds.

2 Securing supply chains and locking in future computing capacity. The current bottleneck is no longer chip availability but grid and space. CSPs are using backstop arrangements to preemptively secure these power and heavy assets and prevent them from being captured by competitors.

How the backstop mechanism works:

1. Lease obligation: nVidia signs long-term take-or-pay agreements with downstream neoclouds or AI startups. Under these agreements, nVidia commits that even if customers fail to materialize or AI application adoption falls short – leaving GPU capacity idle – nVidia will act as a "customer of last resort," leasing back the compute capacity.

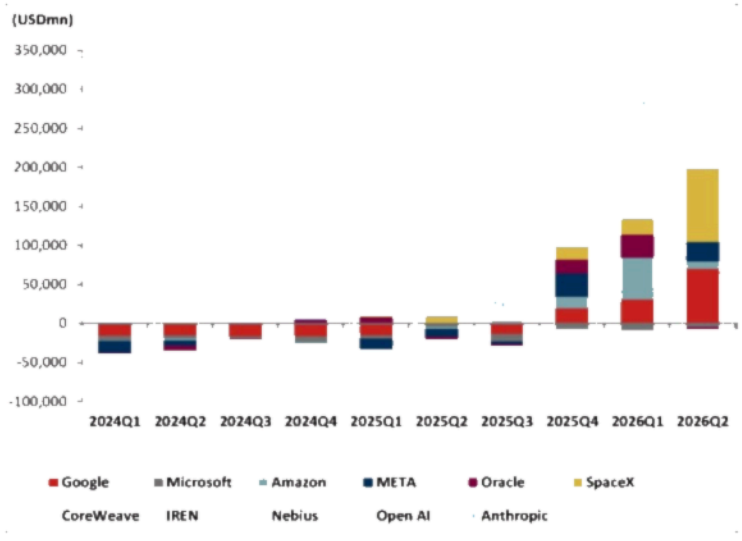
2. Residual Value Guarantee: Chipmakers or CSPs commit to repurchasing used GPU hardware at a contractually agreed "floor price" at the end of the financing period or in the event of borrower default. This effectively mitigates the rapid depreciation and asset impairment risk inherent in GPU hardware.

nVidia provides guarantee for neoclouds to obtain financing, and neoclouds gets the financing and then purchases nVidia's GPU. This "vendor lends to customer, customer buys vendor's product" structure resembles the vendor financing risk Cisco and Lucent aggressively deployed in the lead-up to the 2000 dotcom crash.

The current AI financing cycle does not yet resemble the sub-prime mortgage crisis of 2007, but in terms of capex-to-revenue, it has already exceeded the dotcom bubble's peak.

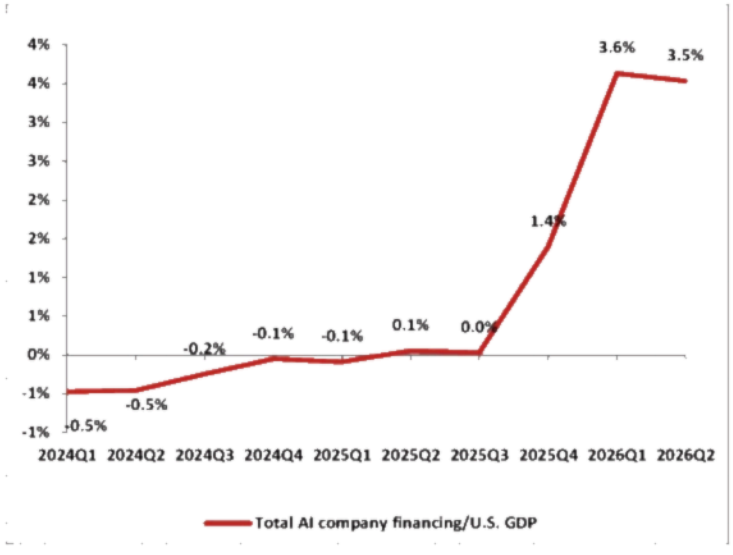
Major AI companies' financing as a percentage of GDP climbed to 3.6% in 1Q26 and 3.5% in 2Q26. This is well below the subprime mortgage origination-to-GDP ratio of 5.4% in 1Q07 (ahead of the housing bubble burst), and far below the total mortgage origination-to-GDP ratio of 21.2% in 1Q07.

Fig. 21: Major AI companies' financing



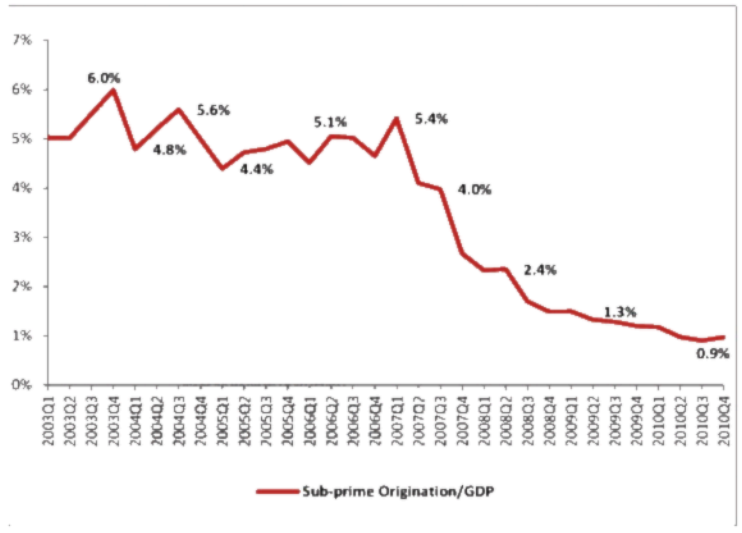
Source: Company data, Nomura research

Fig. 22: Major AI company financing / US GDP



Source: Company data, Nomura research

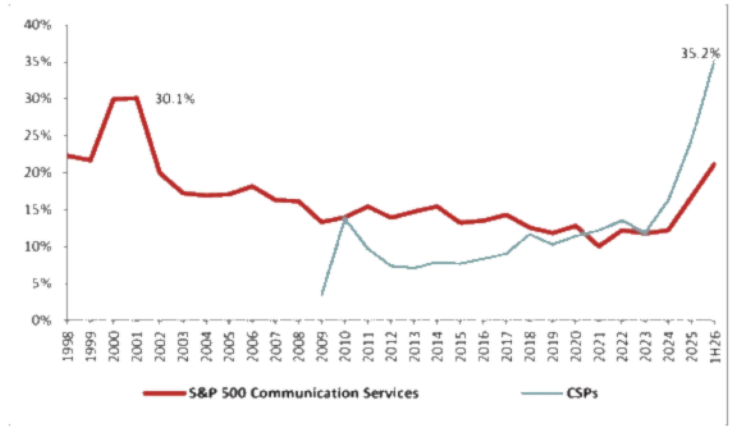
Fig. 23: Subprime origination / GDP



Source: Federal reserve bank of New York, Nomura research

However, CSPs' capex-to-revenue ratio has now surpassed the peak levels recorded during the dotcom bubble.

Fig. 24: Capex/Revenue



Source: Bloomberg Finance L.P., Nomura research

Backstop cases

Here we list some backstop cases we observed from company official announcements, press release, SEC files, and media:

nVidia

nVidia is the most aggressive backstop provider in the AI infrastructure cycle. Its arrangements span two structures, including lease-obligation backstops (rent-back of unsold compute) for neoclouds, and residual-value guarantees (RVGs) for OpenAI large-scale project. Aggregate disclosed exposure was USD108.5bn as of 2Q26.

Case 1 — nVidia × Wall Street financing platform (undisclosed; $500bn financing, 25% of RVG, 3Q26)

nVidia signed an MOU with Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs and KKR to establish a financing platform targeting the mobilization of USD500bn of third-party capital for AI data center construction. The platform provides a dedicated capital pool at attractive rates for nVidia customers to purchase compute equipment and build AI factories; nVidia will backstop up to 25% of the residual value of the chips. While nVidia's exact exposure is undisclosed, a 25% residual-value share of platform-financed hardware could scale into tens of billions of dollars

Case 2 — nVidia × CoreWeave (USD6.3bn, Lease Obligation Backstop, 3Q25)

nVidia guarantees the purchase of any CoreWeave cloud capacity that is not sold to end customers through 2032, in an agreement sized at USD6.3bn. CoreWeave can secure better credit terms.

Case 3 — nVidia × Lambda Labs (USD1.5bn, Lease Obligation Backstop, 3Q25)

Structured similarly to the CoreWeave agreement, nVidia committed to purchase unsold compute capacity from Lambda Labs for up to USD1.5bn.

Case 4 — nVidia × Sharon AI ($4.88bn, Lease Obligation Backstop, 2Q26)

nVidia rents back unsold GPUs at a fixed rate plus a revenue share. The six-year agreement covers up to 40,000 GB300 GPUs for an Australian sovereign cloud buildout, with a backstop size of approximately USD4.88bn.

Case 5 — nVidia × Firmus (undisclosed, Lease Obligation Backstop, 3Q26)

Firmus's Batam, Indonesia project uses the same structure as Sharon AI. The backstop size is undisclosed, but Firmus has projected up to roughly USD30bn of committed offtake over six years.

Case 6 — nVidia × OpenAI — Stargate Ohio ($105bn, Residual Value Guarantee, 3Q26)

For the PORTS-Pike Technology Campus in Pike County, Ohio ("Stargate Ohio"), nVidia agreed to support limited portions of lease and power payments together with a specified residual-value commitment, sized at approximately USD105bn. This is the single largest disclosed AI backstop to date and marks nVidia's move from capacity backstops into direct RVG of GW-scale campuses tied to its flagship customer.

Google

Google pioneered the current wave of lease backstops, routing its support through Fluidstack. Its reported credit-derivative exposure reached USD43.8bn in 2Q26.

Case 7 — Google × Fluidstack → TeraWulf (USD3.2bn, Lease Obligation Backstop, 2Q25)

Google provided a USD3.2bn backstop in support of project-related debt financing for TeraWulf's Lake Mariner (New York) and Abernathy (Texas) campuses, where Fluidstack is the AI cloud tenant. This was the first major lease backstop of the cycle and established the three-party template.

Case 8 — Google × Fluidstack → Cipher Mining (USD1.73bn, Lease Obligation Backstop, 3Q25)

Google backstopped USD1.73bn of Fluidstack's lease obligations at Cipher Mining's Barber Lake site in Texas to support project financing – an initial USD1.4bn covering 168 MW plus a USD333mn expansion.

Case 9 — Google × Fluidstack → TeraWulf, Abernathy expansion (USD1.3bn, Lease Obligation Backstop, 4Q25)

Google added a USD1.3bn backstop for the 168 MW Abernathy expansion of the TeraWulf–Fluidstack joint venture in Texas, bringing its cumulative disclosed support for TeraWulf-related projects to USD4.5bn.

Case 10 — Google × Fluidstack → Hut 8 (undisclosed, Lease Obligation Backstop, 4Q25)

For Hut 8's River Bend campus in Louisiana, Hut 8 issued USD3.25bn of bonds on the strength of Google's backstop. The backstop size was undisclosed.

Case 11 — Google × Fluidstack → Next Frontier JV (undisclosed, Lease Obligation Backstop, 2Q26)

The Next Frontier project in New Lebanon, Sullivan County, Indiana is held by Meridian Arc HoldCo LLC, a JV between Fluidstack Indiana (49.9%) and Frontier Holdings Indiana (50.1%). Google's backstop assumes the lease or pays a termination fee if Fluidstack defaults. On the strength of this support, the JV issued USD5.7bn of five-year senior secured notes.

Case 12 — Google × Anthropic — Nexus DC (undisclosed; USD15bn financing, Lease Obligation Backstop, 3Q26)

For Anthropic's Nexus DC campus in Hubbard, Texas, Morgan Stanley led a USD15bn debt package in which Google guarantees Anthropic's lease and power obligations and takes roughly 20% project equity.

Meta

Meta uses residual-value guarantees (RVGs). RVG thresholds total of $41bn.

Case 13 — Meta × Blue Owl Capital (USD28bn, Residual Value Guarantee, 4Q25)

For the Hyperion campus in Richland Parish, Louisiana, Meta provided Blue Owl Capital a residual-value guarantee with an aggregate threshold of approximately USD28bn, declining over time: if the facility's value falls below the threshold at lease end, Meta covers the shortfall.

Case 14 — Meta × BlackRock (USD13bn, Residual Value Guarantee, 3Q26)

Meta replicated the Hyperion structure with BlackRock for its El Paso, Texas campus, with an RVG aggregate threshold of approximately USD13bn, declining over time. Together, Meta's two RVGs bring its maximum gross exposure to roughly USD41bn by 3Q26.

AMD

AMD's disclosed backstop exposure of USD4.1bn in 2Q26, which is relatively small in dollar terms.

Case 15 — AMD × Crusoe Energy (USD0.3bn, Lease Obligation Backstop, 1Q26)

AMD guaranteed a USD300mn Goldman Sachs loan to Crusoe Energy – collateralized by AMD Instinct chips deployed at Crusoe's Springfield, Ohio campus – and agreed to rent back the chips if Crusoe could not find outside customers.

Broadcom

Broadcom's single disclosed backstop has the second-largest exposure, which is tied to Anthropic's TPU deployment.

Case 16 — Broadcom × Anthropic (USD29bn maximum, Lease Obligation Backstop)

Broadcom partnered with Apollo (lead) and Blackstone in a structure where an investor partner purchases AI racks and leases them to Anthropic, while Broadcom backstops the five-year lease payments. Exposure increases as the AI racks are deployed and decreases as Anthropic pays down its lease obligations, with a maximum exposure of USD29bn, which is an example of a backstop being used to secure a strategic silicon design win.

Oracle

Oracle disclosed guarantees of lessor borrowings of USD2.2bn in 1Q26, which rose to USD3.3bn in 2Q26 before falling to zero in 3Q26.

Fig. 25: AI companies backstop announcement

Provider	Counterparty / Beneficiary	Size ($bn)	Project / Data Center	Backstop Type	Date
NVIDIA	CoreWeave	6.3	CoreWeave (cloud capacity)	Lease Obligation	3Q25
NVIDIA	Lambda Labs	1.5	Lambda	Lease Obligation	3Q25
NVIDIA	Sharon AI	4.88	Sharon AI Australia	Lease Obligation	2Q26
NVIDIA	Firmus	Undisclosed	Batam, Indonesia	Lease Obligation	3Q26
NVIDIA	OpenAI	105	Stargate Ohio (Pike County, OH)	Residual Value Guarantee	3Q26
NVIDIA	Apollo / BlackRock / Blackstone / Brookfield / GS / KKR platform	Undisclosed (Mobilizing $500bn, 25% RVG)	Global AI Data Center Buildout	Residual Value Guarantee	3Q26
Google	Fluidstack → TeraWulf	3.2	Lake Mariner, NY + Abernathy, TX	Lease Obligation	2Q25
Google	Fluidstack → TeraWulf (expansion)	1.3	Abernathy, TX (168MW)	Lease Obligation	4Q25
Google	Fluidstack → Cipher Mining	1.73	Barber Lake, TX	Lease Obligation	3Q25
Google	Fluidstack → Hut 8	Undisclosed	River Bend, LA	Lease Obligation	4Q25
Google	Fluidstack → Next Frontier JV	Undisclosed	New Lebanon, IN	Lease Obligation	2Q26
Google	Anthropic (Nexus DC)	Undisclosed ($15bn debt package)	Hubbard, TX	Lease Obligation	3Q26
Meta	Blue Owl Capital	28	Hyperion, Richland Parish, LA	Residual Value Guarantee	4Q25
Meta	BlackRock	13	El Paso, TX	Residual Value Guarantee	3Q26
AMD	Crusoe Energy	0.3	Springfield, OH campus	Lease Obligation	1Q26
Broadcom	Anthropic	29	N.A.	Lease Obligation	2Q26

Source: Company data, Nomura research

