Nvidia’s Backstop Universe – Heads I Win, Tails Who Loses?

Nvidia’s Backstop Universe – Heads I Win, Tails Who Loses? The $11T AI Buildout, Nvidia’s Backstop Economics, and the Limits of Nvidia’s Balance Sheet

SEP 11 https://newsletter.semianalysis.com/p/nvidias-backstop-universe-heads-i

View in browser

Nvidia’s Backstop Universe – Heads I Win, Tails Who Loses?

The $11T AI Buildout, Nvidia’s Backstop Economics, and the Limits of Nvidia’s Balance Sheet

DANIEL NISHBALL, OLIVER KENNON, AND TERENCE ONG

SEP 11 • PAID

QQ:1413221435



READ IN APP

Follow the money behind the AI buildout with our AI Compute, Capital and Markets Model - understand who is funding the expansion, how deals are structured, and where the risks sit. Contact our team here to find out more.

We’re also hiring for SemiAnalysis’s Compute, Capital & Markets team. We have four openings across New York and Singapore:

- *Senior Credit Markets Specialist - New York: 5-7 years in credit sales, sales trading, or trading, with strong pricing knowledge and investor relationships.*
- *Credit Analyst - New York: 5-7 years in credit analysis, ideally covering high-yield, distressed, or complex debt.*
- *Two Junior Credit Analysts - New York and Singapore: 0-3 years of experience, with strong financial modeling skills. Recent graduates are welcome; prior specialization in credit, semiconductors, or AI is not required.*

QQ:1413221435

If you want to combine financial analysis with a detailed understanding of the technology and economics behind AI infrastructure, we’d like to hear from you -

apply here.

An expanded view of Nvidia's Balance Sheet

As of 2Q F1/27 (second quarter of fiscal year ended Jan 2027) Nvidia in its 10-Q disclosed $530B of gross off-balance sheet guarantees across six different line items – a jump from $184B of obligations disclosed in the prior quarter due mainly to moves in three items:

1. Supply and capacity commitments went from $119B to $279B, primarily memory as explained by the CFO, with 96% of these supply commitments due by F1/29.
2. Guarantees and LPS Guarantees went from $3.5B to $108.5B on SB Energy's PORTS-Pike campus in Ohio, representing 4.25 GW leased to OpenAI for twenty years with Nvidia guaranteeing the land, power and shell.
3. Two line items appeared for the first time: $36B of AI cloud agreements, the take-or-pay floors under Neocloud capacity, and $20B of datacenter leases Nvidia has signed as a tenant but expects to reassign to third parties.

QQ:1413221435

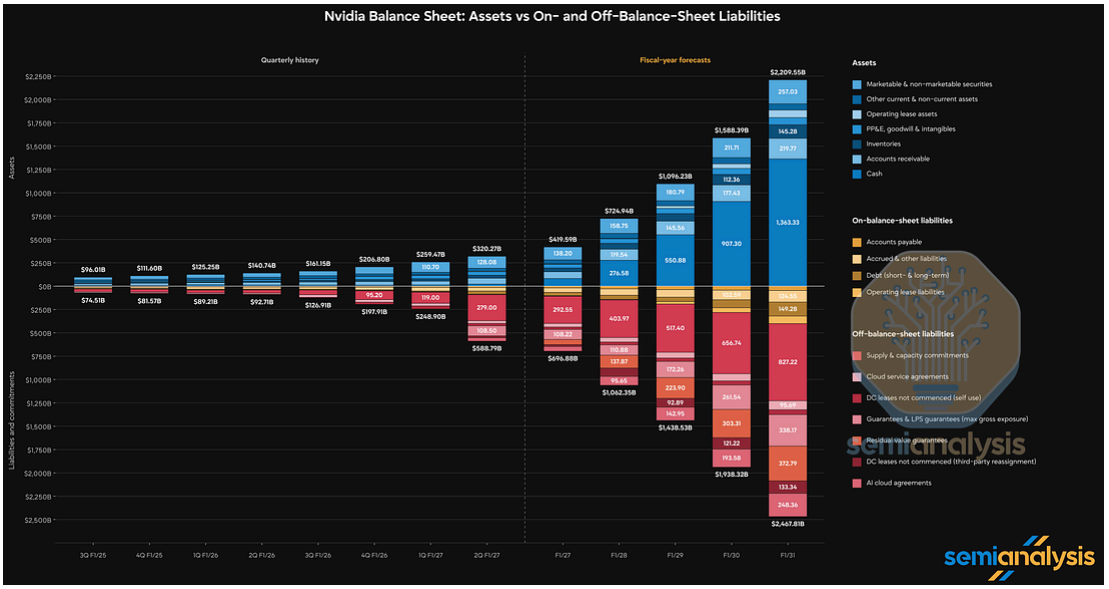
The $530B of commitments dwarf on-balance sheet liabilities of $91B, which includes $33.4B of total debt with $25B of senior notes issued in the most recent quarter. On the asset side, Nvidia's cash balance has increased from $12B to $22B in the past few quarters, but its investment book (marketable and non-marketable securities) has leapt from $45B in 2Q F1/26 to $128B by 2Q F1/27. The $99B equity book generated $23.7B of gains in the first half, and $25B more of equity investments are committed.

The question on everyone's mind is this: given Nvidia's on-balance sheet and off-balance sheet obligations, what is its capacity to meet all of these commitments in the case of an industry slowdown or downturn, and how much can Nvidia grow these obligations over time while still enjoying a margin of safety? How much support does the industry need and how much support **should** Nvidia offer?

In the chart below, we illustrate the history of these off-balance sheet items as well as our estimates of how much additional obligations Nvidia can reasonably support given its balance sheet and future cash flows. To be clear, this does not represent our base case or a definitive forecast for where we think Nvidia will go with its various backstops. Rather it focuses on the extent of Nvidia's ability to support obligations

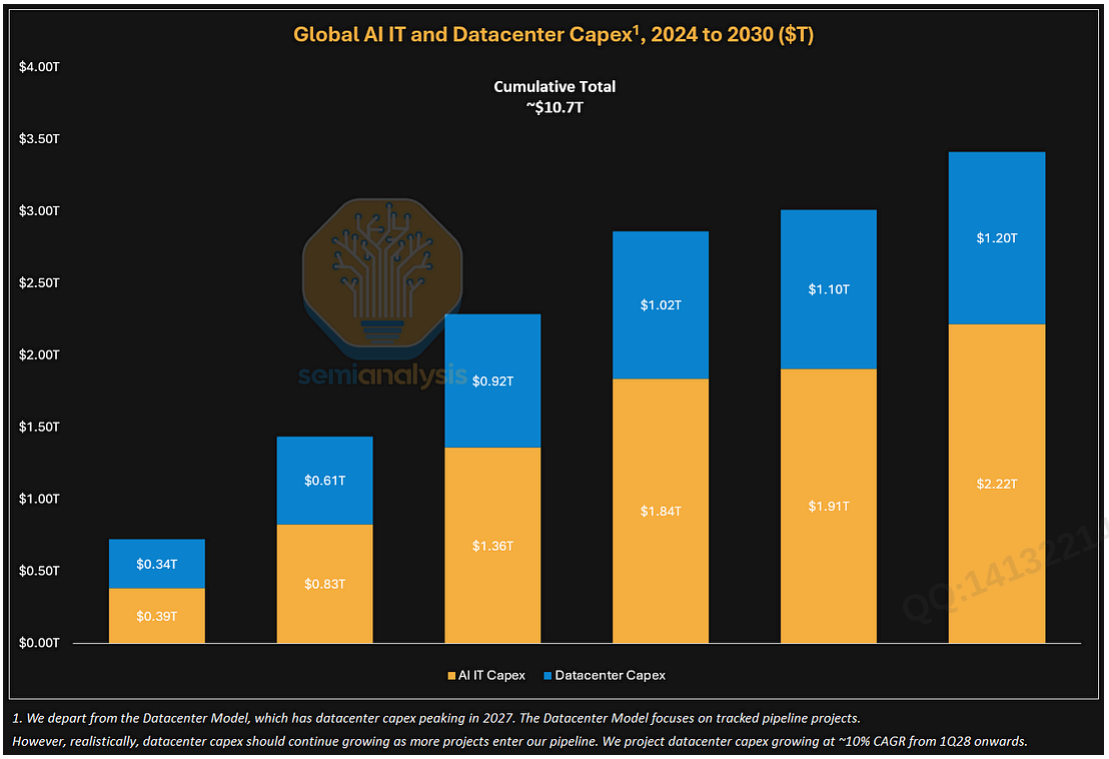
QQ:1413221435

given its business profile, and we argue given this considerable firepower, that it **should** have the willingness to continue expanding these commitments significantly.



Source: [SemiAnalysis AI Compute, Capital and Markets Model]

What is obvious from the diagram above is that Nvidia is a true behemoth when it comes to cash generation – consensus estimates have it generating ~$441B of EBITDA in F1/28, and this is why we reflect this via a growth of the cash balance to $1.4T by F1/31 – though in reality Nvidia will most likely reinvest this cash into various securities. These are huge numbers, yet the scale of overall AI investment is larger – we estimate ~$11T of cumulative capex from CY24 through CY29!. This note will explore how far Nvidia can take each form of support, and what impact this could make on the massive AI investment needs ahead of us.



1. We depart from the Datacenter Model, which has datacenter capex peaking in 2027. The Datacenter Model focuses on tracked pipeline projects. However, realistically, datacenter capex should continue growing as more projects enter our pipeline. We project datacenter capex growing at ~10% CAGR from 1Q28 onwards.

QQ:1413221435

Source: SemiAnalysis AI Compute, Capital and Markets Model

Why all the Guarantees and Backstops?

The mechanism that gives rise to these off-balance sheet guarantees and backstops is the same one we described in The Front End Gets Crowded (institutional article). The Gigascalers (Amazon, Microsoft, Google, Meta, Oracle) have been gatekeeping the investment grade capital that the AI buildout runs on and have benefitted tremendously from the high spread between offtake price they are being charged, and spot/short-term price that they are selling at as well as services that sit on top. In response, Nvidia has been manufacturing alternative credit anchors so Neoclouds and Neolabs can raise money without the Gigascalers. In these support arrangements, Nvidia sells GPUs at full margin on day one, then floors the buyer's revenue (in the case of the AI Cloud Partner (AICP) Program), or guarantees the buyer's landlord, or signs the lease itself, all so that non IG-rated operators can effectively borrow at IG rated pricing.

QQ:1413221435

Heads, demand holds and Nvidia wins twice, once on the GPU sale at full margin and again on the revenue share above the floor. **Tails, demand falls and the Neocloud makes no money but remains solvent, as the revenue floor is set high enough to repay lenders.** Nvidia's backstops and guarantees also act as a breakwater against further contagion as lenders have some shielding from a downturn. Yet how high and broad the breakwater that is the Nvidia backstop/guarantee depends on Nvidia's ability to meet these obligations.

Risk reward is asymmetric in Nvidia's favor as long as its balance sheet remains strong. Nvidia loses only if the backstopped Neoclouds fail to meet their lease and offtake commitments at the same time that Nvidia's own business and cash generation slow, degrading its ability to make good on these obligations. The two are correlated, since a Neocloud that cannot pay its lease is a Neocloud that has stopped buying GPUs. Only if a downcycle is steep enough to overwhelm Nvidia's ability to meet its off-balance sheet obligations, then it will be tails, we all lose.

By our count, Nvidia today backstops ~6.5 GW of capacity, most of which is not yet built. Our datacenter model has Microsoft, Meta, AWS and Oracle leasing roughly 15 GW of third-party capacity in 2026 and we expect them to lease more than 35 GW by 2028. Each of those leases typically run fifteen to twenty years, and developers borrow against them with investment grade pricing. The Gigascalers are the implicit backstop of the buildout and are in aggregate far larger force than Nvidia's backstopped capacity, yet Nvidia's program gets the attention because it is new and

QQ:1413221435

because many investors have Cisco PTSD, contracted vicariously without an appreciation for the dissimilarities.

Nvidia Datacenter Guarantees and Credit Support

Deal name	Capacity (GW)	Nvidia obligation
PORTS-Pike - OpenAI / SB Energy	4.25	$105.00B guarantee cap Land, power and shell
Second frontier AI lab (Anthropic inferred)	2.00	Selective credit enhancement Amount not disclosed
Ward County - Nscale / Ionic Digital	0.23	$0.86B rent guarantee cap First five lease years
Total included	**6.48**	**$105.86B disclosed caps + unquantified credit support**
PORTS-Pike expansion option - excluded	3.80	Option only. No commitment included above.

Source: SemiAnalysis Capital and Markets Model

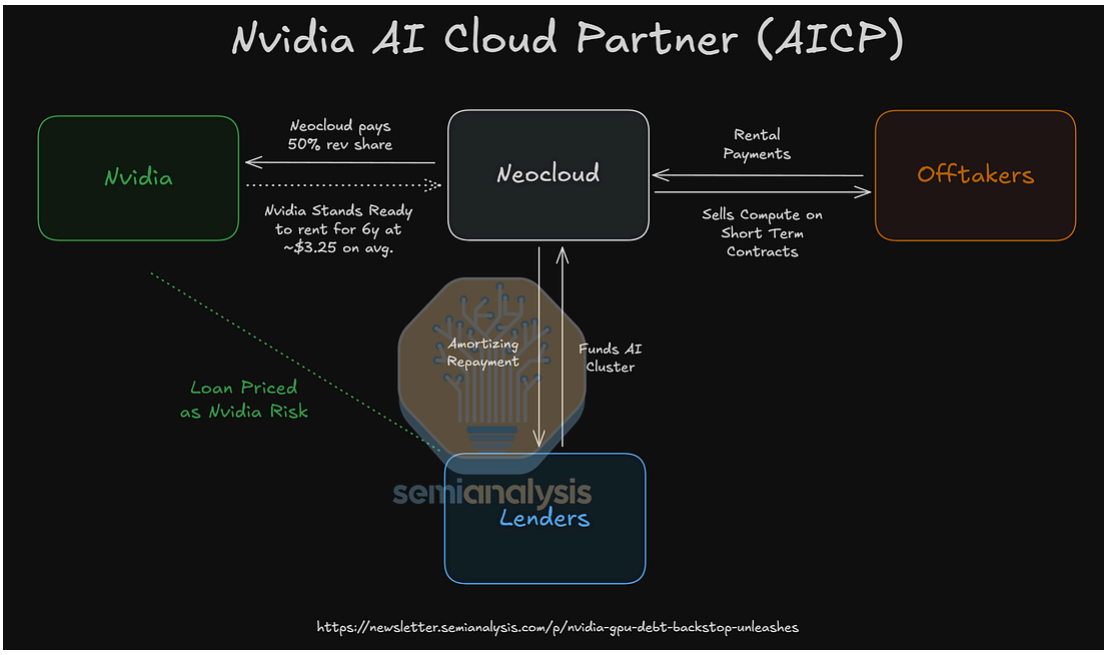
Source: SemiAnalysis AI Compute, Capital and Markets Model

We will now explore how this support could evolve in detail – walking the six off-balance sheet obligation lines in the recent quarter's disclosure, plus a new one we have created to track residual value guarantees under the capital partnership.

AI Cloud Agreements

This line item made its debut in the 2Q F1/27 earnings report released on August 26th, 2026 at $36B of aggregate obligation, spread across six years. It represents Nvidia's AI Cloud Partner (AICP) program, where Nvidia backstops GPU rental contracts for six years at a below market rate, approximately $2.35/hr/GPU on average for the GB300 vs the $4.50 to $4.60 range for 5y contracts, with the intent of making these AI clusters bankable. Neoclouds in this program would then seek to rent out compute on shorter term contracts to AI Natives and other tenants, but can always rent to Nvidia at any time should these contracts fall through or not renew. We published an extensive analysis of this structure in a July 2026 Newsletter Article.

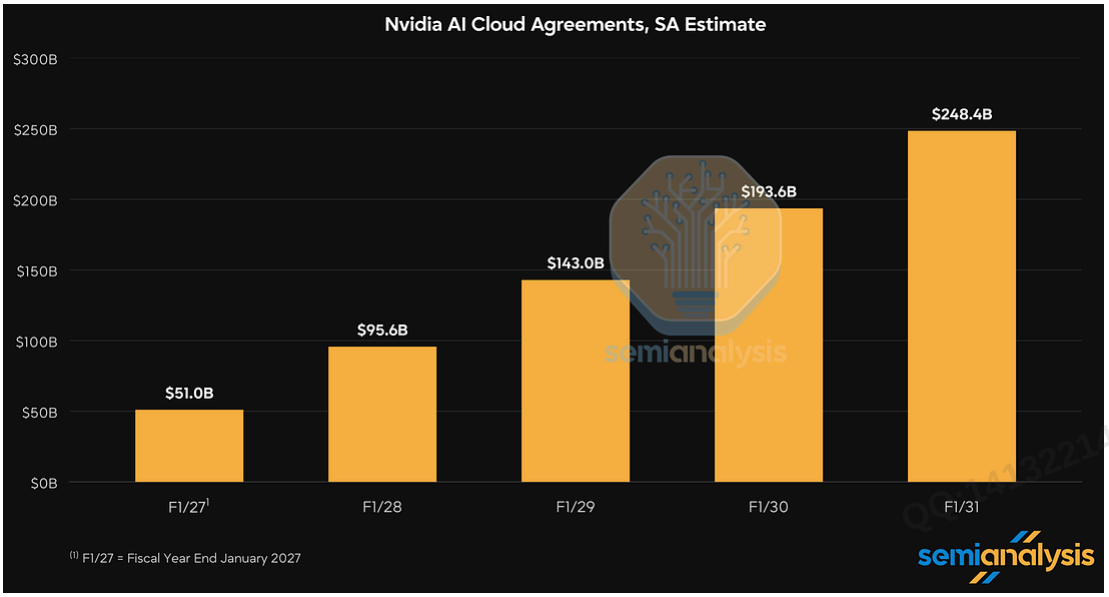
QQ:1413221435
QQ:1413221435



Source: SemiAnalysis AI Compute, Capital and Markets Model

There are three deals announced so far under AICP: Firmus's 360MW deal in Batam corresponds to a $21.1B total obligation, the SharonAI deal means $4.2B of obligations and we estimate the GMI deal corresponds to $2.2B of obligations for Nvidia. We believe there are still a few deals that have been completed before Nvidia reportedly paused AICP two weeks ago and estimate this line item will close F1/27 at $51B.

We do understand that Nvidia is not advancing any further AICP deals at this time, but believe that they still wish to be supportive in one form or another. In this analysis, our forward modeling still assumes considerable growth in this line item to reflect our view that Nvidia should offer support in some form, as well as to explore how far Nvidia could go in terms of explicit support.



Source: SemiAnalysis AI Compute, Capital and Markets Model

QQ:1413221435

Residual Value Guarantees under Capital Partnerships

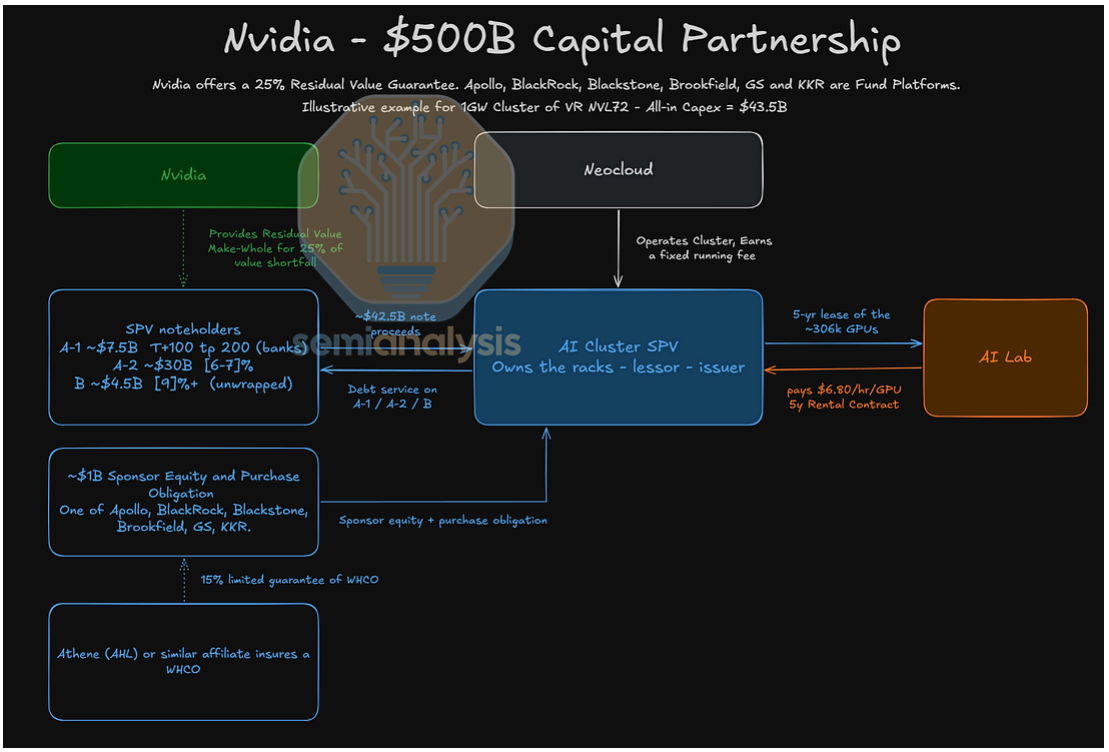
Earlier in August 2026, Nvidia announced its partnership with a number of private equity funds to mobilize over $500B of third-party capital to support the buildout of AI Clusters. Under this program, institutional investors are in the driver's seat and will evaluate and price each proposed project, with the only support from Nvidia being a residual value guarantee for up to 25% of a deal. In his X post detailing the program, Jensen Huang highlighted that this support would be limited and on a project-by-project basis. Though there are few details on what such a program could look like, we drew inspiration from the Google-Broadcom-Anthropic-Apollo TPU structure to surmise how deals could be structured.

In the TPU backstop structure, if Anthropic is unable to continue to service lease payments to the TPU SPV, then the structure can be resolved by Broadcom assuming the lease or by the racks being sold. If the racks are sold, then the residual value guarantee takes effect, as Broadcom must then make the senior noteholders whole after such equipment sale – effectively plugging the residual value gap, if any. The deal gets better – Google will play its part by also backstopping rental payments to the various datacenter operators. The effect is that lenders to the TPU SPV are effectively taking Broadcom credit risk while lenders to the datacenters are taking Google credit risk.

Nvidia's proposed structure would appear to wean investors away from this dual backstop as it not only reduces the residual value guarantee to only up to 25%, but it also is silent on any datacenter backstop.

QQ:1413221435

QQ:1413221435



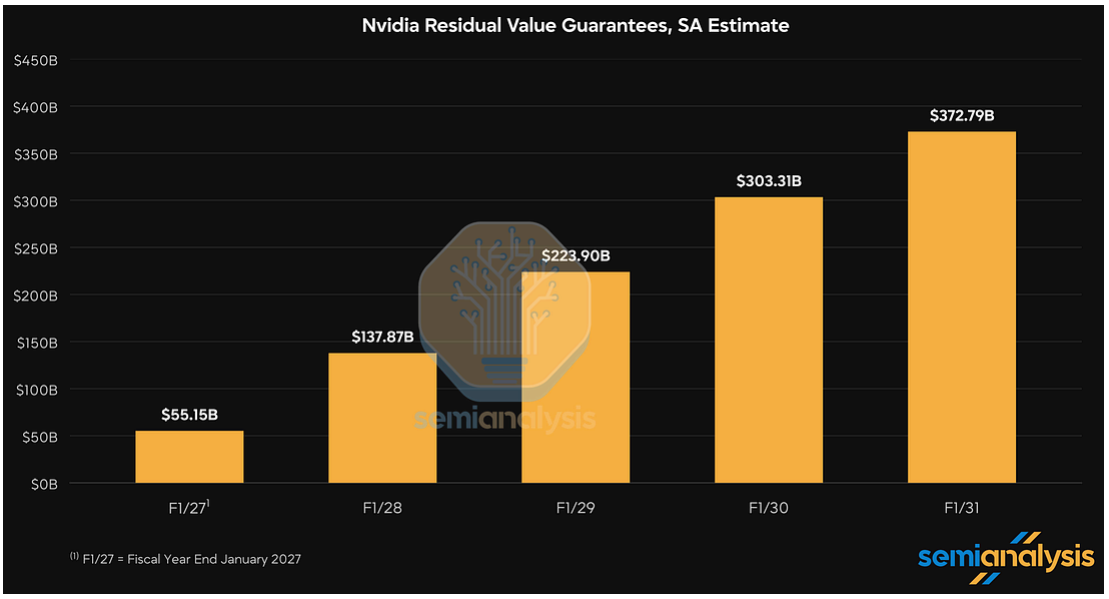
Source: [SemiAnalysis AI Compute, Capital and Markets Model]

We think that the funds could tranche the lending to create different pools of risk, with the first loss being taken by riskier equity/mezzanine style lending and the last loss taken by more senior creditors. Nvidia's 25% residual value guarantee would insert itself above the senior lenders but after the equity/mezzanine investors, adding a meaningful amount of crumple room to the structure, which could place the senior tranche at low enough risk of actual shortfall to attract conservative lenders like banks.

This structure can also provide the most bang for buck when it comes to Nvidia's off-balance sheet obligations. While the AICP creates $59B of obligations per GW of capacity enabled and the PORTS-Pike structure requires $25B of obligations per GW, the residual value structure would only require $9.4B of obligations per GW enabled, though the datacenter leg still needs to be solved for.

While Nvidia has said the program targets over $500B of capital, we think it will likely swell larger due to the strong compute demand. In the spirit of examining the full extent of Nvidia's maneuverability, we take an aspirational stance, assuming $2.5T of total funding under this program through F1/31, with the obligations under this line item ending F1/27 at $55B before growing to $373B by F1/31.

QQ:1413221435
QQ:1413221435



Source: SemiAnalysis AI Compute, Capital and Markets Model

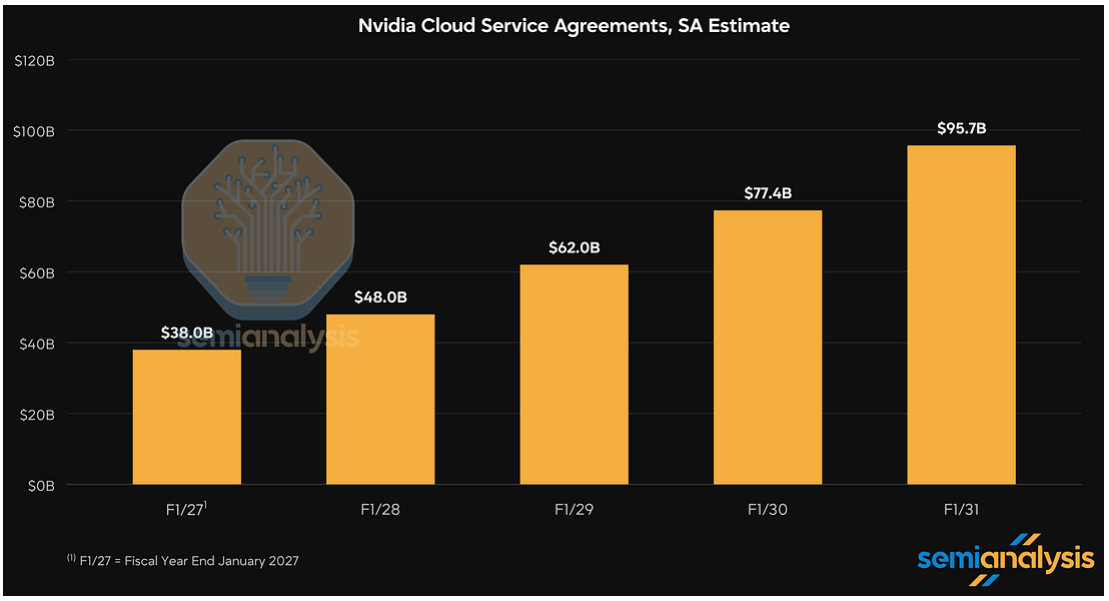
Cloud Service Agreements

Cloud Service Agreements are not a new line item, having first been disclosed in 3Q F1/23 at $1.6B before growing to ~$11B by 1Q F1/26 before jumping to $26B in 3Q F1/26. While AI Cloud Agreements represent Nvidia's backstops to external Neoclouds, Cloud Service Agreements primarily represent Nvidia's commitments related to its own in-house compute needs for model development and Continuous Integration/Continuous Development (CI/CD), though Nvidia has been known to resell this capacity at times.

Among the $29B of cloud service agreements listed, Nvidia has a $6.3B agreement thru 2032 with CoreWeave and we think the balance of the agreements are partially accounted for by ~360MW of compute in Batam, Indonesia via Oracle, ~20MW of capacity at YTL in Malaysia, 10-20k GPUs rented from CoreWeave, and a few other Oracle facilities in the US. Our AI Datacenter Model has more details on the exact facilities Nvidia is an end user at.

We estimate that these cloud service agreements will exit F1/27 at $38B and there is capacity to grow them to $96B by F1/31.

QQ:1413221435
QQ:1413221435



Source: SemiAnalysis AI Compute, Capital and Markets Model

Datacenter leases yet to commence for third-party reassignment

We think datacenter leases signed for reassignment become a large and fast-growing line. In Datacenter Backstop Wars (March 2026), we said Nvidia would lease multiple gigawatts in 2026, and indeed we now count about 1.2 GW signed year to date for both its own use and for third-party reassignment.



Nvidia Direct Datacenter Leases: Signed and Held (as of Sep 2026)

Landlord	Site	Capacity (MW)	Ready for service	Lease term	Nvidia commitment	10-Q line item	Signed	Source
Hut 8	Beacon Point Bldg 1, Calallen TX	352	3Q27	15 years, 3% escalator	$9.80B	Leases for reassignment	Jul 2026	Hut 8 8-K; tenant per WSJ, FT
Hut 8	Beacon Point Bldg 2, Calallen TX	352	2Q28	15 years, 3% escalator	$9.80B	Leases for reassignment	Jul 2026	Hut 8 8-K; tenant per WSJ, FT
Tract (Fleet Datacenters)	Peru Drive, McCarran NV	400	4Q27 to 2Q28	197 months	$12.00B	Own-use leases	2026	SA Datacenter Model; term per press
Vantage	100 Electric Ave, McCarran NV	224	2027 to 2030	n.a.	$6.72B	Own-use leases	2026, first 64MW	SemiAnalysis Datacenter Model
Digital Realty	Manassas VA	96	Live, full 2027	n.a.	$2.88B	Own-use leases	Before 2026	SemiAnalysis Datacenter Model
Mega Data Centres	Ein HaEmek, Israel	64	2Q27	n.a.	$1.92B	Own-use leases	2026	SemiAnalysis Datacenter Model
Aligned	Hillsboro OR	36	2026	n.a.	$1.08B	Own-use leases	Before 2026	SemiAnalysis Datacenter Model; multitenant building
Total		1,524			$44.20B	$45.00B filed ($20.00B + $25.00B)	1,232 MW (1.23 GW) signed in 2026	

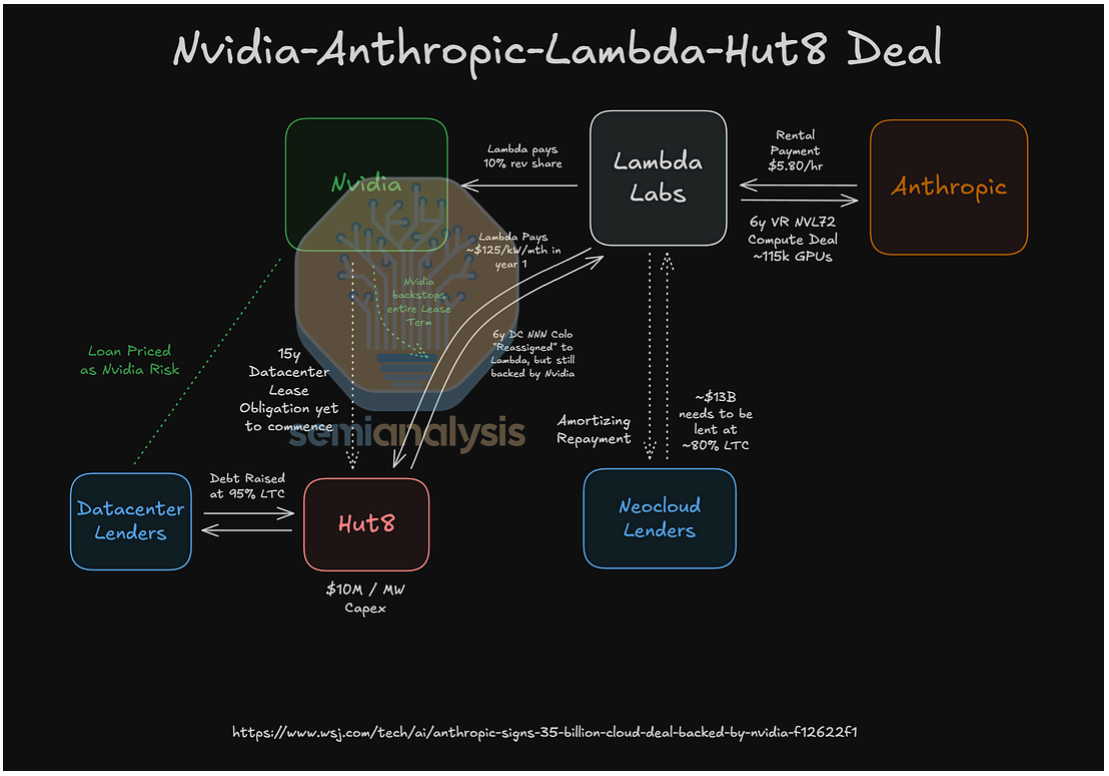
Source: SemiAnalysis Datacenter Model, Hut 8 8-K (Jul 2026), company filings, Tenant at Beacon Point per WSJ and FT.
Own-use estimates use $30M per MW. The $20MW total implies $24.60B, approximately the $25B filed line.

Source: SemiAnalysis AI Datacenter Model

The reassignment structure is effectively a credit substitution. A datacenter gets financed against Nvidia as an AA rated offtaker, then the lease is handed to a Neocloud that buys and operates Nvidia GPUs. Hut 8 funded the first Beacon Point building with $4.25B of senior secured notes at 6.129%, rated Baa2 and amortizing from 2030, which no Neocloud could price on its own name or for that term. Developers want both an investment grade tenant and a long lease, making the datacenter leg the hardest leg to solve for within the AI Project Trinity that was described in Nvidia GPU Debt Backstop Unleashes the AI Project Trinity (July

QQ:1413221435

2026). Nvidia supplies both, and keeps years seven through fifteen after Lambda's term to re-let or use.



Source: SemiAnalysis AI Compute, Capital and Markets Model

We think that this deal comes with a revenue share to Nvidia in order to compensate it for risk taken and the implicit use of its balance sheet. This seems fair enough – after all, we are by no means advocating for communism here. Even after reassignment to Lambda, we don't think Nvidia fully offloads the risk. Lenders that bought Baa2 paper on an AA tenant will not consent to a B-rated one, so Nvidia likely stays primary obligor until the project significantly de-levers and/or clears a DSCR test.

The fifteen year lease against a six year offtake gets described as a duration mismatch, but we would frame it as optionality. At year six, Nvidia can re-tenant with upgraded GPUs and touch only the power equipment, or let the operator renew and run the existing hardware at spot or short-term prices. Both paths are worth more in a world where interconnect queues run four to seven years.

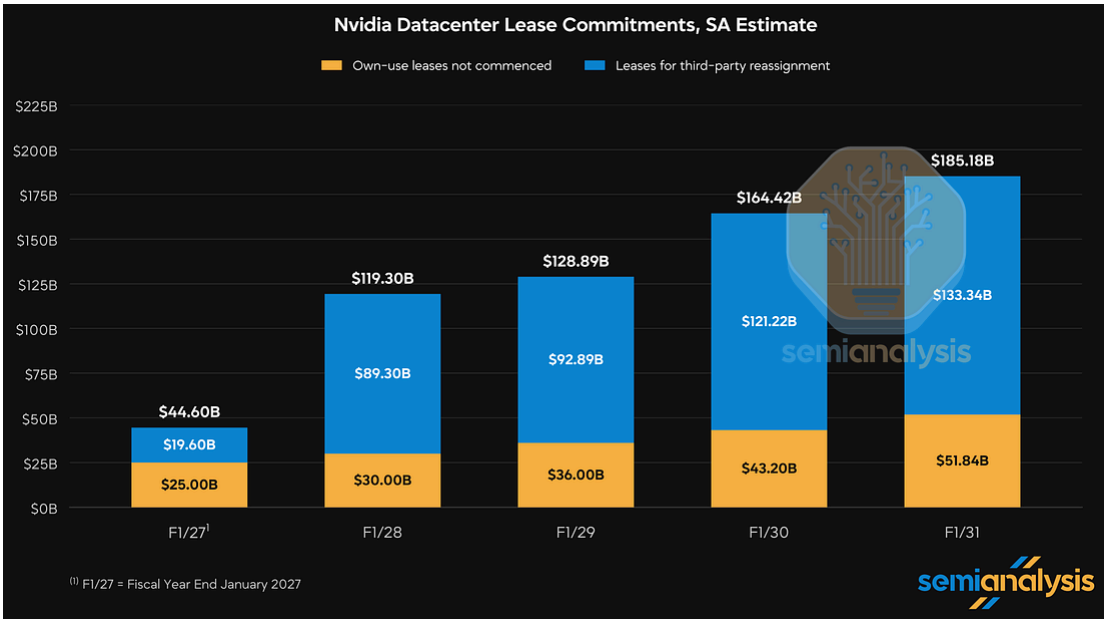
The scale of the above deal is small in the grand scheme of things. 704 MW is one campus. Most of the GWs coming online over the next five years are neither facilitated nor backstopped by Nvidia, and the Gigascalers' implicit backstops are far larger. Microsoft, Meta, Amazon, Oracle and Google lease ~39 GW of third-party capacity by 2028 against Nvidia's 1.5 GW as per our datacenter model.

QQ:1413221435

Datacenter Leases Not Commenced

Among the leases we outline in the table in the prior section, only the two at Hut 8 are for third-party re-assignment. Hut 8's July 8-K shows these two 352 MW leases at Beacon Point worth $9.8B each over fifteen years for $19.6B in total lease payments.

The rest of the leases in the table will be occupied by Nvidia for its own use, and the value of these commitments sit in the separate $25B "Datacenter Leases Not Commenced" line item in Nvidia's disclosure. These leases include: Tract's 400 MW at McCarran, Nevada, the first 64 MW of Vantage's campus, and Mega's 64 MW in Israel. Two older own-use leases, Digital Realty in Virginia at 96 MW, and Aligned in Oregon at 36 MW, takes Nvidia's total as a tenant to 1.5 GW.



Source: SemiAnalysis AI Compute, Capital and Markets Model

Guarantees and land, power and shell (LPS) guarantees

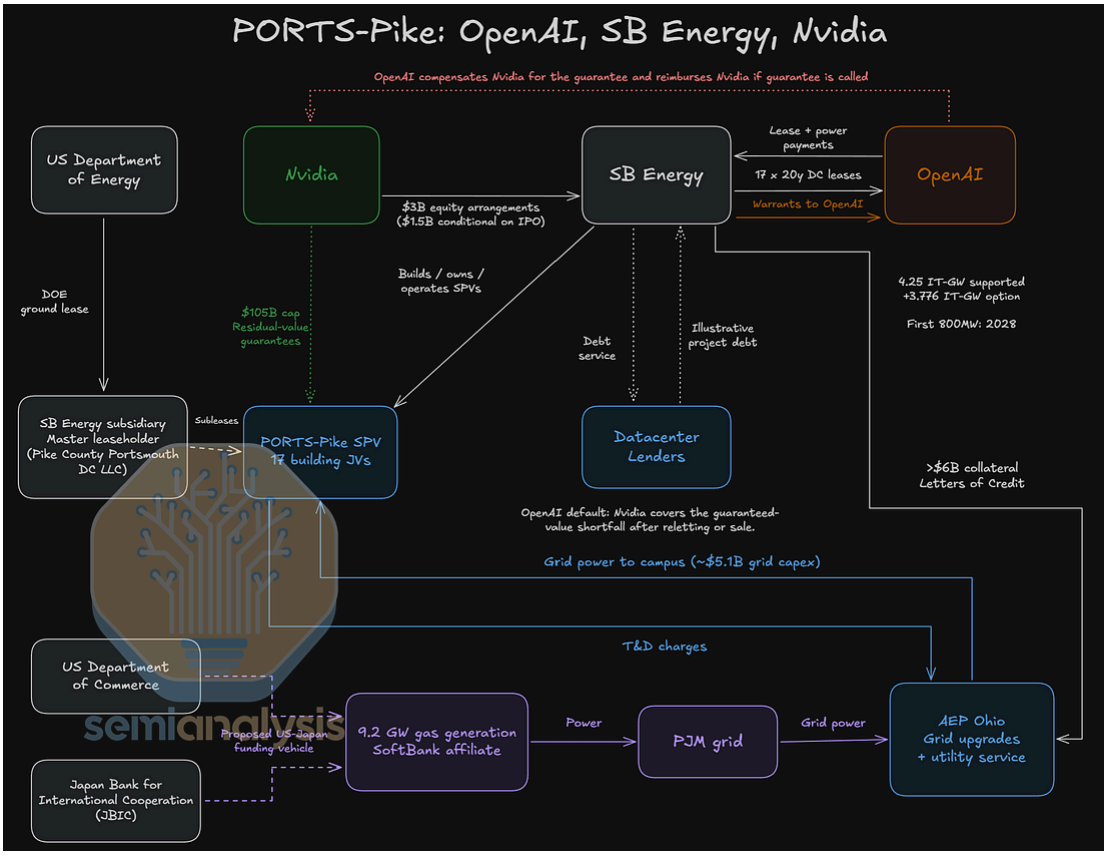
This line moved from $3.5B to $108.5B in one quarter, and we think it is where the program is going. Ionic Digital's S-1 shows Nscale on a 126 month triple net lease for 234 MW of capacity, with Nvidia guaranteeing the first five years of rent up to $860.3M, and Nscale's parent guaranteeing from year six, with Nvidia taking $60M of Nscale warrants as consideration. About $2.6B of the $3.5B from last quarter still has no name against it.

The PORTS-Pike deal added $105B of obligations disclosed in the most recent quarter, and it functions in a similar way, with Nvidia guaranteeing the value of a datacenter lease, but at a much larger scale. The deal includes credit support on land,

QQ:1413221435

power and shell for approximately 4.25 GW at SB Energy's PORTS-Pike campus in Ohio, capped at $105B against twenty year leases to an OpenAI affiliate, such support becoming effective in phases as buildings reach ready for service with the first expected in fiscal 2029. Exposure declines as OpenAI pays and terminates altogether if OpenAI reaches a satisfactory credit rating. The site will exclusively host Nvidia infrastructure, and Nvidia holds an option to backstop a further 3.8 GW in a second phase, and if Nvidia declines to backstop this phase, then OpenAI has to find an alternative guarantor with an A- credit rating or better. The first phase prices in $24.7B/GW of credit support, so exercising the full option on the same terms would add about $93B and take the PORTS cap to roughly $200B.

We diagram the structure of the deal below, but fully unpacking the entire structure is another article onto itself!



Source: SemiAnalysis AI Compute, Capital and Markets Model

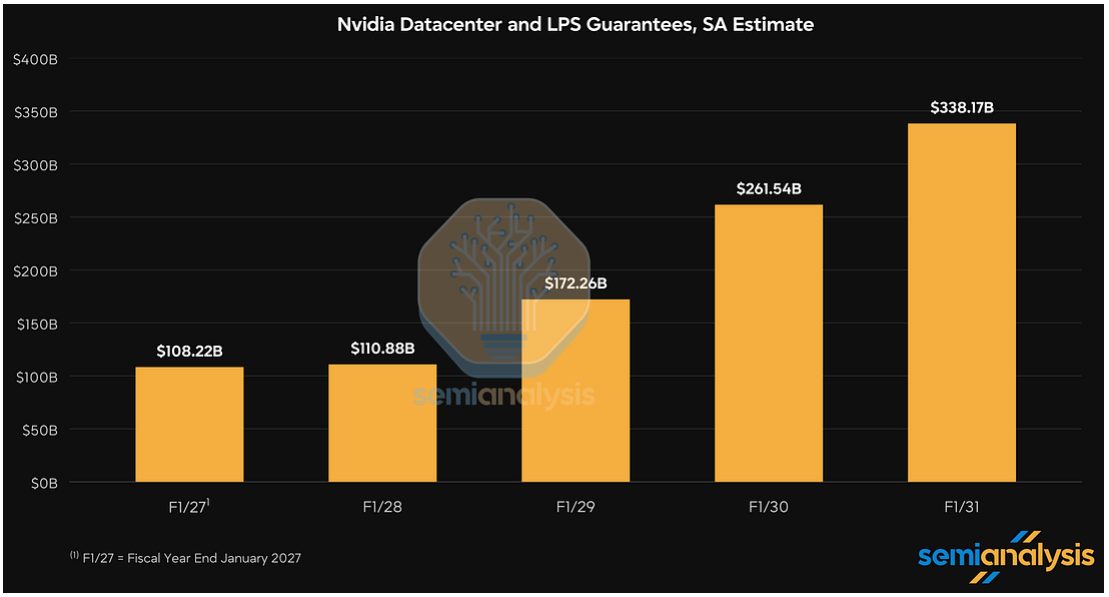
PORTS-Pike uses Nvidia's credit to help turn OpenAI's demand for compute into a campus that is financeable. SB Energy develops the roughly 8GW campus through 17 building-level landlord subsidiaries, each with a separate 20-year triple-net lease to an OpenAI affiliate. From a financing perspective, these project SPVs give lenders a defined set of assets and lease receipts to underwrite, allowing financing to follow through to the construction phase.

If OpenAI defaults on the lease, Nvidia faces a defined shortfall under the guarantee, with recoveries from replacement tenants or a sale reducing its exposure. Nvidia also has options to assume leases and arrange alternative uses for the capacity. OpenAI also compensates Nvidia for the guarantee though the terms and amount of the compensation structure is still unknown.

To further incentivize the counterparties within the deal structure, Nvidia has agreed to invest $3B into SB Energy, including $1.5B conditional on its IPO, while OpenAI receives SB Energy warrants through the broader commercial relationship.

Nvidia estimates that PORTS-Pike could generate $150-200B of revenue per hardware generation, giving it a substantial incentive to provide the guarantee. On a gross commitment basis, LPS guarantees also support more than twice the capacity per dollar as AICP – approximately $24.7M/MW of maximum exposure vs. ~$58.6M/MW for AICP. LPS guarantees address the datacenter financing constraint, the hardest leg of the AI Project Trinity.

Our forecast combines new LPS guarantees, third-party reassignment leases transferring into the guarantee line when they commence, and we assume Nvidia exercises their option to guarantee the remaining 3.75 GW of PORTS. Together, these assumptions take maximum exposure under datacenter and LPS guarantees from $108.5B today to roughly $338B by F1/31.



Source: SemiAnalysis AI Compute, Capital and Markets Model

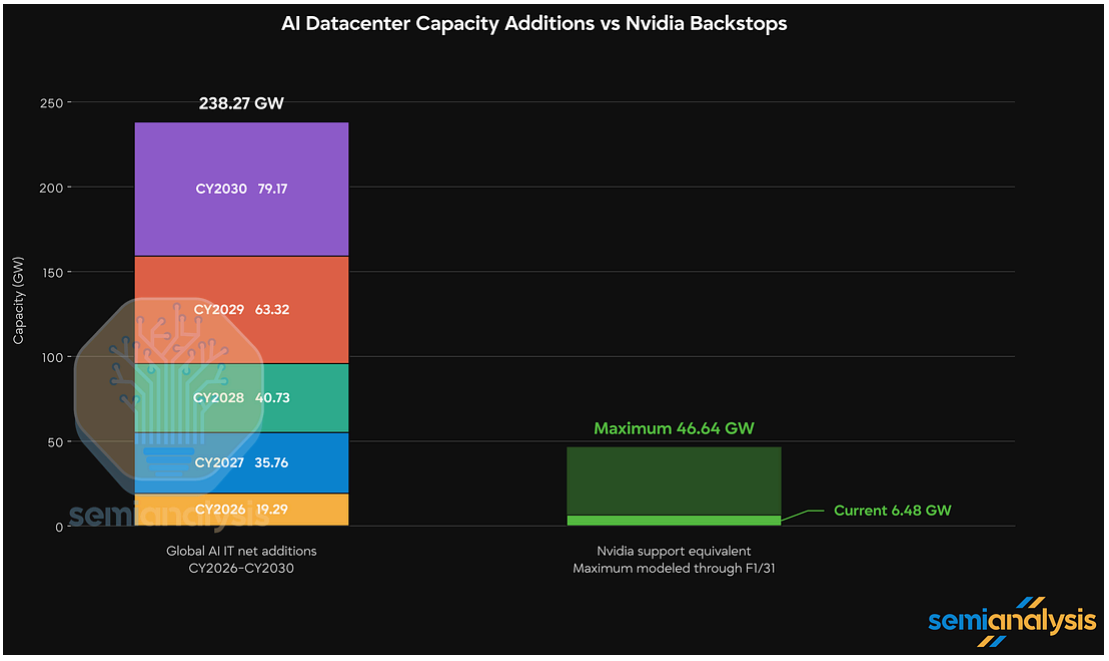
Putting Nvidia's Backstops in Perspective

QQ:1413221435

We believe the circular financing concerns overstate Nvidia's role in funding the overall AI buildout. The ~6.5GW in our current backstop count is equivalent to just under 3% of the approximately 240GW of net growth in global AI IT capacity forecast by our Datacenter Model between year-end 2025 and year-end 2030.

Our expanded support scenario models a maximum potential gross backstop of approximately 46.6 GW, a large amount of capacity, though 28.2 GW of this would be under the AI Capital Partnership, with only a 25% residual value guarantee on the supported lending. To align the comparison with the buildout through F1/31, we include the opening commitments and new commitments through F1/29, allowing two years for deployment. This 46.6 GW represents less than ~20% of total planned AI IT capacity additions through CY2030, and while it definitely can move the needle, Nvidia clearly cannot hold up the entire world on its own.

QQ:1413221435

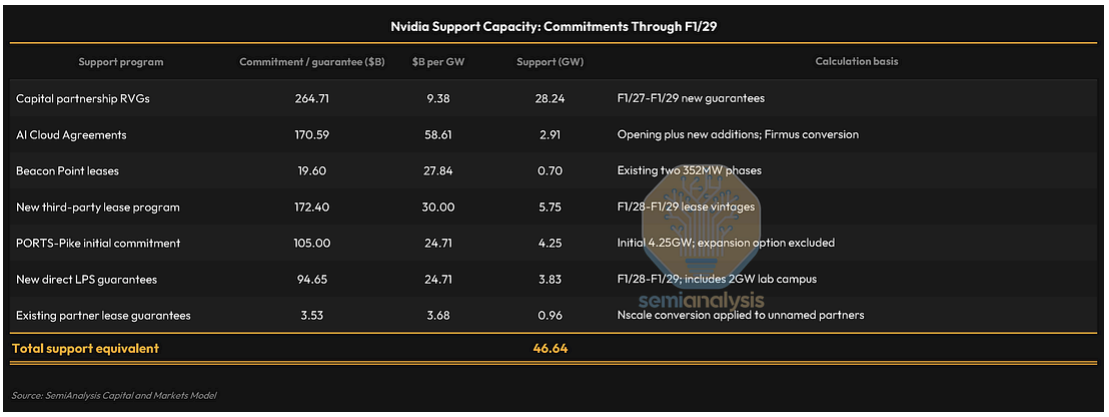


Source: SemiAnalysis AI Datacenter Model, SemiAnalysis AI Compute, Capital and Markets Model

This is also a clear illustration of why the AI industry is simply too big for Nvidia to fund the overall buildout through circular financing. The ~6.5 GW in our current backstop tally corresponds to approximately 2.5-3M GB300-equivalent GPUs, about 25-35% of the number of Nvidia GPUs we forecast shipping in CY2027. The backstopped capacity will be deployed over multiple years however – spread evenly over five years, the current tally would correspond to roughly 0.6M GB300-equivalent GPUs annually, well less than 10% of our CY2027 shipment forecast.

QQ:1413221435

We model an upper-end comparison – while we count each support type discretely, several instruments can support the same project and some announced capacity will be commissioned beyond 2030.



Nvidia Support Capacity: Commitments Through F1/29

Support program	Commitment / guarantee ($B)	$B per GW	Support (GW)	Calculation basis
Capital partnership RVGs	264.71	9.38	28.24	F1/27-F1/29 new guarantees
AI Cloud Agreements	170.59	58.61	2.91	Opening plus new additions; Firmus conversion
Beacon Point leases	19.60	27.84	0.70	Existing two 352MW phases
New third-party lease program	172.40	30.00	5.75	F1/28-F1/29 lease vintages
PORTS-Pike initial commitment	105.00	24.71	4.25	Initial 4.25GW; expansion option excluded
New direct LPS guarantees	94.65	24.71	3.83	F1/28-F1/29; includes 2GW lab campus
Existing partner lease guarantees	3.53	3.68	0.96	Nscale conversion applied to unnamed partners
Total support equivalent			**46.64**	

Source: SemiAnalysis Capital and Markets Model

semianalysis

Source: SemiAnalysis AI Compute, Capital and Markets Model

Even under the full expansion we model, the bulk of the AI buildout would still be financed without Nvidia's backstops. We view these programs as a way for Nvidia to get emerging operators to a scale which they can finance themselves. Successful projects under the backstop give operators a track record and allow lenders to experience underwriting their assets and cash flows at a lower risk. The intended end result is for customer revenues and independent capital to fund subsequent expansion, reducing the need for Nvidia's support. Nvidia directly benefits by broadening the customer base that is able to buy its GPUs at scale – our numbers suggest it can materially accelerate that development while remaining a minority source of support for the overall buildout

QQ:1413221435

How far can Nvidia take backstops and guarantees?

Green light – go, go go! Funded debt of about $36B sits against ~$283B of F1/27 EBITDA on our balance sheet model, so net debt to EBITDA runs at negative 0.1x in F1/27 and reaches -1.6x by F1/31. This is because the cash pile compounds faster than anything Nvidia can possibly borrow money on the balance sheet. Credit analysts we have spoken with put downgrade territory around one and a half turns of leverage (i.e. Net Debt to EBITDA). As of today, none of the off-balance sheet obligations count as leverage.

QQ:1413221435

Nvidia Leverage with Obligations Held Flat

USD billions, except ratios	F1/27[1]	F1/28	F1/29	F1/30	F1/31
OFF-BALANCE-SHEET EXPOSURE HELD FLAT					
AI Cloud Agreements	36.00	36.00	36.00	36.00	36.00
Datacenter leases for third-party reassignment	20.00	20.00	20.00	20.00	20.00
Datacenter and LPS guarantees (maximum exposure)	108.50	108.50	108.50	108.50	108.50
Capital partnership residual value guarantees	n.a.	n.a.	n.a.	n.a.	n.a.
Cloud Service Agreements	29.00	29.00	29.00	29.00	29.00
Datacenter leases not commenced for own use	25.00	25.00	25.00	25.00	25.00
Supply and capacity commitments	279.00	279.00	279.00	279.00	279.00
Total exposure held flat	**497.50**	**497.50**	**497.50**	**497.50**	**497.50**
DEBT, CASH AND EARNINGS					
Funded debt	36.00	44.83	45.90	85.82	149.28
Add operating lease liabilities	4.54	4.54	26.86	50.08	80.39
Less cash and cash equivalents	(82.14)	(276.58)	(550.88)	(907.30)	(1,363.33)
Net debt, including operating lease liabilities	**(41.60)**	**(227.22)**	**(478.12)**	**(771.39)**	**(1,133.66)**
EBITDA	282.64	460.79	539.44	577.10	704.53
Net debt / EBITDA	*(0.15x)*	*(0.49x)*	*(0.89x)*	*(1.34x)*	*(1.61x)*
LEVERAGE INCLUDING FIXED OBLIGATIONS					
Net debt plus fixed obligations	**455.90**	**270.28**	**19.38**	**(273.89)**	**(636.16)**
(Net debt + fixed obligations) / EBITDA	***1.61x***	***0.59x***	***0.04x***	***(0.47x)***	***(0.90x)***

[1] F1/27 = Fiscal Year End January 2027

Source: SemiAnalysis Capital and Markets Model

Source: SemiAnalysis AI Compute, Capital and Markets Model

Yet if we close our eyes, does the world disappear? Clearly not. So we built a harder test to evaluate how well Nvidia could weather the storm of a downturn and still meet its obligations. Add every off-balance-sheet commitment we forecast to net debt as if all were owed immediately, and divide by EBITDA. That ratio is just below 2x for F1/27, inside what the agencies tolerate, and it falls to 1.33x by F1/31, even with considerable expansion of off-balance sheet commitments.

Nvidia Stable Net Worth and Growing Off-Balance-Sheet Obligations

USD billions, except ratios	F1/27[1]	F1/28	F1/29	F1/30	F1/31
EARNINGS AND PAYOUT CAPACITY					
Net income	238.31	409.32	502.85	618.33	772.52
Earnings retained in equity	154.90	266.06	326.85	401.91	502.14
Annual dividends and buybacks (model proxy)	**83.41**	**143.26**	**176.00**	**216.42**	**270.38**
STABLE ASSETS AND NET WORTH					
Cash and cash equivalents	82.14	276.58	550.88	907.30	1,363.33
Receivables after 30% haircut	50.21	83.68	101.89	124.20	153.84
Securities after 80% haircut	27.64	31.75	36.16	42.34	51.41
Stable assets after haircuts	**159.99**	**392.01**	**688.93**	**1,073.84**	**1,568.57**
Less on-balance-sheet liabilities	(107.40)	(146.68)	(191.13)	(281.37)	(400.39)
(A) Stable net worth with planned payouts	**52.59**	**245.33**	**497.80**	**792.47**	**1,168.18**
Add one year of suspended dividends and buybacks	83.41	143.26	176.00	216.42	270.38
(B) Stable net worth including payout suspension	**136.00**	**388.59**	**673.80**	**1,008.89**	**1,438.56**
COVERAGE OF GROWING OBLIGATIONS					
(C) Total forecast off-balance-sheet obligations	589.48	915.67	1,247.40	1,656.95	2,067.42
Surplus / (shortfall) after payout suspension (B - C)	**(453.48)**	**(527.08)**	**(573.60)**	**(648.06)**	**(628.85)**
Coverage with planned payouts (A / C)	*0.09x*	*0.27x*	*0.40x*	*0.48x*	*0.57x*
Coverage including payout suspension (B / C)	***0.23x***	***0.42x***	***0.54x***	***0.61x***	***0.70x***

[1] F1/27 = Fiscal Year End January 2027

Source: SemiAnalysis Capital and Markets Model

Source: SemiAnalysis Compute, Capital and Markets Model

QQ:1413221435

Then take the harshest framing we could construct. Count only stable assets, meaning cash, receivables after a 30% haircut and a fifth of the securities book, and exclude inventory, the rest of the equity portfolio and all of the intangibles. Set those against total liabilities to get a stable net worth, and measure that against the off balance sheet obligations we forecast for each year, which start at the $497.5B filed today and reach $2.1T by F1/31 as supply commitments and the guarantee, lease, AICP and residual value lines grow.

The idea here is to consider the case where Nvidia's revenue evaporated over the following year. Could it still meet its obligations by drawing on its balance sheet and a few quarters of remaining cash flow before the apocalypse arrives? Even though these obligations come due over a long period of time, up to 20 years in some cases, we take a conservative stance and check whether Nvidia could pay them all off right away. Against today's commitments held flat, stable net worth covers 0.11x in F1/27, crosses 1.0x in F1/29 and reaches 2.35x in F1/31.

We push on this test further, and model out how well Nvidia's cash generation will grow as more obligations are piled on in our stress test. Against the growing stack it covers 0.09x in F1/27 and 0.57x by F1/31, or 0.70x if Nvidia suspends a year of dividends and buybacks.

QQ:1413221435

Nvidia Leverage with Growing Off-Balance-Sheet Obligations

USD billions, except ratios	F1/27[1]	F1/28	F1/29	F1/30	F1/31
FORECAST OFF-BALANCE-SHEET OBLIGATIONS					
AI Cloud Agreements	51.00	95.65	142.95	193.58	248.36
Datacenter leases for third-party reassignment	19.60	89.30	92.89	121.22	133.34
Datacenter and LPS guarantees (maximum exposure)	108.22	110.88	172.26	261.54	338.17
Capital partnership residual value guarantees	55.15	137.87	223.90	303.31	372.79
Cloud Service Agreements	37.97	48.00	62.00	77.36	95.69
Datacenter leases not commenced for own use	25.00	30.00	36.00	43.20	51.84
Supply and capacity commitments	292.55	403.97	517.40	656.74	827.22
Total forecast off-balance-sheet obligations	**589.48**	**915.67**	**1,247.40**	**1,656.95**	**2,067.42**
DEBT, CASH AND EARNINGS					
Funded debt	36.00	44.83	45.90	85.82	149.28
Add operating lease liabilities	4.54	4.54	26.86	50.08	80.39
Less cash and cash equivalents	(82.14)	(276.58)	(550.88)	(907.30)	(1,363.33)
Net debt, including operating lease liabilities	**(41.60)**	**(227.22)**	**(478.12)**	**(771.39)**	**(1,133.66)**
EBITDA	282.64	460.79	539.44	577.10	704.53
Net debt / EBITDA	*(0.15x)*	*(0.49x)*	*(0.89x)*	*(1.34x)*	*(1.61x)*
LEVERAGE INCLUDING GROWING OBLIGATIONS					
Net debt plus forecast obligations	**547.88**	**688.45**	**769.28**	**885.56**	**933.76**
(Net debt + forecast obligations) / EBITDA	***1.94x***	***1.49x***	***1.43x***	***1.53x***	***1.33x***

[1] *F1/27 = Fiscal Year End January 2027*

Source: SemiAnalysis Capital and Markets Model

semianalysis

Source: SemiAnalysis AI Compute, Capital and Markets Model

QQ:1413221435

These metrics are not crazy. We would go further and argue Nvidia should be offering more support. The direct marginal cost of a guarantee in the base case is zero, because the tenant pays and, at PORTS, OpenAI also indemnifies Nvidia. The

marginal benefit is a buyer that could not otherwise raise capital and a campus locked to Nvidia silicon for twenty years, which Nvidia itself sizes at $150B to $200B of revenue per generation.

Each backstop turns a B-rated operator into a financeable one and bootstraps the Neoclouds so they can gain operational experience and build a track record and potentially be put on the path towards an investment grade credit rating. The alternative, which we laid out in The Front End Gets Crowded, is a market where four Gigascalers are the only protagonists in the world with access to investment-grade capital, and thus only they can decide who buys compute. Our datacenter team's read of Nvidia's objective is to have twenty buyers at scale rather than just four. This playbook is not lost on Google at $75.5B of guarantees and Meta at $59.0B of obligations. The downside risk taken here is balanced by a floor of its own. Nvidia returned nearly $26B to shareholders in a single quarter, so the dividend and buyback are a $300B lever over three years to cushion a moderate downturn.

QQ:1413221435

Taking a step back, the limits that bind are not balance sheet limits. OpenAI is 97% of the guarantee cap, and the agencies are worried about how much of it points at one customer. Nvidia can spread the next tranche across other labs and across Neoclouds, but it cannot diversify away the fact that two frontier labs are the end payers behind almost every floor, lease and guarantee. That means that a demand shock at the labs hits everything, all at once. These are reasons that, if possible, the next tranche should go to new counterparties.

A demand air pocket marks down everyone's chips while Nvidia keeps its margins and its cash. The real risk to Nvidia is not from a cyclical downturn in AI demand relative to expectations, instead it is a scenario where competing chips take leadership, and Nvidia enters a spiral where it must price down to hold share. In this scenario, Nvidia's EBITDA falls, and the higher productivity of competing chips mean that Nvidia is now backstopping projects that are now uncompetitive. Both sides of the ratio – stable net worth, and likelihood of obligations being called upon move in a bad direction for the same reason. These obligations amplify that outcome without causing it.

That scenario is also the exact reason to be aggressive now. Nvidia leads today, AI demand is strong for at least the next three years, and power is the binding constraint on the buildout, so every powered site is scarce and whoever locks the site up owns that capacity for a generation of hardware or more, tying a campus down and creating a generation of buyers that have to run Nvidia GPUs. As long as Nvidia leads the market, backstops cost nothing, at least not explicitly. This is a play to win large parts of the market and we believe Nvidia should go for it while the window is open.

QQ:1413221435

Like Comment Restack

© 2026 Dylan Patel
548 Market Street PMB 72296, San Francisco, CA 94104
Unsubscribe

Start writing

QQ:1413221435

QQ:1413221435