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 12, 2026 • PAID

40 1 Share

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.*

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.

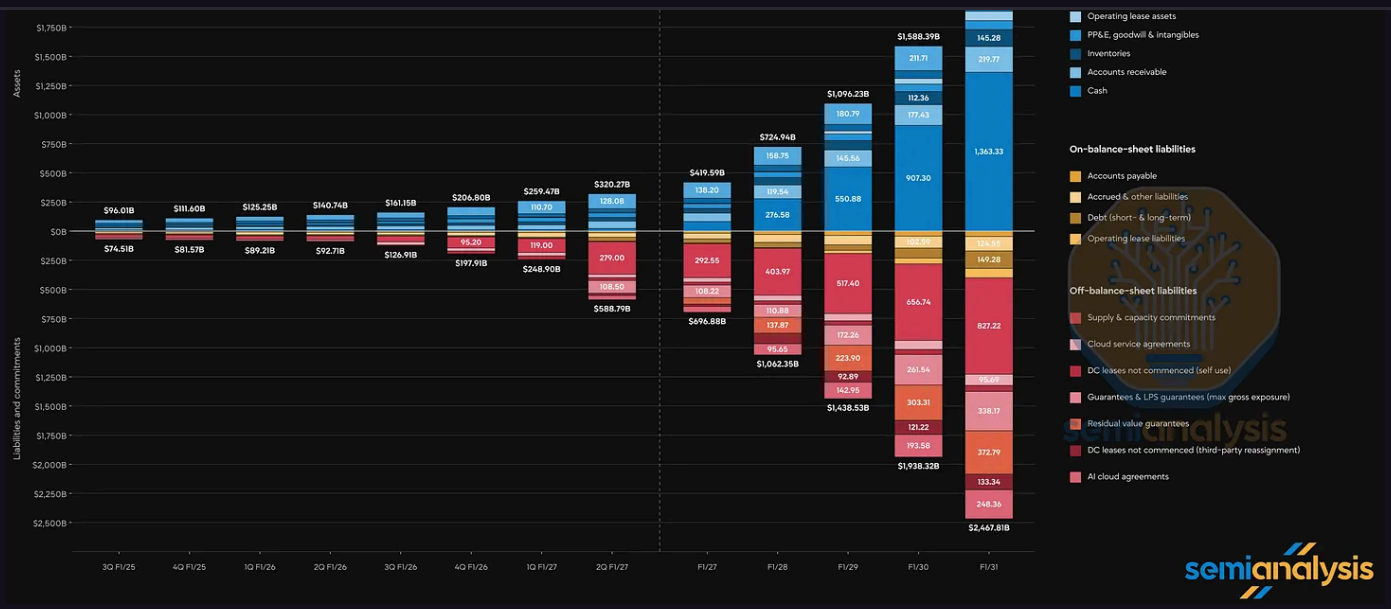
Nvidia guaranteeing the land, power and shell.

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.

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 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.

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.

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

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

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

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.

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 thi
believe that they still wish to be supportive in one form or another. In this an
forward modeling still assumes considerable growth in this line item to refle
that Nvidia should offer support in some form, as well as to explore how far I
in terms of explicit support.





Source: SemiAnalysis AI Compute, Capital and Markets Model

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

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

Source: SemiAnalysis AI Compute, Capital and Markets Model

Cloud Service Agreements

This post is for paid subscribers

Subscribe

Already a paid subscriber? Sign in

← Previous

© 2026 Dylan Patel · Privacy · Terms · Collection notice

Start your Substack | Get the app

Substack is the home for great culture