


THE MACHINE · ANNOUNCED AUGUST 10, 2026 · SUBJECT TO FINAL AGREEMENTS



THEMATIC RESEARCH · NVDA · AI INFRASTRUCTURE CREDIT

The Fifth Seat

What's Really Happening With NVIDIA's \$500 Billion Compute-Financing Machine — and Who Ends Up Holding the Risk

TON618 CAPITAL RESEARCH · AS OF AUGUST 14, 2026 · NVDA \$225.30 CLOSE, AUGUST 13, 2026 · MARKET CAP ≈\$5.46T

A follow-up to The Losing Bets That Built the World. Balance-sheet figures are read from NVIDIA's Form 10-Q as filed; CDS levels are press-reported prints, not terminal-verified, and are labeled as such throughout.



§1 The Verdict

13×

THE GROWTH OF NVIDIA'S
PRIVATE-INVESTMENT BOOK,
YEAR OVER YEAR, WHILE IT
ARRANGED \$500B OF OTHER
PEOPLE'S CAPITAL

NVIDIA is not selling its risk. It is recruiting co-investors so it can run a bigger book. The August 10 financing platforms create a seat no prior infrastructure wave has had: the originator, who earns on the risk without holding it. At the end of the pipe — through investment-grade NAIC designations into insurance general accounts — sits the American annuity holder, long the two things nobody can price: **GPU residual value and take-or-pay counterparty credit.** The first of those NVIDIA has now agreed to floor itself, up to 25% of a deal.


The risk-offload read — that NVIDIA built this machine to get AI-cycle exposure off its own balance sheet before the music stops — fails on the filings. A seller reduces exposure. NVIDIA grew its non-marketable investment book from **\$3.2 billion to \$42.3 billion in four quarters** — \$17.9 billion added in the April quarter alone, with **\$27 billion more already committed** through fiscal 2027 — *before* arranging \$500 billion of third-party capital. The motive is not balance-sheet fear; it is that NVIDIA's binding constraint has migrated. Not fab capacity. Not end demand. **Its customers' cost of capital.** Labs and neoclouds cannot fund a ~\$750-billion-a-year buildout from cash flow, so the supplier is manufacturing the financing channel that keeps its own order book — \$119 billion of supply commitments, \$95 billion of it due within this fiscal year — clearing. Risk transfer is the mechanism. Demand generation is the motive. And the collateral's largest scheduled loss-driver is NVIDIA's own product roadmap.

As of this morning the counterparties have names. NVIDIA's Form 13F-HR, filed **August 14** for the quarter ended June 30, discloses **\$63.4 billion** across eight holdings — a different population from the balance-sheet captions above, and two months later, so not a like-for-like comparison, but the first public naming of who is on the other side. **Intel is the largest at \$30.0 billion:** 214,776,632 shares, which is exactly the \$5.0 billion NVIDIA committed at \$23.28 in September 2025, now worth **six times what it paid.** Then SpaceX at \$21.0 billion. And then the two that matter here — **CoreWeave at \$4.7 billion and Nebius at \$0.3 billion,** equity in the neoclouds whose GPU-backed debt is the working template for the asset class these platforms will originate. The circularity, in two line items.

That same filing costs us a piece of our own argument, and we would rather pay it than not. The investment marks running through NVIDIA's income statement are dominated not by AI customers but by **Intel** — a strategic and partly political stake in a rival-turned-foundry-partner, which on our arithmetic accounts for something like 60–64% of the quarter's public-equity gain (workings in §9). The circular-marks reading of that line is weaker than the headline percentage suggests.

EXHIBIT 1 · THE BOOK THE "SELLER" GREW

NVIDIA non-marketable equity securities, carrying value at period end, with committed additions (NVDA)



\$17.9B

added to NVIDIA's private-investment book in a single quarter (10-Q, Note 6). The machine's designer is buying the asset class, not exiting it — the balance was \$3.2B a year earlier, and another \$27B is committed.

STATION 0 · THE PRESSURE GAUGE · READ BEFORE THE MACHINE



\$2 The Gauge

You may have arrived at this note the way most readers did: wondering what NVIDIA's credit-default swaps are trying to say. Four facts, in order of how much they explain:


- 1. The market is nine months old, and NVIDIA only recently gave it something to trade.** NVDA 5-year CDS began actively trading in **November 2025**. Before that there was no meaningful market in insuring NVIDIA's debt at all. Then on **June 18, 2026, NVIDIA printed its largest bond deal ever: \$25.0 billion across seven tranches maturing 2028 to 2056**, its first issuance since 2021. That roughly tripled total debt, to about \$38 billion including leases, and handed the CDS market two things at once — deliverable paper, and \$25 billion of fresh new-issue risk for dealers to hedge. But the raise added the same \$25 billion to the cash pile. Pro forma, NVIDIA still holds roughly **\$68 billion of net cash**, and it earned \$58.3 billion in a single quarter. In solvency terms there remains almost nothing to insure. In market-structure terms, there is now something to trade — and the widening began weeks after the deal priced.
- 2. The move is violent but the level is benign.** Roughly 40bps in late June; a record in the low 80s in late July, the largest single-day widening since trading began; about 80bps into the August 10 announcement. At one point the spread crossed *above* Alphabet's (~69bps): a spread ordering no credit analyst would defend on

fundamentals. Even at the record, ~80bps prices roughly a 6–7% cumulative five-year default probability at standard recovery assumptions — high-yield-tourist pricing on a AAA-shaped balance sheet.

3. **Each widening leg has coincided with NVIDIA adding obligations its old balance sheet never carried.** First the \$25 billion funded issuance in June; then the contingent stack — the financing platforms, the ecosystem investments, the lease guarantees. The CDS is repricing NVIDIA's *entanglement with its customers' leverage*, and hedging its newly minted paper. It is not pricing standalone solvency.
4. **The dominant flow is proxy hedging.** Funds warehousing unrated neocloud, data-center, and AI-adjacent private paper cannot buy protection on the actual borrowers, so they short the one liquid, correlated name. NVDA CDS is functioning as the AI complex's credit-hedging instrument — the tell being that Oracle, Alphabet, and even SpaceX CDS printed records in the same tape.

EXHIBIT 2 · THE GAUGE: A NINE-MONTH-OLD INSTRUMENT AT ITS RECORD

NVDA 5-year CDS — four dated press-reported prints, straight segments, nothing interpolated (NVDA)



The market has noticed the same structure this note describes — Michael Burry called the platforms "a sign of desperation" — but the desperation read requires NVIDIA to be *shedding* exposure, and the filings show the opposite (§1). The honest reading: the CDS is not a solvency signal; it is the **price of the missing instrument**. There is currently no way to short AI credit directly, so its price is being discovered in the nearest liquid proxy — and the moment rated compute ABS exists, that flow has a better home. Hence audit point 4 (§8): transfer is real only if NVDA protection tightens *relative to* the levered AI names once the platforms issue.

Three days after the record print, NVIDIA answered — though not in the way the protection buyers were positioned for.

STATION I · INTAKE · ANNOUNCED AUGUST 10

§3 The Machine

The announcement was not the launch of a machine. Goldman, Blackstone and Apollo had been working for **months** to draw up the debt deals that would let AI developers pay for NVIDIA chips. It was *slow progress* on those negotiations that prompted Huang to change tack and go public, and the stated aim was defensive: to assure NVIDIA's own investors that deep-pocketed firms stand ready to finance its customers, Anthropic and OpenAI among them. The \$500 billion itself is, in Bloomberg's characterization, "a round figure with no obvious provenance": no set time frame, and a combination of deals already under discussion with forecasts of demand

still to come. Read correctly, August 10 is not an origination event. It is a **confidence signal, issued because origination was stalling** — which is the demand-generation motive of §1 stated in the sponsor's own sequence of moves.

On August 10, NVIDIA announced six independent financing platforms — with **Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs, and KKR** — to mobilize **over \$500 billion of third-party capital** for AI infrastructure. Compute is the collateral. The capital is deployed, per press accounts of the structure, through private offerings and bonds issued by special-purpose entities, all of it "subject to execution of the final agreements." Jensen Huang's pitch is the collateral case verbatim: NVIDIA compute is *"broadly adopted, flexible across models and workloads, fungible and transferable across customers and operators, and continuously improved through CUDA"* — and the release's own copy, a sentence ahead of the quote and not Huang's words, adds the superlatives: the lowest token cost, the highest revenue, and the **longest life**. Apollo's Jim Zelter supplied the asset-class case: modern compute is *"a scarce, mission-critical asset class with compelling investment characteristics."* The supplier asserts the collateral is durable; the originator asserts it is investable; neither will hold it.

Separately — and deliberately not NVIDIA's — the price-discovery venue already exists. **Compute Exchange**, launched January 2025 by DRW's Don Wilson and Suna Said of Woodside AI, runs reserved-capacity auctions across 100-plus providers, offers forward contracts on future capacity, and in July 2026 opened a **secondary market in used H100 and A100 systems** (by the venue's own account, used B200s are not yet even available). The rails and the tape are being built by different hands — because the supplier does not want to own the printed mark on its own hardware's residual value.

And the collateral case did not stand on its own. In the days after the announcement, Huang disclosed the credit enhancement that makes the paper work: NVIDIA will provide **residual-value support of up to 25% of an opportunity** — if the resale or reuse value of the installed hardware falls short at the end of a financing term, NVIDIA covers part of the gap. In NVIDIA's own words the mechanism is "assessed carefully on a project-by-project basis," is "substantially lower than other compute-financing arrangements," and is designed to "complement — not replace — independent underwriting." No definitive agreements exist yet — these are memoranda of understanding — and neither the documentation, the loss order, nor the assumed useful life of the hardware has been stated. And the sequence matters more than the number. The announcement itself *unnerved* credit investors, who read it as NVIDIA absorbing open-ended leverage; the 25% figure was published the following day and the concern eased — which means the cap is doing two jobs at once. It is a **credit enhancement** that makes the collateral financeable, and simultaneously a **disclosed ceiling** that bounds the sponsor's exposure for the benefit of the sponsor's own creditors. Both sides of this trade needed convincing. But the direction is now on the record either way: **the supplier is writing a floor under the one variable this note identifies as unpriceable.**

One more filing fact frames why the machine had to be built. Three direct customers now account for **21%, 17%, and 16% of NVIDIA's revenue — 54% across three checks** — up from two customers at 16% and 14% a year earlier. A financing channel that broadens the buyer base is not a convenience for this order book; it is a structural necessity.

The distribution machinery moved within days. Goldman — the only bank among six otherwise-asset-management partners, and among the lead underwriters on NVIDIA's own June bond sale alongside JPMorgan and Morgan Stanley — was reported on August 14 to be already in talks with investors, positioned to lead the public debt offerings. The bank that helped place NVIDIA's paper is now placing the paper written against NVIDIA's chips.

And it all had a name before it had a machine. Six weeks before any of this, the Bank for International Settlements set out the pattern in its June 28 Annual Economic Report: circular financing in which chip manufacturers and cloud providers invest in AI laboratories and smaller computing firms that then commit to buying hardware from those same investors. The five largest hyperscalers are set to spend over \$1 trillion across 2025–26, outrunning

their own earnings and free cash flow, and the sector's net economic surplus, on the BIS's estimate, is declining and could turn negative. Direct lending funds, it notes, "have quadrupled their lending to the AI and IT sectors in the past five years, to about 15% of their portfolios." Its warning, in its own words: a repricing of risk "has the potential to be **similarly disruptive** by triggering a corporate credit freeze with wider implications for aggregate investment." The comparison is to the "large, synchronised corrections" it invokes two sentences earlier — of which 2008 is the cited precedent. An earlier BIS bulletin (No. 120, January 7, 2026, Aldasoro, Doerr and Rees) had already named the transition: financing the AI boom must shift "from operating cash flows to debt, with private credit playing a rapidly increasing role." Equity prices, it added, were running "far ahead of debt market pricing." A footnote supplies the scale of the channel now being widened — funds originated over \$40 billion of loans to AI-related companies in 2025, against about \$3 billion in 2010. The August 10 platforms are that shift, institutionalized, eight months later. We claim no novelty for the diagnosis — only that the apparatus built to answer it routes the risk into insurance general accounts and, per its own architects, possibly to retail.

Which raises the question this note exists to answer. A machine that moves risk moves it *to* someone.

 THE OPERATORS' MANUAL · WHO SITS WHERE

\$4 The Seats

The Losing Bets That Built the World established the recurring cast of every infrastructure wave: **builders** lose in aggregate, **suppliers** are paid in cash during the boom, **consolidators** buy the wreckage, **users** keep the surplus. In every prior wave, the financier was one of the four in costume — Insull *was* the builder, financing the grid through his own holding-company pyramid; Morgan was the consolidator, financing rail *after* failure at the reorganization price. The August 10 structure introduces a seat with no historical occupant: the **originator** — six firms whose economics are fees on origination, structuring, and management, with the risk held only in transit. The closest rhyme is not Insull or Morgan but the mortgage originators of 2003–2007 — originate-to-distribute, industrialized — with one improvement (four of the six own or manage the end-holder, so they eat some of their own cooking) and one aggravation (the collateral depreciates on the supplier's own product cadence).

Seven seats at the machine — what each gets, and what each is actually exposed to

SEAT	OCCUPANT	WHAT THEY GET	WHAT THEY'RE ACTUALLY EXPOSED TO
Supplier	NVIDIA	Customers' purchasing power underwritten by the deepest capital pools in the system, without consolidating their debt	Its "longest life / fungible" collateral claims are now load-bearing under \$500B of paper; a P&L already ~23% marks on its own investees; residual-value support of up to 25% per deal — a written put on the obsolescence its own product cadence schedules
Builders	Neoclouds	Cheaper term capital; a refinancing bid for their debt ladders	Disintermediation. Their moat is privileged access to allocation and capital, rented at a spread; a platform that funds anyone's cluster at a rated spread industrializes the scarcity away (see <i>The Thin Floor</i>)
Builders (split)	Hyperscalers	Little directly — they don't need SPE capital	Their capital advantage compresses; SPE-financed GPUs that must be filled become marginal price-cutters in spot compute if demand disappoints

SEAT	OCCUPANT	WHAT THEY GET	WHAT THEY'RE ACTUALLY EXPOSED TO
The Fifth Seat	The six GPs	Origination, structuring, and management fees on \$500B; risk held only in transit	The 2003–07 originate-to-distribute seat, now stated outright — Apollo expects to sell larger chunks of what it originates. Goldman is the only one with a banking arm. Four of six own the end-holder, so repudiation of the asset class strikes their AUM franchise
End holder	Insurers → annuity savers	Rated spread product; NAIC-designated ABS for general accounts (Athene, Global Atlantic, Brookfield Wealth Solutions, insurance SMAs)	Not "does AI work." GPU residual value at term, and take-or-pay counterparty credit — the two unpriceable variables that tranching makes look ratable, the first now partially floored by the supplier itself
The reader	Passive investors	—	The same household is now long this collateral twice: equity via NVDA at ~8% of the S&P 500, credit via the annuity general account — and executives have floated a third route, with some of the money potentially coming from retail investors directly . A correlation on no risk report anywhere in the chain
Sell side	Dealers, agencies	A genuine new asset class, rated inventory, and at last a direct instrument for AI credit	The curve they build is the mark-to-market machine (\$5)

The end-holder claim is not a forecast. The channel opened in March 2026, when CoreWeave's \$8.5 billion delayed-draw term loan closed with an A3 from Moody's and A-low from DBRS — the first GPU-backed facility to reach investment grade, and therefore the first to become eligible for the investment-grade mandates that govern insurance general accounts and pension portfolios. Press expectation is that the platform securitizations distribute among insurers, pension funds, and sovereign wealth funds on the same logic. What has *not* been established is that the NVIDIA platforms themselves can clear that bar: no rating standard for this paper has been resolved, and the useful-life assumption that would drive it remains unstated by every party.

The seat is already generalizing beyond this issuer. Apollo and Blackstone are anchor investors in a parallel Broadcom structure financing more than 20 gigawatts of compute, with roughly \$35 billion of financing already in hand, and Blackstone has sounded out investors on a further transaction of more than \$30 billion. Apollo's own economics are described plainly: it stands to unlock more fees as it expands its trading operation, **selling larger chunks of the loans it originates to other investors** — originate-to-distribute, named. JPMorgan's asset-management arm is reported to be discussing how to participate. The fifth seat is not an NVIDIA arrangement; it is a role the AI supply chain is now building at multiple issuers at once.

Seats one through four are as ranked in the prior note. The fifth is new — and it is the seat that, in every prior cycle that had even a partial version of it, won.

Every seat above is priced off a single assumption: that the collateral holds its value. That assumption has an adversary, and it is not the market.



§5 The Contradiction

A liquid compute market is the **precondition** for the asset class and the **detonator** under it. Rating agencies cannot rate compute-backed paper without observable liquidation values — which is precisely what Compute Exchange's forwards and secondary tape provide. But pre-mark, collateral is worth what the model says; post-mark, it is worth what the screen says. ABX did not cause the subprime losses; it made them instantaneous, public, and undeniable. The compute forward curve will do the same for GPU residuals.

And the residuals have a scheduled adversary: **the supplier's own roadmap**. NVIDIA needs H100 and H200 spot rates high to support the residual assumptions in the paper its platforms will originate — and it needs to ship Rubin, which obsoletes them, on an annual cadence it cannot slow without ceding the frontier. NVIDIA's product roadmap is the largest single loss-driver in the collateral pool being financed on NVIDIA's collateral narrative. Early price discovery already leans the wrong way: Compute Exchange reports secondary demand concentrated in **H100 and A100 systems rather than Blackwell** — the residual curve being discovered in real time, sloping against the "longest life" claim. Our depreciation series (*The First Invoice*) documented the useful-life debate; that debate now has a \$500-billion credit structure built on one side of it. And NVIDIA's roadmap is not the only force pressing down: China's domestic accelerator buildout is reported as a second bid-side threat to the same resale market, compressing the very re-rental economics the lenders are relying on.

For context on what the rating agencies will do with this: through 2025 S&P capped data-center ABS at A+, on the sector's short history. That ceiling has since given way, with senior GPU-backed tranches reaching higher and Moody's signaling no hard cap — and saying in February 2026 that it would run its own probability assessments on AI data-center leases. Residual-value insurance and guarantees are already a routine enhancement in these structures. So the 25% support is not an exotic concession — it is a standard technique, applied at unprecedented scale by the party that also builds the asset.

The lenders' own stated mitigant runs straight into this. Collateral is expected to be some combination of the chips and the offtake agreements, and the answer to a borrower default is that **the chips can simply be rented to someone else** — which is to say the downside protection *is* the re-rental rate, and the re-rental rate is precisely what each new generation compresses. Skeptics quoted in the same reporting make the other half of the point: chip valuations are inflated by record demand today, and the buildout risks an oversupply of compute years out. The mitigant and the risk are the same variable.

Which is what makes the 25% residual-value support the most consequential detail in the entire announcement. NVIDIA has agreed to cover part of the shortfall between expected and actual resale value of accelerators whose obsolescence it schedules. **It has sold a put on its own product cadence**. Every Rubin launch that strengthens the franchise weakens the collateral and moves that option closer to the money. The annuity holder is short an option nobody controls. NVIDIA is short an option it controls entirely — and has every strategic reason to exercise it against itself. Held flat across a \$500 billion target the exposure would be on the order of \$125 billion — an upper bound that compounds two separate hypotheticals, since it assumes both that the entire uncommitted target is reached and that every dollar of it is support-eligible, neither of which is established. Nobody should quote it as a commitment: the support is deal-by-deal, covers a gap rather than a principal amount, and no agreement has been signed. The order of magnitude is the point, not the figure.

The guarantee question is not hypothetical either. NVIDIA's current filings already disclose a template: beginning in fiscal 2026, the company **guarantees partners' facility-lease obligations in the event of their default, in exchange for warrants** — maximum gross exposure \$3.5 billion, terms of five to seven years, \$712 million escrowed by the partners, the guarantees classified as credit derivatives with fair-value changes running through Other income. This is the exact shape of the question that decides the whole thesis — and Huang has now answered its first half in public. Residual-value support of the kind the lease guarantees already demonstrate is

stated policy for the platforms. What remains open is how it is documented, where it sits in the loss order, and whether it lands on NVIDIA's balance sheet as a guarantee, a credit derivative, or nothing at all until it is drawn. To the extent it exists, **the \$500 billion is distribution, not transfer** — the risk rides third-party balance sheets while a defined slice of it economically remains NVIDIA's.



§6 The Return Line

Everything above rests on filings and announcements. This section is conjecture, labeled as such — our best guess at how the machine behaves, stated before the SPE documents exist so it can be scored later.

The clean-agency story is the no-free-lunch story: the terminal holder of the paper earns the spread because they accepted the downside — you do not get the cash flow from the hottest asset class in the world without owning its risk. We think that story is directionally right and wrong in three compounding ways.

The buyer will own the risk but not be paid for it. "Accepting the downside" and "being compensated for the downside" are different claims, and the securitization apparatus exists to open a gap between them. The 2006 AAA CDO buyer also accepted the downside — for 30 basis points over Libor. Tranching, an investment-grade NAIC designation, and a collateral class with five quarters of operating history are how a fat, unpriceable risk is sold at a thin, priceable spread. The yield will be calibrated to the rating; the risk is calibrated to reality. The gap between those two calibrations *is* the fifth seat's margin.

The buyer is short an option the counterparty controls. The feature with no historical precedent. A 2006 mortgage-bond buyer took house-price risk that nobody controlled. The compute-paper buyer takes obsolescence risk whose timing is set by one company — NVIDIA decides when Rubin ships and how hard it deprecates the pool backing the paper — and that same company sponsors the financing channel and sells the collateral. The downside's trigger sits in the seller's product-roadmap meeting. No rated spread prices that correctly, because ratings have no cell for "the sponsor schedules the impairment."

In stress, the risk boomerangs to the center. Structured investment vehicles were off-balance-sheet until December 2007, when Citigroup took them back rather than let the franchise die. Every sponsor here has the same compulsion: Apollo cannot let Athene-held compute paper impair without damaging the machine that funds its entire model, and NVIDIA cannot let a major offtaker default and print a loss onto the collateral curve under \$500 billion of paper — it is *already* backstopping partner leases in exchange for warrants. So our base case is a three-layer ownership: **the annuity holder owns the risk contractually, the sponsors own it reputationally, and NVIDIA owns it economically** — because any default in the pool impairs the collateral narrative its order book runs on. The 25% residual-value support (\$3) means the third layer is no longer purely economic: a defined slice is contractual too, by the sponsor's own statement, before a single agreement has been signed. That is the boomerang arriving early — not in stress, and not, as we first wrote, simply as the price of getting the paper sold. The reporting is more specific: the cap was published to calm credit investors alarmed by the sponsor's own exposure. The risk returning to the center and the sponsor advertising a limit on how much can return are the same act, four days after the announcement and before a single agreement is signed. The music-stops scenario is not a clean loss eaten by a distant bondholder; it is the discovery that the risk was never severed from the center. Every prior cycle ended with the same discovery.

Questions we cannot yet answer, listed so they can be checked off as the documents surface: Who holds the **equity and first-loss tranches** — retained by the GPs, sold to yield-hunters, or quietly taken by NVIDIA in the warrants-for-guarantees shape it already uses? First-loss placement is the honest map of who believes what, and the 25% residual-value support raises rather than settles it: support that pays only at term-end on a resale shortfall is a very different instrument from first-loss on the pool, and the two are easy to conflate in a press release and hard to conflate in a waterfall. What **residual-value assumption** clears a rating committee for an asset whose secondary tape already trades off "longest life"? Is there any structure in which NVIDIA can *credibly* decline to support a failing offtaker — and if not, is the \$500 billion economically recourse to NVIDIA regardless of what the documents say? If the paper prices wide, the platforms fail to scale; if it prices tight, the risk is mispriced at \$500 billion of scale. **Is there any spread at which this paper is both sellable and fairly priced?** And

when the first impairment prints, does the NAIC re-designate the class — and what does a forced insurance-general-account seller do to a collateral curve with a nine-month history?



THE OTHER SIDE · THE BEST PUBLISHED COUNTER-CASE

§7 The Case Against This Note

Everything above is a reading, and readings can be wrong. A note this opinionated owes its best counterargument a fair hearing before it asks anyone to watch tripwires on its behalf — and the best one on the tape is Bank of America's. Vivek Arya: *"The burden sits with the consortium, not NVDA's balance sheet. NVDA guarantees asset quality, not the debt, turning the bear's depreciation worry into the enabling feature."* The argument is genuinely strong, and it runs like this: NVIDIA is not underwriting anyone's credit, only the quality of the machine — and the party best placed to warrant that a GPU still works and still rents is the party that built it and controls its software stack. On BofA's numbers, prior vendor-financing exposure ran to roughly 15% of NVIDIA's ~\$470 billion of projected 2026–27 free cash flow; shifting that load to third-party capital releases cash for buybacks. On this reading the platforms are a *de-risking*, and the 25% cap is the proof: NVIDIA kept the sliver it is uniquely qualified to price and gave away the rest.

We think that is right about the debt and wrong about the residual — and the distinction is the whole note. Guaranteeing "asset quality" is precisely the guarantee whose cost NVIDIA determines, because obsolescence is not a defect in the machine but a decision about the next one. A warranty that the GPU still functions is cheap; a warranty on what it still *rents for* is a short position in your own product cycle. Arya is correct that the consortium carries the debt. Our claim is narrower and survives his: NVIDIA retained the one exposure nobody else could price, and it retained it in the one variable its own roadmap moves.

The dissent is not ours alone. Wells Fargo and Mizuho have flagged NVIDIA's continuing exposure, UBS the circular-financing question, and Seaport's Jay Goldberg the risk that investors eventually ask whether the spending cycle makes economic sense. The equity's own verdict on announcement day was a decline.

**§8 Audit Points**

Five falsifiable checks, pre-registered August 13, 2026 — the first scoring at NVIDIA's Q2 FY27 10-Q, expected in late August 2026.

Registered to the trigger monitor; what each outcome would mean

#	THE CHECK	IF IT FIRES
1	The commitments/guarantees footnote quantifies the 25% residual-value support (§3): maximum exposure, accounting treatment (guarantee vs credit derivative vs undisclosed until drawn), and whether it extends beyond the existing \$3.5B facility-lease program	Distribution - not - transfer confirmed and sized. Huang has stated the policy; the footnote is where it becomes a number. Silence in the Q2 filing is itself informative — it means the exposure is contingent on agreements not yet signed
2	The non-marketable investment book exceeds ~\$55B (vs \$42.3B reported plus \$27B committed)	Leverage - not - exit confirmed — the co-investor read of §1 holds
3	The first rated compute ABS prints: its residual-value assumption, tranche NAIC designations, and whether NVIDIA appears anywhere in the waterfall	Sizes the pipe and names the bag
4	NVDA CDS tightens <i>relative to</i> Oracle and levered AI-adjacent credit after the platforms issue	Transfer is real; the proxy-hedge flow found its home. NVDA and ORCL staying locked together = nothing moved
5	Direct-customer concentration (21/17/16 = 54%) broadens as SPE-financed buyers appear	The machine is generating new demand rather than funding the same three checks

The three-part falsification stands, though point 1 has already moved against it: if the residual-value support proves narrow and rarely used, point 5 broadens, and point 4 decouples, the transfer is genuine, the annuity holder is being paid roughly fairly, and §6's boomerang is wrong. We would rather record that outcome than have hedged the claim.

§9 What Is Observed, Computed, and Conjectured

Observed (primary sources, independently verified): every balance-sheet and footnote figure — the non-marketable investment ramp and its \$17.9B single-quarter addition, the \$27B of commitments, the \$30.2B marketable equity book and its lockups, the \$13.4B and \$2.6B unrealized-gain components of Other income, the \$3.5B facility-lease guarantee and its \$712M escrow, the 21/17/16 customer concentration, the \$119B supply and \$30B cloud commitments, net cash and the undrawn \$25B commercial paper program — all from NVIDIA's Form 10-Q for the quarter ended April 26, 2026, as filed. The June 18, 2026 \$25.0B seven-tranche senior notes offering is from NVIDIA's Form 8-K as filed. The 25% residual-value support comes from NVIDIA's own newsroom post and Jensen Huang's public statement, both August 11, 2026, reported contemporaneously by Bloomberg — a company disclosure rather than press inference, but one not yet filed with the SEC and not yet reduced to a definitive agreement. The August 10 announcement facts and both executive quotes are from the NVIDIA and Blackstone press releases as published. An independent claim-verification pass was run before publication; it corrected three items in our own draft — the June bond deal had been omitted and total debt stated at the stale 10-Q figure, a commercial paper program had been mislabeled a revolver, and B200/H200 had been wrongly included in the used-secondary-market claim — and confirmed every 10-Q-sourced figure exactly, including recomputing the DSO and inventory-day table to the decimal from four separate primary filings.

Corrections. This note was published and then revised three times on substance the same day, and each revision fixed something we had got wrong rather than something the world had changed. They are recorded together because the pattern is more useful than any one of them. **(i)** Version 1.0 stated total debt from the April 10-Q and omitted NVIDIA's June 18 issuance of \$25.0bn of senior notes — public, filed, and easily checked; it also mislabelled the company's commercial paper programme a revolver and wrongly included B200 and H200 in a claim about the *used* secondary market that its own cited source contradicted. **(ii)** Version 1.1 presented NVIDIA's residual-value support as an open question when Huang had answered it publicly on August 11, three days before we published; we had not swept the company's own statements between the announcement and our publication date. **(iii)** Version 1.2 corrected one of our own framings rather than an outside fact: v1.1 said the 25% cap "arrived as the price of getting the paper sold," attributing it to buy-side demand, when the reported sequence is that the announcement alarmed NVIDIA's *own* creditors and the ceiling was published to bound its exposure. **(iv)** Version 1.3 carried two BIS figures — a "\$40 billion, fivefold year-on-year" private-credit rise and a warning of "force comparable to the 2008 global financial crisis" — taken from press summaries. Neither appears in the Annual Economic Report. The \$40bn figure is Bulletin 120's and compares 2025 with 2010, not with the prior year; the AER's own statistic is the five-year quadrupling of direct-lending exposure quoted in §3, and its warning is that a repricing "has the potential to be similarly disruptive." What made that error feel safe is what should have raised the alarm: half a dozen outlets carried identical wording, which reads as corroboration but is one paraphrase propagating. All four are fixed in the text above. We disclose them for the same reason we count everyone else's.

Observed (filed today): NVIDIA's Form 13F-HR for the quarter ended June 30, 2026, filed August 14, 2026 — eight reportable positions totalling \$63,439,974,569, read from the SEC information table directly. The Intel line (214,776,632 shares, \$29,989,261,126) reconciles exactly to a \$5.0 billion commitment at \$23.28 per share, which is how the 6.0x and the \$25.0 billion unrealized gain are derived rather than assumed.

Computed (ours, method stated): the Intel contribution to the quarter's public-equity mark referenced in §1 — Intel's own closes across NVIDIA's February–April quarter (roughly \$45 in late January against roughly \$83–85 in late April) applied to the 214,776,632-share position give about \$8.0–8.6 billion, or some 60–64% of the \$13.4 billion public mark. The mark belongs to the quarter ended April 26 and the 13F to June 30, so the two are computed separately rather than one inferred from the other; as a units check on the 13F, Intel's June 30 close of \$139.63 reproduces the reported \$29,989,261,126 to the dollar. Days sales outstanding and inventory days (formula in §10); the ~23% share of pretax income attributable to investment marks (\$13.4B + \$2.6B against

\$69.9B); the ~6–7% cumulative default probability implied by an ~80bps five-year spread (method in §10); pro-forma total debt and net cash after the June issuance (10-Q balance plus the 8-K's \$25.0B of notes, cash up by the same proceeds).

Press-reported (not independently verified): all CDS levels and dates, including the November 2025 trading start, the late-July record prints (accounts disagree between roughly 82 and 84bps — itself a comment on the market's depth), the Alphabet crossover, and the Oracle/Alphabet/SpaceX records; the SPE/bond structure of the platforms (the releases themselves specify almost nothing structurally); Compute Exchange's founding details and its reported H100/A100-concentrated secondary demand; the Burry quotation.

Reported (Bloomberg, August 14, 2026, and its syndications): the months of prior work by Goldman, Blackstone and Apollo and the slow progress that prompted the public announcement; the "round figure with no obvious provenance" characterization of the \$500 billion and its lack of a time frame; the unnerved-then-eased sequence of credit-investor reaction around the 25% disclosure; the expected chips-plus-offtake collateral and the re-rental mitigant; the Broadcom parallel structure and Blackstone's further \$30 billion sounding; Apollo's origination-and-distribution expansion; the sovereign wealth, pension, insurance and possible retail buyer base; JPMorgan asset management's discussions; and the quoted skepticism on inflated chip valuations and future compute oversupply.

Institutional and sell-side commentary (attributed, not ours): the Bank for International Settlements' Annual Economic Report of June 28, 2026 (circular financing, the >\$1tn hyperscaler capex, the declining net economic surplus, the direct-lending-quadrupled statistic, and the "similarly disruptive"/corporate-credit-freeze warning — all quoted from the report's own text) and BIS Bulletin No. 120 of January 7, 2026 by Aldasoro, Doerr and Rees, including its footnoted \$40bn-in-2025-against-\$3bn-in-2010 origination figure; Bank of America's Vivek Arya (quoted verbatim, with the ~15%-of-~\$470bn-FCF vendor-financing estimate as BofA's, not ours); the reservations attributed to Wells Fargo, Mizuho, UBS and Seaport Global's Jay Goldberg; the reported state of rating-agency treatment of data-center and GPU-backed ABS; and the reported competitive threat to GPU resale values from China's domestic accelerator buildout.

Conjectured (labeled): all of §6. The three-layer ownership claim and the boomerang are stated as speculation and pre-registered against §7's audit points rather than asserted as findings.

§10 Sources & Method

SEC filings. NVIDIA Corp. Form 10-Q for the fiscal quarter ended April 26, 2026, filed May 20, 2026, accession 0001045810-26-000052, retrieved via SEC EDGAR (edgartools), pulled August 13, 2026. NVIDIA Corp. Form 8-K filed June 18, 2026, accession 0001193125-26-275783 (Item 8.01) — the \$25.0 billion senior notes offering in seven tranches maturing 2028–2056, verified tranche-by-tranche; pulled August 14, 2026. Figures read from the notes to the condensed consolidated financial statements: cash equivalents and marketable securities including the lockup disclosures (Note 5); non-marketable securities, the adjustments rollforward, equity-method infrastructure funds, and investment commitments (Note 6); the facility-lease guarantee disclosure; commitments and contingencies including supply, cloud, and other vendor commitments (Note 10); concentration of revenue (direct-customer percentages); and the condensed statements themselves (revenue, pretax income, Other income (expense) net, cash and investments, debt). Prior-year comparatives from the same filing.

Announcement. NVIDIA newsroom and Blackstone press releases, both dated August 10, 2026, retrieved August 13, 2026. Both quotes verbatim from the NVIDIA release. The releases do not themselves describe SPEs, bonds, or tranching; that structural description is attributed to contemporaneous press coverage (Reuters explainer; Benzinga) and is labeled press-reported in §9.

Market data. NVDA closing price \$225.30 and market capitalization ≈\$5.46T as of August 13, 2026 (NASDAQ close, via Financial Modeling Prep). Quarterly income-statement and balance-sheet aggregates used for the DSO/inventory computation from Financial Modeling Prep's standardized statements for the four quarters ended July 2025 through April 2026, pulled August 13, 2026. $DSO = \text{accounts receivable} \div \text{quarterly revenue} \times 91$; $\text{inventory days} = \text{inventory} \div \text{quarterly cost of revenue} \times 91$. S&P 500 weight for NVDA (~8%) as of August 6, 2026, per SPY holdings, carried from the prior note.

CDS. All credit-default-swap levels are press-reported prints (Benzinga, July 2026; Yahoo Finance/Reuters explainer, August 2026; Global Markets Investor, August 2026) — no terminal source was used, and the levels should be treated as approximate. The implied-default arithmetic assumes spread $\approx (1 - R) \times \lambda$ with recovery $R = 40\%$: $\lambda \approx 80\text{bps} \div 0.6 \approx 1.33\%/yr$, cumulative five-year default probability $\approx 1 - e^{(-5\lambda)} \approx 6.4\%$. The chart in §2 connects four dated press prints with straight segments and interpolates nothing between them.

Residual-value support. NVIDIA corporate newsroom/blog post and Jensen Huang's own public statement, both August 11, 2026 — the source of the verbatim phrases quoted in §3 ("up to 25% of an opportunity," "assessed carefully on a project-by-project basis," "substantially lower than other compute-financing arrangements," "complement — not replace — independent underwriting") — first reported by Bloomberg the same morning and corroborated across independent outlets. This is a company statement, not press inference; but it appears in no SEC filing (EDGAR shows no NVDA 8-K between July 2, 2026 and this note's date) and no definitive agreement exists — only memoranda of understanding. The ~\$125B order-of-magnitude figure in §5 is our own arithmetic ($25\% \times \text{the } \$500\text{B platform target}$) and is explicitly not a commitment: the support is per-deal, covers a residual gap rather than a principal balance, and the platform target is itself uncommitted capital. Neither the assumed useful life of the hardware nor the loss order has been stated by any party.

Investment-grade precedent and distribution. CoreWeave's \$8.5B delayed-draw term loan, closed March 31, 2026, rated A3 (Moody's) and A-low (DBRS) — the first GPU-backed facility to reach investment grade; reported. Expectation that the platform securitizations distribute among insurers, pension funds, and sovereign wealth funds: reported (Axios, August 11, 2026, and contemporaneous coverage), and framed in §4 as precedent and expectation rather than as achieved by the NVIDIA platforms, whose rating standard is reported unresolved. Goldman Sachs's role: Reuters, August 14, 2026 ("among the lead underwriters" on the June bond sale, in talks with investors on the platforms); the joint bookrunning of that \$25.0B deal with JPMorgan and Morgan Stanley is separately reported. Goldman's position as the only one of the six with a banking arm is Bloomberg's characterization (see above).

The announcement's provenance and the deal pipeline. Bloomberg, "Nvidia's \$500 Billion Plan Envelops Wall Street in Its AI Frenzy," August 14, 2026; the original is paywalled and was read via the Yahoo Finance syndication of the same piece — the source for the months of prior negotiation, the slow-progress-prompted-publicity account, the "round figure with no obvious provenance" characterization, the credit-investor reaction sequence around the 25% cap, the chips-plus-offtake collateral expectation and re-rental mitigant, the Broadcom/Apollo/Blackstone parallel (20+ gigawatts, ~\$35B in hand), Blackstone's further \$30B+ sounding, Apollo's origination-and-distribution expansion, the buyer base including possible retail, JPMorgan asset management's discussions, and the quoted skepticism. Goldman "the only firm with a dedicated banking arm" is Bloomberg's characterization, not ours (v1.1 carried it as our own inference; corrected here). No structural terms — tranching, ratings, advance rates, tenor, or pricing — are disclosed in any source; the announcement is reported as deliberately vague.

Compute Exchange. compute.exchange (platform descriptions, forward contracts); SiliconANGLE, July 17, 2026 (secondary-market launch, founding details, demand concentration). Press-reported, per §8.

Form 13F-HR. NVIDIA Corp., period ended June 30, 2026, filed August 14, 2026, accession 0001045810-26-000065; figures read from `information_table.xml` on SEC EDGAR, pulled August 14, 2026. Values are as-reported in whole dollars; per-share figures are ours ($\text{value} \div \text{shares}$) and were used to validate the units. 13F-reportable holdings are a different population from the balance-sheet captions "marketable" and "non-

marketable equity securities" used elsewhere in this note, and the two dates differ (June 30 vs April 26) — they are deliberately not presented as comparable.

Institutional and sell-side sources. BIS Annual Economic Report 2026, published June 28, 2026 (chapter on the AI boom and financial stability), and BIS Bulletin No. 120, "Financing the AI boom: from cash flows to debt," Aldasoro, Doerr and Rees, January 7, 2026 — both bis.org, full text extracted and searched directly, August 14, 2026 (see Corrections, §9). Bank of America (Vivek Arya) via contemporaneous coverage, August 2026 — the quotation is verbatim as reported and the free-cash-flow estimate is BofA's. Analyst reservations (Wells Fargo, Mizuho, UBS, Seaport Global) and the announcement-day share-price decline: contemporaneous coverage, August 11–12, 2026. Rating-agency treatment of data-center and GPU-backed ABS, and the reported China domestic-accelerator threat to resale values: contemporaneous trade and press coverage, August 2026, reported rather than independently verified against agency methodology documents.

Prior TON618 notes relied upon: The Losing Bets That Built the World (the four-seats framework, index-concentration figures); The Neocloud Fuse (neocloud credit); The Thin Floor (CoreWeave's intermediation economics); The First Invoice (GPU useful-life debate); The AI Capex Payback Clock (hyperscaler free-cash-flow rollover); The Private Credit Reckoning (the insurance general-account channel); The Round Trip (Oracle). Chart sourcing lives in this appendix by design; charts carry no source captions.

Assumptions and estimates. The ~\$750B/yr hyperscaler datacenter-spend figure is a press-reported industry estimate, not a computation of ours. The 2003–07 originate-to-distribute comparison and the ABX and SIV histories are interpretive frames, sourced from the public record of those episodes. Historical spread references (the 30bps AAA CDO spread) are order-of-magnitude period figures.

§11 Disclosures

Information only. This material is published by TON618 Capital for information purposes only. Nothing here is an offer to sell or a solicitation of an offer to buy any security, fund interest, or digital asset, and nothing here is personalized investment advice or a recommendation regarding any instrument.

Publisher's exclusion. All research is published solely as general, impersonal information of regular circulation. It is not tailored to the objectives or circumstances of any individual and is not issued in connection with compensation from any client. The Fund has no clients and distributes all research free of charge. On that basis it publishes in reliance on the publisher's exclusion from the definition of "investment adviser" under the Investment Advisers Act of 1940 (§202(a)(11)(D); cf. *Lowe v. SEC*, 472 U.S. 181 (1985)).

Registration & conflicts. TON618 Capital is not registered as an investment adviser or broker-dealer in any capacity. The Fund is a Bitcoin fund and may hold or transact in the securities or digital assets it discusses; material conflicts are disclosed where they exist. This note discusses, among others, NVIDIA, Oracle, CoreWeave, Alphabet, Microsoft, Amazon, Meta, Broadcom, Apollo Global Management, BlackRock, Blackstone, Brookfield, Goldman Sachs, and KKR; as of the date of this note the Fund holds **no position, long or short, in any single security named herein**. It maintains a strategic Bitcoin allocation, and to the extent this note bears on liquidity conditions, credit formation, and capital-cycle dynamics that inform that allocation, the reader should weigh that as a material interest. The Fund receives no compensation from any party in connection with its research.

Use of AI. Artificial intelligence is used in the creation of this research. **A material conflict is disclosed here: this note analyzes the financing of AI compute infrastructure, and the AI tooling used to produce it is made by a company that is among the larger buyers of the compute whose financing this note discusses.** All methodology and data integrity are reviewed and approved before publication by TON618 Capital's Chief Investment Officer, **Keyth Beck**; errors may nonetheless occur, and readers should verify independently.

CFA. This report was prepared to align with CFA Institute analytical standards (methodology only). CFA® and Chartered Financial Analyst® are registered trademarks owned by CFA Institute. That reference describes the analytical framework applied; it does not imply the report was prepared, reviewed, or authored by a CFA charterholder, and the report is not issued, reviewed, endorsed, certified, or approved by — nor affiliated with — CFA Institute.

Risk & feedback. Past performance is not indicative of future results. Digital assets and equities are volatile and may result in total loss of capital. Corrections and feedback are welcome — please direct them to CIO **Keyth Beck** at keyth@ton618capital.com.

Version 1.4 · published August 14, 2026; revised the same day by an editorial pass for structure and cohesion, and before that three times on substance: v1.1 incorporated NVIDIA's disclosed 25% residual-value support (§3, §5, §6); v1.2 added the announcement's reported provenance and corrected our own framing of why the cap was published (§3, §6); v1.3 adds NVIDIA's 13F-HR filed the same morning (§1), the BIS systemic frame quoted from primary text, the rating-agency and China context, and the strongest published case against this note's reading, with our answer; v1.4 restructured — the systemic frame moved to §3 where it sets up the machine, the counter-case given its own section (§7) ahead of the audit points, and §1 returned to being a verdict · a follow-up to The Losing Bets That Built the World (August 12, 2026) · independently claim-verified before publication; the v1.1 omission is disclosed in §9.