




By: TA Analysis

# Who will pay when AI chips lose value?



Amazon is investing more in artificial intelligence this year than its business generates, so it is looking for ways to shift some of that burden to investors. The company is negotiating the sale and leaseback of about \$8 billion of **Grace Blackwell** chips from the US manufacturer Nvidia.

The sum is modest compared with Amazon's total investment, but the transaction raises the question of who bears the loss when expensive computing equipment loses value before it pays for itself.

The answer depends on what the wider AI sector will pay for the capital it borrows to fund further expansion.

The chips are already operating in more than a dozen Amazon data centres in the United States.

Under the arrangement, Amazon would sell them to a special-purpose entity established solely for the transaction, then lease them back for continued use.

That entity would fund the purchase with capital borrowed from investors. Because the lease payments would be made by Amazon, which carries an **AA credit rating** – the second highest awarded by rating agencies – the entity's debt is expected to secure a comparable rating.

Insurers and pension funds, which primarily acquire low-risk bonds, would then also be able to invest. Final terms remain undisclosed and Amazon has declined to comment.

## Amazon's capital expenditure exceeds its earnings

Amazon has raised its projected capital expenditure for 2026 to about \$220 billion, most of which is allocated to data centres and chips.

**Amazon Web Services** (AWS), the division that leases computing capacity to enterprise

clients over the internet, increased its revenue by 37 per cent in the second quarter.

Customers have already contracted **services** worth \$496 billion that have yet to be delivered.

However, infrastructure build-out costs exceed incoming revenue. Over the past 12 months, after funding its capital investments, Amazon recorded **negative cash flow** of \$7.6 billion for the first time since 2023, as equipment investment over the year increased by \$66.1 billion.

### By selling the chips, Amazon secures capital without issuing new bonds

The timeline of Amazon's decisions also indicates the immediate rationale for the transaction.

The company issued \$25 billion in **bonds** on 7 July, stating that the offering met its dollar-denominated borrowing requirements for the year.

Investor demand was weaker than in March; consequently, Amazon had to offer a **premium** of 0.18 to 0.21 percentage points on the longest-dated maturities.

Three weeks later, it increased its **investment target** from \$200 billion to \$220 billion.

By selling the chips, Amazon secures capital without issuing new bonds, having already signalled to the market that it does not need them this year.

However, its commitments do not vanish: under **US accounting standards**, the lease is classified as a liability, and **credit rating agencies** treat it as debt.

Amazon secures immediate liquidity while attempting to transfer the risk of accelerated chip obsolescence to investors.

## Nvidia builds a market for chip financing

Amazon's transaction is part of a broader effort to finance computing hardware as conventional infrastructure.

In August, **Nvidia signed** letters of intent with investment houses Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs and KKR to establish entities expected to raise more than \$500 billion to fund AI infrastructure. Definitive **agreements** have not yet been executed.

The objective is for pension funds, insurers and private credit funds to finance chips in the same way they have financed power stations, aircraft and telecommunications networks for decades.

### The distinction lies in asset lifespan

The distinction lies in asset lifespan. A power station or an aircraft depreciates gradually over decades.

An AI processor, by contrast, can operate without fault yet lose most of its economic value within three years, displaced by a successor architecture that delivers multiples of the throughput for the same energy consumption.

Nvidia is already introducing its next-generation platform, **Vera Rubin**. The older chip is then relegated to lighter workloads, such as running pre-trained models, but its yield per hour contracts far more quickly than creditors anticipated when financing the purchase.

## Nvidia counts on ten years, banks on three

Nvidia argues that premium chips can generate income for up to a decade.

**Commercial banks**, by contrast, typically underwrite them over a three- to four-year lifespan, and prospective lenders earmarked for Nvidia's infrastructure funds are already pressing for more robust guarantees than the company has offered.

S&P Global acknowledges that legacy processors have operated well beyond five years, but applies conservative valuations when modelling their worth.

Nvidia is not a disinterested party in this debate. Beyond supplying chips, it actively arranges customer financing.

The company holds **strategic equity positions** valued at \$99 billion in third parties, including cornerstone chip customers OpenAI and CoreWeave.

It has entered into commitments to purchase computing capacity from selected clients, effectively enhancing their credit profile when they secure third-party debt to buy its processors.

In Ohio, Nvidia has underwritten **guarantees** capped at \$105 billion to back campus development and power infrastructure leased by OpenAI.

For the chips within the new funds, however, its commitment is materially lower, underwriting residual values of only up to 25 per cent of individual transactions.

The more readily banks and institutional investors accept elongated depreciation cycles, the lower Nvidia's contingent liabilities will remain and the faster it can clear inventory.



*If Amazon enters into a long-dated payment schedule, the investor is essentially extending credit to Amazon, rendering the chip value immaterial*

Amazon's internal estimate contradicts Nvidia's. Beginning in January 2025, the company amended its accounting policies to assume that certain servers and networking infrastructure have a useful life of five years rather than six, citing the accelerating pace of AI development.

That change alone increased its 2025 depreciation and amortisation charge by \$1.4 billion. In projecting to private credit investors that the chips should last at least five years, Amazon is simply applying its internal accounting assumptions.

Consequently, any long-term investor assuming a longer life cycle is relying on an assumption more optimistic than Amazon's own audited disclosures.

Investor risk is increased by Amazon developing its own processors: Trainium for AI workloads and Graviton for general-purpose compute, a unit that already generates more than \$25 billion in annual revenue.

Upon lease expiry, Amazon retains the flexibility to transition to next-generation Nvidia hardware or its own chips, leaving the owner holding physical equipment that cannot readily be sold at the contracted value.

For this reason, the lease terms are far more critical than the \$8 billion transaction value.

If Amazon enters into a long-dated payment schedule, the investor is essentially extending

credit to Amazon, rendering the chip value immaterial.

If the lease is shorter, includes early-termination options, or ties capital recovery to secondary-market chip values, residual obsolescence becomes an immediate credit risk.

## Meta singles out the building, Amazon singles out the chips

Meta has formed a joint venture with investment manager Blue Owl for the Hyperion data centre campus in Louisiana, valued at around \$27 billion.

The entity's bonds secured an **A+ rating** from S&P Global, underpinned by Meta's financial backing, with the liabilities remaining off **Meta's balance sheet**.

Meta, however, carved out the physical facilities, designed to last for decades, into a separate entity.

**The capital cost of developing AI infrastructure will escalate well before demand for artificial intelligence begins to weaken**

Closer to Amazon's structure is Elon Musk's xAI transaction from October 2025, in which a **special-purpose entity** acquired Nvidia processors and leased them to xAI for five years.

What distinguishes Amazon's proposition is that rapidly depreciating equipment is being packaged for insurers and pension funds, underwritten by one of the world's highest-rated companies.

This is unfolding against a backdrop of tightening credit conditions. Amazon, Alphabet, Meta and Oracle had issued approximately \$194 billion in **corporate bonds** by 7 July – a 79 per cent surge compared with the whole of 2025.

According to data from the investment company Apollo, in February there were nearly five dollars in orders for every dollar of those bonds offered, whereas in July there were fewer than two.

If investors insist on wider credit spreads, stronger guarantees and shorter deadlines, the capital cost of developing AI infrastructure will escalate well before demand for artificial intelligence begins to weaken.

## Risk emerges when the shortage ends

While the shortage of computing capacity persists, legacy chips continue to generate cash flow.

Amazon has announced that from 7 October it will increase, by around 15 per cent, the tariff for Nvidia chips leased to customers, spanning the 2020 A100 architecture through to the current B300.

The market value of Blackwell chips will therefore not depreciate immediately upon the debut of a successor generation.

Instead, depreciation will accelerate once structural shortages ease, whether driven by a deceleration in demand or by secondary capacity catching up with market requirements.

In that scenario, asset owners face a dual compression: deteriorating cash yields alongside a collapse in residual equipment value. Given this, commercial lenders underwrite chip lifespans over considerably more conservative horizons than Nvidia.

**Amazon's high corporate credit rating ensures that a portion of this risk is transferred to insurers and pension funds**

Amazon's high corporate credit rating ensures that a portion of this risk is transferred to

insurers and pension funds.

Apollo, one of Nvidia's six funding partners, controls the major life insurer Athene.

As early as April, the **Bank of England** cautioned that technology conglomerates are financing AI capital expenditure less through retained earnings and equity, and more through rapid, highly leveraged and complex borrowing structures. A subsequent collapse in chip valuations would therefore feed directly through to pension savings.

## Nvidia will have to guarantee the value of its own chips


The fundamental problem with this model is the difference in the rate at which components of the AI infrastructure age.

The data centre building, connections and electrical substations can last for decades, whereas the servers housed within the same facility must be replaced every few years.

Debt capital markets will therefore have to finance them separately. Physical structures and power infrastructure will attract low-cost, long-tenor debt – as demonstrated by Meta – while chip financing will incur wider credit spreads, shorter maturities and mandatory guarantees.

The Amazon transaction is likely to proceed, but structured so that the bulk of the residual exposure remains on Amazon's balance sheet.

A separate entity's debt can only achieve an investment-grade rating if institutional investors underwrite Amazon's corporate covenant rather than take direct risk on accelerated hardware obsolescence.



*Nvidia will be forced to increase its residual value guarantees beyond the current 25 per cent threshold per transaction - Jensen Huang*

The lease will therefore require a tenor long enough for Amazon to amortise most of the equipment's capital cost through its payments.

Comparable transactions over the next 12 months should be expected from capital-constrained operators, notably Oracle and smaller companies leasing computing capacity.

Without Amazon's pristine balance sheet rating, these entities will face considerably higher financing costs.

The biggest shift will occur at Nvidia. As prospective lenders demand more stringent risk mitigation, Nvidia will be forced to increase its residual value guarantees beyond the current 25 per cent threshold per transaction.

The prime corporate beneficiary of the AI build-out will thus gradually absorb a larger share of the loss should its chips lose value.

Amazon's transaction demonstrates that AI capital expenditure increasingly relies on the thesis that chips generate returns over a horizon of five years or longer – an assumption validated to date only under conditions of acute supply deficit.

The decisive test will arrive at the turn of the decade, when peak-cycle leases mature alongside multiple successive architectures displacing Blackwell on the secondary market.

Until then, lease tenors, the scope of

guarantees, and secondary liquidation prices for refurbished chips will be far more informative barometers of sector viability than aggregate unit shipments.

If chips depreciate faster than debtors have modelled, credit losses will be borne by the debt markets that funded the build-out, driving up the cost of capital for subsequent investment cycles.

Those acquiring chip-backed debt today are, in substance, wagering that the capacity deficit will outlast the lease.