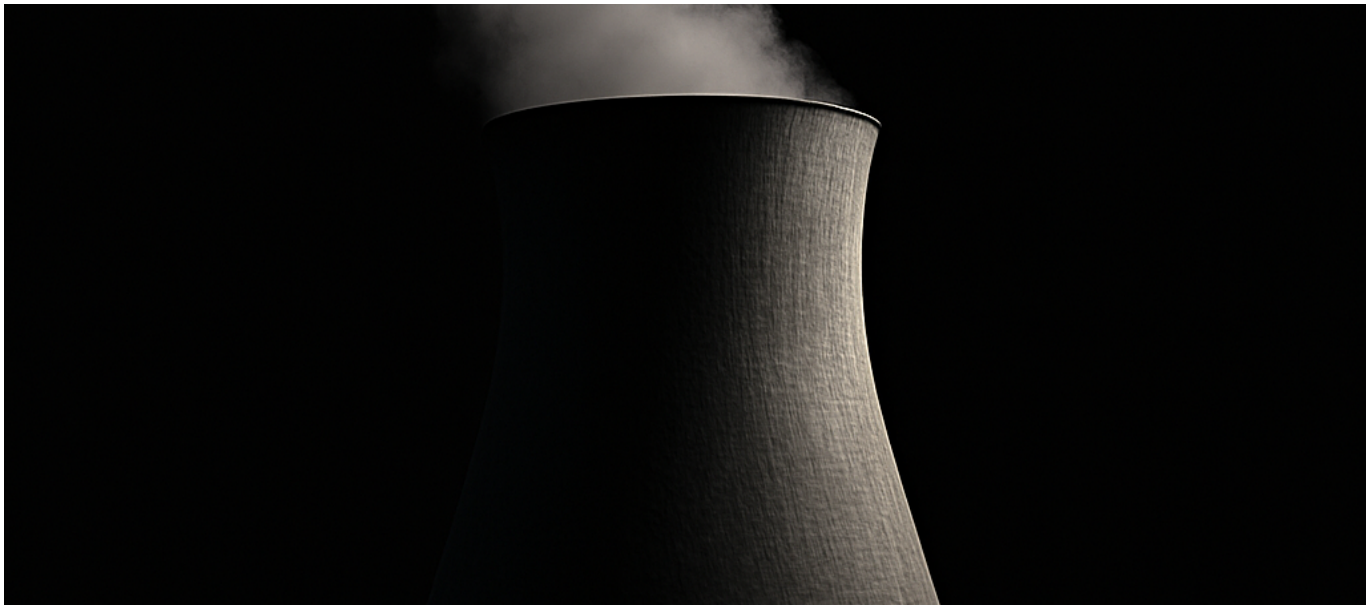

Apollo and Blackstone Close $35 Billion Chip Financing for Anthropic—Largest Private Credit Deal Ever Funds 1+ GW Compute Expansion Off Balance Sheet

Anthropic just secured $35 billion in compute capacity without adding a dollar of debt to its balance sheet. Wall Street has arrived in AI infrastructure, and the playbook looks like aircraft leasing.

The Deal: $35 Billion in Private Credit, Zero Balance Sheet Debt

On June 9-10, 2026, Apollo Global Management and Blackstone finalized a $35 billion private credit package—the largest private financing ever executed. The money doesn’t go to Anthropic directly. Instead, it funds a special-purpose vehicle

(SPV) that purchases Google TPUs, which Anthropic then leases.

The structure breaks down into three tranches with distinct risk profiles:

- **$6 billion Senior A1:** Treasury rate plus 100 basis points—essentially government-adjacent pricing for AI chips
- **$24 billion A2:** 5.75% coupon, the bulk of the financing
- **$4.5 billion Class B notes:** 8.5% coupon for investors willing to take subordinated risk

Apollo committed an additional $800 million in equity through its Atlas SP Partners vehicle, giving the firm skin in the game beyond just collecting fees.

The financing supports more than 1 GW of compute capacity starting mid-2026, according to reporting from Axios. For context, that's roughly the power consumption of a city with 750,000 homes—dedicated entirely to training and running AI models.

Broadcom provides the structural glue that makes this deal work: credit endorsement and residual value guarantees on the $30 billion A1 and A2 tranches. Translation: if Anthropic walks away from the lease or the TPUs become worthless, Broadcom is on the hook, not Apollo or Blackstone's investors.

This deal closed just weeks after Anthropic completed a $65 billion Series H round at a $965 billion post-money valuation. The company raised $100 billion in capital—$65 billion equity, $35 billion in SPV-financed compute—without materially changing its debt profile.

Why SPV Financing Changes the AI Infrastructure Game

The aircraft leasing comparison isn't just clever framing. It's literally the same financial engineering.

When airlines need 737s, they rarely buy them outright. Instead, leasing companies purchase the planes and rent them to carriers. The airline gets the asset, the lessors get steady returns, and the airline's balance sheet stays clean enough to borrow more money or go public.

Anthropic just applied this model to TPUs. The implications ripple across three dimensions: capital efficiency, competitive dynamics, and market structure.

Capital Efficiency: Turning CapEx into OpEx

Traditional AI infrastructure spending forces a brutal choice: raise enormous equity rounds that dilute founders, or take on debt that spooks investors and limits future borrowing capacity. Anthropic found a third path.

By keeping $35 billion in hardware off its balance sheet, Anthropic preserves its debt-to-equity ratios while accessing compute that would otherwise require massive capital commitments. The lease payments show up as operating expenses—significant, but predictable and manageable.

For a company reportedly considering an IPO, this matters enormously. Public market investors scrutinize balance sheets obsessively. A company with $35 billion in debt looks fundamentally different than one with $35 billion in lease obligations, even if the economic substance is similar. Accounting rules, specifically ASC 842, require lease disclosure, but the optics differ.

Competitive Dynamics: Who Can Play This Game?

This deal required three things most AI companies lack: a $965 billion valuation to attract Wall Street’s attention, a strategic partner in Broadcom willing to guarantee residual values, and the negotiating leverage to secure near-Treasury rates on senior tranches.

OpenAI, with Microsoft’s backing, can likely replicate this structure. Mistral, Cohere, and the next wave of foundation model companies cannot—at least not at this scale or cost of capital.

The financing creates a two-tier market. Companies with access to SPV-style compute financing can scale infrastructure without dilution or debt constraints. Everyone else faces the traditional CapEx wall.

> The largest private credit deal in history wasn’t for real estate, energy, or infrastructure. It was for AI chips. That reordering of capital allocation priorities tells you where the economy is heading.

Market Structure: Private Credit Enters AI

Apollo and Blackstone managing $35 billion in AI chip financing signals a structural shift in how compute infrastructure gets funded. Private credit has spent the past decade eating market share from traditional banks in leveraged buyouts, real estate, and corporate lending. Now it's coming for AI.

The economics work for everyone involved. Institutional investors in Apollo and Blackstone's funds want yield in a world where Treasury returns feel inadequate. AI companies want compute without dilution. Chip manufacturers want guaranteed purchase volumes. The SPV structure aligns all three.

Expect more deals. Bloomberg's reporting suggests the Broadcom AI XPV Platform targets over 20 GW of compute capacity by 2028—implying potential financing needs that dwarf even this record-breaking deal.

Technical Architecture: What 1+ GW of Compute Actually Means

Let's ground the financial engineering in physical reality. One gigawatt of power dedicated to AI compute represents a step-function increase in training capacity.

Power and Cooling at Scale

Modern data center efficiency runs around 1.2-1.4 PUE (Power Usage Effectiveness), meaning 1 GW of compute power requires roughly 1.2-1.4 GW of total facility power when cooling and overhead are included. The TPUs themselves consume most of that, but thermal management at this scale requires industrial-grade infrastructure.

Google's TPU v5p clusters run at about 400W per chip. One gigawatt of compute power translates to roughly 2.5 million TPU chips, though actual configurations depend on interconnect, cooling, and redundancy requirements. The real number is probably lower given infrastructure overhead, but the order of magnitude illustrates the scale.

Training Implications

Current frontier models—GPT-4 class and beyond—require training runs measured

in tens of millions of GPU/TPU hours. With 1+ GW of dedicated compute, Anthropic can run multiple frontier training experiments simultaneously, accelerating their research cadence without the scheduling constraints smaller compute pools impose.

More importantly, this capacity supports the inference demands of a consumer-scale AI product. Claude's API usage has grown exponentially. Each query requires compute. Dedicated gigawatt-scale infrastructure ensures Anthropic can scale inference without competing for resources against training workloads.

The TPU Choice

Anthropic's selection of Google TPUs over NVIDIA GPUs reflects both economic and technical calculations. Google offers TPU access at competitive rates to counter NVIDIA's dominance. The resulting pricing makes TPU-based training economically attractive for organizations willing to adapt their software stack.

From a technical standpoint, TPUs excel at the dense matrix operations that dominate transformer training. Their systolic array architecture processes large batch sizes efficiently, and Google's XLA compiler has matured enough to support complex training workflows. Anthropic's engineering team has clearly invested in TPU optimization—a sunk cost that makes continued TPU usage rational even if NVIDIA's newest GPUs offer theoretical performance advantages.

The Broadcom involvement adds another layer. Broadcom manufactures custom silicon and networking equipment that integrates with hyperscale compute deployments. Their residual value guarantees suggest confidence that TPUs will retain meaningful value over the lease term, either through continued AI workloads or secondary market sales.

The Contrarian Take: What the Coverage Gets Wrong

Most reporting on this deal focuses on the record size and clever financial engineering. That framing misses the more significant dynamics.

Overhyped: Balance Sheet Cleanliness

The “off-balance-sheet” framing makes great headlines but overstates the accounting benefit. Under current lease accounting standards (ASC 842 in the US, IFRS 16 internationally), operating leases still appear on balance sheets as right-of-use assets and lease liabilities. Anthropic won’t show zero impact from this deal.

The real benefit isn’t accounting treatment—it’s capital structure flexibility. Lease obligations differ from traditional debt in covenants, terms, and investor perception. Anthropic gains compute access without the restrictive covenants typical of large corporate loans, preserving operational flexibility.

Underhyped: The Broadcom Guarantees

Buried in the deal structure is the detail that matters most: Broadcom guarantees residual values on $30 billion of the financing. This isn’t a small commitment. It means Broadcom believes TPUs will hold significant value at lease termination, and they’re willing to backstop that belief with their own balance sheet.

Why would Broadcom do this? Several possibilities: they expect secondary markets for AI chips to develop, they want the volume commitments to justify their own manufacturing investments, or they see strategic value in being the infrastructure partner for AI’s capital formation.

Whatever the motivation, Broadcom’s guarantee enabled the near-Treasury pricing on the senior tranches. Without it, Apollo and Blackstone would have demanded much higher yields to compensate for chip depreciation risk. The guarantee effectively subsidizes Anthropic’s cost of compute.

Underhyped: Regulatory Risk

The deal “raised regulatory concerns” according to multiple reports, but coverage glosses over what those concerns actually involve. Several angles deserve scrutiny:

Off-balance-sheet treatment: Post-Enron accounting reforms specifically targeted SPV structures that hid obligations. While this deal appears to comply with current standards, aggressive use of SPVs to separate compute liabilities from operating companies could invite regulatory attention—especially if the structure becomes widespread.

Systemic risk: $35 billion in AI chip financing from two private equity firms concentrates risk in ways regulators typically monitor. If AI demand drops or the models underperform, the cascade effects through private credit markets could be substantial.

Antitrust implications: A financing structure that only companies valued at nearly $1 trillion can access raises competition concerns. If SPV-style compute financing becomes essential for frontier AI development, barriers to entry grow substantially.

None of these risks are immediate, but they suggest the deal's structure faces more scrutiny ahead than current coverage acknowledges.

Practical Implications: What This Means for Your Organization

Unless you're running a frontier AI lab, you won't replicate this deal. But the underlying dynamics create second-order effects worth considering.

For AI Startups: Compute Financing Is Coming

This deal proves institutional capital is hungry for AI infrastructure exposure. Smaller versions of this structure—$100 million, $500 million, $1 billion—will emerge for mid-market AI companies. If your startup faces compute constraints, begin conversations with lenders who understand AI workloads. The financing options available in 12 months will look different than today.

Practical step: Document your compute utilization, unit economics per GPU hour, and revenue per compute dollar. Lenders will want these metrics.

For Enterprise AI Teams: Lease vs. Buy Analysis Matters

The Anthropic deal validates leasing as a viable model for AI infrastructure. Enterprise teams evaluating cloud vs. on-premise vs. colocation should add structured financing to their analysis. The total cost of ownership calculation changes when you can finance hardware through third-party SPVs rather than CapEx budgets.

Practical step: Request lease financing options from hardware vendors and

hyperscalers. Several offer programs modeled on this structure at smaller scales.

For Infrastructure Investors: AI Hardware Has Entered Asset-Backed Finance

TPUs, GPUs, and AI-optimized chips now have the characteristics that attract asset-backed investors: predictable cash flows (from lease payments), quantifiable residual values (with the right guarantees), and essential-use status (AI companies need compute to survive).

This asset class didn't exist two years ago. By 2028, AI hardware-backed securities could rival aircraft or equipment leasing in market size.

Practical step: If you manage institutional capital, engage with the private credit desks at Apollo, Blackstone, and competing firms. AI infrastructure financing is where they're deploying resources.

For Technical Leaders: Understand the Capacity Implications

Anthropic's 1+ GW expansion means their model capabilities will accelerate. Claude 4 or whatever follows will benefit from training runs impossible at smaller compute scales. Plan your AI integration strategies assuming frontier model capabilities improve faster than the 2023-2025 period suggested.

Practical step: Build abstraction layers in your AI-dependent systems. When Claude's capabilities improve, your architecture should let you leverage those improvements without rewriting integration code.

Forward Look: Where This Leads in 6-12 Months

More Deals, Larger Deals

The Broadcom AI XPV Platform targets 20+ GW by 2028. Anthropic's 1+ GW represents roughly 5% of that target. Expect announcements of additional financing tranches—potentially $50-100 billion across multiple AI companies over the next year.

Microsoft and OpenAI will likely announce a similar structure. Google and DeepMind already have internal compute access, but Google Cloud's external customers may

see SPV-style financing options emerge. Amazon has the balance sheet to finance AI infrastructure directly, but AWS enterprise customers will demand competitive financing terms.

Secondary Markets for AI Chips

Broadcom's residual value guarantees only make sense if secondary markets for used AI chips develop. Those markets require price discovery, standardization, and buyers willing to operate non-current-generation hardware.

Watch for announcements of AI chip marketplaces, refurbishment programs, or cloud providers offering discounted compute on older hardware. The infrastructure for chip resale is a necessary complement to the financing structures now emerging.

Regulatory Clarification

The SEC, FASB, or international equivalents will likely issue guidance on SPV-style AI infrastructure financing within the next year. That guidance will either validate the current approach or force structural modifications that change deal economics.

Companies planning similar financing should build flexibility into their structures. The rules may shift.

Model Economics Improvement

More compute at lower effective cost means Anthropic can offer better model performance at competitive prices. Expect API pricing pressure across the foundation model market as companies with advantaged infrastructure compete on cost.

For enterprises, this is unambiguously good news. For startups building on foundation model APIs, it validates the "don't train your own model" approach—the leaders are investing billions precisely to ensure their models remain superior to anything a startup could build internally.

The Structural Transformation Underway

Step back from the deal specifics and recognize what this moment represents. The largest private financing in history funded AI infrastructure. Not real estate, not energy, not traditional technology infrastructure—AI chips.

Capital flows where returns exist. $35 billion in private credit flowing to AI compute represents a massive vote of confidence in the economic returns from frontier AI development. Apollo and Blackstone have armies of analysts modeling these investments. Their willingness to deploy at this scale reflects genuine conviction, not hype.

The structure—SPVs, tranched debt, residual guarantees—brings Wall Street's financial engineering sophistication to AI infrastructure. This is how mature industries finance capital-intensive growth. Its arrival in AI signals the industry's transition from venture-backed experimentation to industrial-scale capital formation.

For technical leaders, the implication is clear: frontier AI capabilities will accelerate. The compute constraints that bounded model development are loosening. Assume the models available in 2027 are substantially more capable than today's, and build systems flexible enough to leverage those improvements.

For finance professionals, a new asset class has emerged. AI infrastructure-backed securities will need analysis, trading, and risk management. The skills developed in aircraft leasing, equipment finance, and structured credit apply directly.

For everyone watching AI's trajectory, this deal marks an inflection point. When Wall Street finances an industry at this scale, that industry stops being speculative and becomes essential infrastructure. AI has crossed that threshold.

The largest private financing in history went to AI chips—that's not a tech story, it's a signal that AI compute is becoming core economic infrastructure on par with energy and transportation.