From hardware to asset class: unpacking GPU financing

By John Hwang, Derek Poon and Caleb Adebayo, Linklaters LLP*

GPUs have emerged as a nascent asset class in the AI buildout; quickly evolving from a mass-deployed component of data centres to standalone collateral against which banks and private capital are willing to finance. GPUs – or graphics processing units – have been around since 1999, when Nvidia – then a mid-sized graphics chipmaker – launched the GeForce 256 to reshape the computing and gaming industry.

Introduction

For the past few years, data centres have been the predominant focus of AI infrastructure financing, and that asset class has experienced an unprecedented wave of growth and financing primarily driven by AI development. Nvidia's guarantee of up to US$105bn for the Ohio data centre campus housing OpenAI as anchor tenant in August 2026 and Vantage Data Centres' record US$13bn incremental funding in FY 2024, to name a few, are illustrative. As data centre debt has matured into an established, rated market, GPUs have begun to evolve from prosaic hardware into an asset class of their own.

The combination of surging GPU demand, the surging cost of GPUs themselves, and deep pools of capital that hyperscalers, private credit funds, insurers and banks have made available to finance GPUs and other AI accelerators has prompted market participants to look to develop financing structures tailored to the distinctive features of the asset class.

This article unpacks the financing mechanisms that the market has seen in the last few years and examines other key considerations for investors and other market participants seeking a piece of the "GPU action".


John Hwang


Derek Poon


Caleb Adebayo

John Hwang, Partner
Email: john.hwang@linklaters.com
Tel: +1 212 903 9190

Derek Poon, Partner
Email: derek.poon@linklaters.com
Tel: +1 212 903 9180

Caleb Adebayo, Associate
Email: caleb.adebayo@linklaters.com
Tel: +1 212 903 9309

GPUs, TPUs and the AI buildout

GPUs are a key component in the AI buildout, serving as the workhorses behind Large Language Model (LLM) training and real-time inference. Having started out as a gaming technology in the early 2000s, their ability to simultaneously run thousands of mathematical processes to render graphics quickly made them relevant to general-purpose computing, as the graphics demands of desktop computing increased.

Through the remainder of the early 2000s, GPUs continued to be seen primarily as graphics chips, and their general-purpose compute potential was underappreciated. That began to change in 2007, when NVIDIA released CUDA and researchers realised that the same parallel processing capabilities that had previously been deployed for graphics rendering could equally be useful for scientific and machine-learning workloads.

In 2012, the AlexNet image-recognition model, trained on two Nvidia gaming GPUs, showed that GPUs could make deep learning practical. The transformer architecture and, eventually, ChatGPT built on that foundation, in 2017 and 2022 respectively. Today, they are a central component of the global AI competition and demand for them has soared as a result.

Tensor Processing Units (TPUs) are a closely associated, but different, product. Given the huge cost of AI infrastructure and the increasing scarcity of space and power to fuel it, the industry moved from merely demanding raw computational power for AI models and handling large data sets, towards demanding efficiency, performance density and sustained throughput.

Although Nvidia's data centre GPUs have included AI-specific Tensor Cores since 2017, GPUs still retain their graphics-oriented designs. That's where TPUs come in: in 2016, Google began publicly deploying TPUs as proprietary, application-specific integrated circuits (ASICs), with particular attention to deep learning.

TPUs strip out most of the overhead general-purpose features of GPUs to focus purely on matrix multiplication and deterministic execution – two of the key workflows required for developing and interrogating AI models. Although TPUs remain far less widely deployed than GPUs, they are increasingly being used side by side with GPUs and are beginning to run outside Google-only infrastructure.

From silicon to asset class

J.P. Morgan estimates that approximately US$3 trillion of financing could be required for GPUs and custom AI accelerators over the next five years[1]. From the US$2.3bn Magnetar Capital/Blackstone GPU financing obtained by CoreWeave to the US$500m Macquarie-led GPU financing obtained by Lambda and the US$35bn TPU financing by Apollo and Blackstone for lease by Anthropic, large amounts of capital have been used to finance GPUs and TPUs in the last few years. In many cases, these chips and the contracts through which their benefits are provided, are being used as the primary basis on which such financing rests.

Prior to 2022, GPUs did not exist as standalone collateral or independent asset class. Today, the GPU market is valued at more than US$100bn[2] with a compound annual growth rate of roughly 26% to 32%[3]. So, what led this shift from simple IT hardware to financial asset class?

One driver was that demand for these chips far exceeded supply. The launch of ChatGPT in late 2022 and the immediate hyperscaler capex response created a supply-demand imbalance for GPUs. A new category of neocloud operators then emerged, and an entire business model evolved, which was buying GPUs and renting them out on medium term contracts. Additionally, GPUs are expensive – an Nvidia H100 at launch cost anywhere between US$25,000 and US$40,000, while its more recent Blackwell GPUs (B300 Ultras) cost anywhere from US$38,000 to US$45,000 per chip – creating an incentive for the industry to address the financing need for these assets.

This demand for GPUs and their steep price point happened to coincide with the growth of a private credit market and the development of complex financing

structures able to address the structural complexities and bespoke nature of this hardware.

TPUs have equally attracted significant financing, but these are less common compared to GPUs. TPUs' close integration with Google's software stack, together with their historical availability through Google Cloud on a rental basis, has shaped the particular secured-financing structures typically built around them, although, as noted above, deployments outside Google-controlled infrastructure are now beginning to appear.

The methods for GPU financing weren't entirely novel when they kicked off in 2023; the equipment finance and aviation finance worlds had utilised these methods for years, but what didn't exist was financing at that scale that relied on an asset class with shorter lifecycles, higher depreciation schedules and lower liquidity, and the customer contracts for those assets. Unlike what had been financed by these methods prior, these chips had less of a mature resale market and had a less certain residual value.

They are also not standalone equipment. Rather, they are installed in a cluster inside a specific rack, often in a third-party data centre campus. This means that if, for instance, the data centre loses power, has its cooling systems fail or is not energised on schedule because of grid interconnection delays, the GPUs cannot run, providing a different risk profile for contracts predicated on these GPUs.

These are some of the considerations that have underscored the various approaches in financing and underwriting this asset class. As demand for GPUs and GPU as a service (GPUaaS) has continued to grow, investors, underwriters and originators alike have explored a combination of different secured lending and structured finance concepts to properly solve for the risks inherent in the asset class.

Financing mechanisms for GPUs/TPUs

GPU-backed term loans

In 2023, Crusoe, a company that designs and builds AI infrastructure, borrowed US$200m from Upper90 to buy thousands of Nvidia H100 and A100 GPUs.[4] While Crusoe's business model is to provide GPUaaS infrastructure, the loan was an asset-backed loan, backed by the chips themselves and their related cashflows, with Upper90 relying on the Nvidia GPUs as primary collateral.

In August 2025, Lambda used a similar financing structure for its US$275m facility collateralised by GPU fleet and long-term offtake contracts, which it upsized to US$1bn in May 2026, although this time, the customer contracts for the GPUs' capacity were added to the collateral pool alongside the chips.

Under this financing structure, the borrower pledges a pool of GPU (or TPU) hardware (and typically, the associated contracts as well) as collateral to secure a fixed principal amount advanced at closing. The borrowing could be done by a bankruptcy-remote SPV or on the books of the operating company and operates like any other ABL in the market. The collateral pool will usually consist of the hardware, customer contracts, deposit accounts and any insurance proceeds. Revenue from the customer contracts will usually be paid into a collection account controlled by the collateral agent and subject to a deposit account control agreement.

The contracts that back these (and many other) GPU loans routinely include take-or-pay (or "must-pay") obligations designed to generate scheduled cash flow regardless of the customer's actual utilisation. The practical implication is that, although the GPUs are the equipment collateral on paper and backstop recovery on enforcement, the credit that is actually underwritten is the stream of contracted take-or-pay cash flows from creditworthy counterparties.

Delayed Draw Term Loans (DDTLs)

ABLs in the form of DDTLs (whether from financial institutions or private credit lenders) have become a more common structure for GPU financings since CoreWeave's US$2.3bn GPU DDTL from Magnetar and Blackstone in 2023, and many others have followed since.

As long, uncertain lead times are a common characteristic of GPU production and delivery, DDTLs are often a more

well-suited alternative than a single draw straight term loan that fully funds at closing. Financing documents will typically contain borrower covenants that the collateral will be properly perfected, however for many of these GPUs, delivery times can take up to 12 months for volume orders[5] and in some cases they are not installed or are not yet at a known location. DDTLs usually serve well for this timeline.

Each draw under a GPU DDTL is conditioned on satisfaction of a cumulative set of conditions precedent, including that the specific batch of chips being funded must have been delivered, installed and placed in service (and, where the master services agreement so requires, formally accepted by the customer); the lender's security interest in the newly funded collateral must have been perfected through supplemental UCC filings and any related control or access arrangements; no default or event of default may be continuing; and, where the facility has one, the borrowing base or advance-rate test must be satisfied.

Facility amortisation is sized to pay down the drawn principal inside the contract term and within the economic useful life and obsolescence window of the chips so the loan does not have to rely on residual GPU value at maturity. CoreWeave's August 2026 US$2.6bn DDTL 5.5 facility, whose credit agreement has been publicly filed,[6] departed from this template however and carried an approximate five-year maturity against underlying customer contracts averaging approximately three years, with lenders expressly underwriting the renewal risk on the back half of the loan[7].

Sale-leasebacks

A GPU sale-leaseback is another lease financing mechanism being used for GPUs. As with typical sale-leaseback structures used in other asset class financings, the owner (the seller-lessee) of the GPUs sells them to the financing counterparty (the buyer-lessor) with a simultaneous lease by the buyer-lessor of the GPUs to the seller-lessee under a long-term lease agreement. An example of this structure is the September 2025 US$1.5bn Nvidia deal with Lambda to lease approximately 18,000 GPUs over four years.[8]

Sale-leasebacks are a useful method of financing GPUs especially where the seller needs working capital and still wants to continue to make use of the chips. Upon the expiry of the lease, the seller-lessee would typically have the option of purchasing the GPUs at their fair market value or the lessee can elect to renew the lease for an additional term, typically at a renegotiated rental fee, given the depreciation of the chips that will have occurred.

European lease-receivables ABS (G-SPV model)

This is largely a conceptual framework, but one that has been discussed by market participants in Europe. The G-SPV (or GPU-focused special purpose vehicle) model sits as a separate SPV alongside (but separate from) a regular data centre SPV.

While the data centre SPV finances the physical data centre shell (building, power, cooling and connectivity), the G-SPV finances the GPUs themselves by acquiring them and leasing them back to the AI company (thus employing the sale-leaseback structure as part of the model). The G-SPV will aggregate a number of these lease contracts and issue ABS Notes backed by them. The idea of the structure is that GPUs as an asset class (which is a peculiar asset class from a data centre) is isolated and properly addressed according to its risk profile.

GPU-backed asset-backed securities (ABS)

No GPU-backed transaction has yet been publicly rated as a standalone ABS. The rated GPU financings to date have been structured as senior secured delayed-draw term loans while GPU loans have begun appearing as a minority component of broader equipment ABS pools.

A standalone GPU ABS presents some challenges: (i) isolating the collateral pool from the credit of the originator and the offtaker, which is more difficult for operator-dependent assets with concentrated customer bases, and (ii) the absence of a multi-vintage performance history on which to calibrate default, loss and residual-value assumptions. A third-party servicer market for GPUs is also less established compared to other asset classes.

That said, a GPU-backed securitisation is one of the mechanisms by which GPUs may be financed. A GPU-backed securitisation would be collateralised by a pool of GPUaaS contracts or GPU leases sold into a bankruptcy-remote SPV while the GPUs will be pledged as security to the trustee or collateral agent for the benefit of the investors.

Challenges of financing GPUs and TPUs

As GPUs as an asset class continues to develop, market participants are looking to address some of the challenges to financing them, including:

Obsolescence and useful life

One challenge with financing GPUs and TPUs as an asset class is the brevity of their useful life, and their quick race to obsolescence compared to other asset classes. New generations of GPUs are developed in as little as 12 to 18 months. Given that each new generation could perform up to 30x the performance of the previous one[9], the value of the previous-generation GPUs could be significantly affected when a new one is unveiled.

The result is that the contracts under which offtakers of GPUaaS services receive the benefits of installed GPUs typically have initial terms of five to six years, which is very short when compared to the 10 to 15 year initial terms frequently obtained for cloud / CPU deployments from hyperscale customers.

Some potential solutions have included a creditworthy third party, often the chip vendor, providing residual value support or loss-protection guarantee for lenders, something Europe has not seen as much as the US[10] Broadcom, for example, provided residual value support on the senior notes in the US$35bn Anthropic TPU financing.

Nvidia's financing framework of up to US$500bn announced in August 2026 is supported by a Nvidia guarantee of up to US$125bn to protect lenders against loss. Insurance companies particularly focused on AI have also offered residual value protection that could be worked into a transaction[11]. Risk leasing programmes that allow GPUs to be returned or upgraded in line with new hardware cycles may also be a potential solution.

Secondary market liquidity considerations

Both GPUs and TPUs have lower liquidity compared to other more-established asset classes. This affects TPUs more particularly given their specialisation and relatively smaller market. As the market continues to evolve and GPUs and TPUs become a more established asset class, it is likely that a secondary market will be more established.

Vendor concentration

The vendors involved in producing AI GPUs and TPUs are limited, as are the parties along the value chain (including fabrication, system assembly, software and even the neoclouds that rent compute capacity), presenting more concentrated risk.

Physical custody and recovery

The borrower in a GPU financing rarely owns the data centre facility. Instead, the chips are deployed in a third-party operator's facility, exposing the lender's collateral to the operator's contractual liens, statutory landlord-type liens in some states, and the practical risk that power and access are withheld on default. GPU lenders may mitigate these risks through a colocation waiver and access agreement covering lien subordination, cure rights, continuity of services during cure, and access rights to operate and remove the chips on enforcement and, where the operator is itself a lessee, non-disturbance from the operator's upstream landlord or mortgagee.

Recovery over GPU collateral is not as straightforward as it may appear. In the normal operating case, the chips are running customer workloads continuously, and pulling them out of service terminates live processes, exposes the operator to damages claims and service level agreement breaches under compute agreements, and begins to erode the collateral's residual value from the moment power is cut. These are typically addressed through contractual protections like step-in rights for the lender, a qualified replacement operator, cure and standstill periods before enforcement, an operational-IP licence that follows the

collateral, and (for TPUs) a direct agreement with Google covering continued access to the software stack.

A separate risk is that the operator's bankruptcy sweeps the chips and the customer contracts into its estate and holds them up under the automatic stay. Recent GPU and TPU financings mitigate this friction by placing title to the chips in a bankruptcy-remote SPV from day one, with the operator (whether a neocloud or an AI developer) holding the chips as service provider or lessee under a contractual arrangement rather than as owner. The result is a cleaner enforcement path on the lender's collateral, as the operator's bankruptcy does not sweep the chips or the customer contracts into its estate, and lender enforcement proceeds against the SPV's assets.

No asset-specific rating agency standard

At least six AI GPU financings have been rated by either Moody's, Fitch or DBRS, including the IREN/Microsoft US$3.65bn deal, CoreWeave's DDTLs 4.0, 5.0 and 5.5, Lambda's US$926m term loan and Nscale's US$1.83bn DDTL. S&P is yet to issue a rating on GPU collateral as at the time of writing.

In two of these deals, Fitch issued ratings using its corporate finance criteria, rating the neocloud on the deal as a corporate entity, then adding two notches to account for the structural protections and senior secured position and in one case has used a structured finance methodology in which it looked through the SPV borrower to the cashflow and modelled in the structural protections.

Moody's has also used project finance rating methodologies for some of these financings. KBRA is to date the only rating agency to have proposed a dedicated framework for rating GPUs (and other AI accelerators) as an asset class in its July 2026 research piece, *AI Credit: AI Compute Infrastructure — Cross-Sector Credit Considerations for GPU-, TPU-, and Accelerator-Intensive Assets*.[12]

While KBRA has not yet rated a deal in the asset class, and no agency has published a formal GPU/TPU rating methodology, this may well prove a precursor to broader adoption. Dedicated data centre rating methodologies began to emerge with S&P's criteria in June 2024 and rolled out across the following 18 months after years in which data centres were rated under adjacent methodologies. GPUs/TPUs may follow a similar arc toward more GPU/TPU-specific frameworks and, in time, asset-specific rating methodologies.

UCC and security taking

GPUs and TPUs could be classified a few different ways under UCC Article 9. There remains the question as to whether they are more appropriately classified as equipment under UCC § 9-102(a)(33) or inventory (when they are leased out to customers) under UCC § 9-102(a)(48).

The method of perfecting security interests under the UCC varies depending on the type of collateral, so this distinction becomes important when classification affects collateral descriptions, perfection and priority. For instance, a lender may have an existing purchase money security interest (PMSI), which may have super senior priority, depending on the timing for perfecting the PMSI and the classification of the asset under the UCC.

A third classification may arise where the chips become so related to a particular property that an interest in them arises under real property law. In that instance, they are considered fixtures under UCC § 9-102(a)(41).

GPU lessor's rights to lease payments under a contract for lease of GPUs will separately be considered chattel paper (UCC §9-102(a)(11)), while where such contract is for compute-services instead of a lease of the hardware, the receivables will be considered accounts under UCC §9-102(a)(2). These should be taken into consideration when taking security for deals of this nature.

Software and the software licences associated with the chips, by contrast, are considered general intangibles (UCC § 9-102(a)(42)), unless the software is embedded in the chips in a manner that brings it within the definition of goods under UCC § 9-102(a)(44). CUDA toolkit entitlements from Nvidia and TPU software/support rights from Google therefore fall outside the equipment classification and are typically perfected as general intangibles, although

firmware and microcode physically burned into the GPU or TPU travel with the hardware as part of the goods.

As noted earlier, TPUs are architecturally different from GPUs with their economic utility depending heavily on Google's software, firmware and support ecosystem. A lender taking security over the physical TPUs may consider pairing the Article 9 security interest with direct contractual access to that software ecosystem, whether through a direct agreement with Google covering licence continuity, cure rights, and assignability or through credit support from Google and/or the chip co-designer.

Export control and sanctions

As advanced AI chips are now a focus of national security policy, lenders must be familiar with the legal framework for sanctions and export control that apply to the GPUs/TPUs they are financing. Export control and sanctions apply to every stage of GPU/TPU financing – origination, operation, enforcement and exit and increasingly, to ultimate parent ownership structures that go beyond the immediate counterparty.

At inception, structuring must address the classification of the chips under the Bureau of Industry and Security (BIS) Export Administration Regulations (EAR), the licensing and end-use conditions that attach to them, the identity and jurisdiction of permissible end-users, and the geographic footprint of the deployment sites (which drive the choice of borrower vehicle, the location of the collateral, and the covenant package). Collateral value, mobility and even recovery can be affected if subject to tightened controls.

On recovery, export controls and sanctions constrain the universe of eligible transferees and remarketing counterparties. A chip that cannot be exported to, or operated by, large parts of the global market is a less liquid asset, and realisable collateral values and remarketing timelines must be analysed on that basis.

Given how critical AI hardware has become to many countries, the laws around them are also changing rapidly. In the US, market participants must be up to date with the BIS rules which touch routinely on licence reviews and restrictions on exports, re-exports and in-country transfers of chips. Under the EAR, advanced AI GPUs are classified primarily under ECCN 3A090, and the servers containing them under ECCN 4A090. These classifications carry varying degrees of export licence requirements depending on the destination country and end-user.

In January 2025, BIS issued the "AI Diffusion" rule, which would have grouped countries into three tiers with different licence requirements; BIS rescinded it in May 2025, before its compliance date, and has since relied on targeted rules and guidance, including a January 2026 rule moving certain H200-class exports to China to case-by-case review.

Additionally, persons in jurisdictions subject to comprehensive OFAC sanctions, and parties on OFAC's SDN List, generally cannot be dealt with absent a licence, which affects both chip sales and payment flows. The EAR's General Prohibition 10[13] for instance, states that no person may *"sell, transfer, export, reexport, finance, order, buy, remove, conceal, store, use, loan, dispose of, transport, forward, or otherwise service"* any item subject to the EAR with knowledge that a violation has occurred, is about to occur, or is intended to occur.

A lender that finances the acquisition of controlled GPUs knowing (or having reason to know) that the hardware will be diverted to a restricted end-user or destination could face liability under this provision. Given that "knowledge" under the EAR includes having reason to know, lenders should diligence the supply chain, end users and deployment sites, and monitor them through the life of the loan.

US export controls apply to the items themselves, not just the parties. This means that even non-US operators may face liability under the EAR if restricted GPUs, controlled technology, or sanctioned end users are present in their facilities, even indirectly through tenants or sub-tenants. If implemented, European G-SPV structures with GPUs in third-party data centres across multiple jurisdictions would face a chain of compliance obligations.

European structures and entities are of course required to comply with EU sanctions and export controls

(administered through Council Regulations and national authorities) and also with US EAR extraterritorial jurisdiction, which applies to items subject to the EAR even after they leave the US.

A recent regulatory development on export controls is the Remote Access Security Act (RASA), passed by the US House of Representatives in January 2026, and being considered by the US Senate. If signed into law, RASA would extend US export control law to include remote access, bringing a foreign entity's remote access to export-controlled items within BIS' enforcement authority.

Notably, the version of the bill currently before the US Senate focuses on remote access used to train AI models that could facilitate weapons of mass destruction and offensive cyber operations, and human-rights-undermining surveillance which could introduce compliance and diligence challenges to GPUaaS providers and may result in limited GPUaaS access by certain customers. While RASA is not yet passed into law, market participants have worked on a number of fallbacks to address this potential risk, such as token provision and termination compensation.

Conclusion: Where do we go from here?

The AI boom is at an all-time high and is being pulled into institutional capital markets at rapid speed. One projection puts the number at US$7.6 trillion of cumulative AI infrastructure capex through 2031.[14]

With GPUs constituting 50-65% of the all-in cost of a data centre (even counting power and cooling infrastructure), the demand for GPUs is expected to continue to grow. GPUaaS market alone is projected to grow to US$6.10bn in 2026 and approximately US$37.10bn by 2035[15].

The yield profile on GPU and TPU debt has also drawn significant investor attention, and the sector's trajectory indicates continued expansion and relatedly, more capital deployment. The expectation is that this capital will take diverse forms and a range of innovative financing structures tailored to the characteristics of the collateral.

The rating agencies are an important anchor point, as their commentaries, frameworks and methodologies can lead to broader investor pools and wider adoption of the asset class across the market. Should more GPU- and TPU-specific frameworks/methodologies develop, we can expect greater standardisation and a faster maturation of the collateral class, both of which will drive further investment.

Equally significant, as AI workloads shift from training to high-throughput production inference, general-purpose chips like GPUs will be increasingly complemented by ASICs like TPUs, with hyperscalers running both GPUs and their own ASICs in parallel.

Demand for Neural Processing Units (NPUs), the on-device AI accelerators optimised to run AI tasks locally, is also projected to grow materially as on-device AI adoption deepens, further accelerated by the planned arrival of commercial 6G networks from 2029 and broader IMT-2030 deployments around 2030. These shifts will not replace the GPU as the central AI compute asset, but it is important to understand how rapidly the surrounding technology stack is changing and how that will affect the collateral pool in financing structures.

The direction of travel is clear. While GPUs and TPUs are subject to the same development cycle as any other maturing asset class, the need for financing is such that we should expect to see structures tailored to address their particularities and the asset class to continue to grow. The question for 2027 and beyond is not whether the asset class matures, but how fast it does and under what types of structures.

* *With additional thanks to Elisabeth Johnson, Rich Jones, David Martin, Philip W. Lee and Sean Solomon for their contributions.*

Notes:

1 Benzinga, "The AI Boom is Becoming a US$4.1 Trillion Debt Story, JP Morgan Says.", Surbhi Jain, June 16, 2026 https://www.benzinga.com/markets/tech/26/06/53233210/the-ai-boom-is-becoming-a-4-1-trillion-debt-story-jpmorgan-says

2 S&S Insider, "Graphics Processing Unit Market (2026-2035)" July 1, 2026 https://www.snsinsider.com/reports/graphics-processing-unit-market-6783

[3] Fortune Business Insights, "Graphic Processing Unit (GPU) Market", September 14, 2026, https://www.fortunebusinessinsights.com/graphic-processing-unit-gpu-market-109831

[4] Crusoe newsroom, October 23, 2023 https://www.crusoe.ai/resources/newsroom/crusoe-announces-significant-expansion-of-cloud-business-with-new-capacity

[5] DCD, "Nvidia's Blackwell GPUs are Sold Out for 12 Months", Charlotte Trueman, October 11, 2024 https://www.datacenterdynamics.com/en/news/nvidias-blackwell-gpus-sold-out-for-12-months/

[6] SEC Edgar, https://www.sec.gov/Archives/edgar/data/1769628/000176962826000357/ex101creditagreement.htm

[7] Blockspace, "CoreWeave closes US$2.6bn GPU-backed loan" August 11, 2026, https://blockspace.media/insight/coreweave-closes-2-6-billion-gpu-backed-loan/; SEC EDGAR, Coreweave's 8k – Exhibit 99.1, https://www.sec.gov/Archives/edgar/data/1769628/000176962826000357/ex991pr.htm

[8] DCD, "Nvidia signs US$1.5bn deal to lease its GPUs back from Lambda – report", Georgia Butler, September 8, 2025, https://www.datacenterdynamics.com/en/news/nvidia-signs-15bn-deal-to-lease-its-gpus-back-from-lambda-report/

[9] Nvidia Developer, K Moseley, K Perez and P. Mahajan, "Scaling Token Factory Revenue and AI Efficiency by Maximising Performance per Watt", March 25, 2026, https://developer.nvidia.com/blog/scaling-token-factory-revenue-and-ai-efficiency-by-maximising-performance-per-watt/

[10] TMT Finance, Clara Oliver-Amorim, "Inside European GPU-secured debt; terms, tenors and appetite", October 1, 2026, https://www.tmtfinance.com/analysis/115669?utm_campaign=10511182-TMT%20Finance%20Deal%20Digest&utm_medium=email&_hsenc=p2ANqtz-9a8FzckQfvruleT-quPLTliPViguoMp6cQqcZ5vter13ArgRQMq7PBN4yl-Gae88KptnR9Y1pJXuYalPfWehlPJpPGV5599G_3TRJffT7MGX4Ch5-A&_hsmi=443367000&utm_content=443366677&utm_source=hs_email

[11] The Insurer, "GPU residual value: Insurance paper used to tackle AI financing risk" George Abbott, August 26, 2026, https://www.theinsurer.com/ti/analysis/gpu-residual-value-insurance-paper-used-to-tackle-ai-financing-risk-2026-08-26/

[12] KBRA, "AI Credit: AI Compute Infrastructure — Cross-Sector Credit Considerations for GPU-, TPU-, and Accelerator-Intensive Assets", July 27, 2026, https://www.kbra.com/publications/KPZcbgKd/ai-credit-ai-compute-infrastructure-cross-sector-credit-considerations-for-gpu-tpu-and-accelerator-intensive-assets?format=web

[13] 15 C.F.R. § 736.2(b)(10)

[14] Goldman Sachs, G. Lee and L. Greenbaum, "Tracking Trillions: The Assumptions Shaping the Scale of the AI Build-Out", May 1, 2026, https://www.goldmansachs.com/insights/articles/tracking-trillions-the-assumptions-shaping-scale-of-the-ai-build-out

[15] Precedence Research, "GPU as a Service Market Size", Shivani Zoting, April 14, 2026, https://www.precedenceresearch.com/gpu-as-a-service-market

Contact us:

Linklaters LLP

1290 Avenue of the Americas,
New York, NY 10104
US
Tel: +1 212 903 9000
Web: www.linklaters.com