ICELT.

MARKET INTELLIGENCE · 2026

Renting the Pickaxe

The operator's side of GPU lending. Thin coupon, thick risk, and where the "low-risk" pitch breaks.

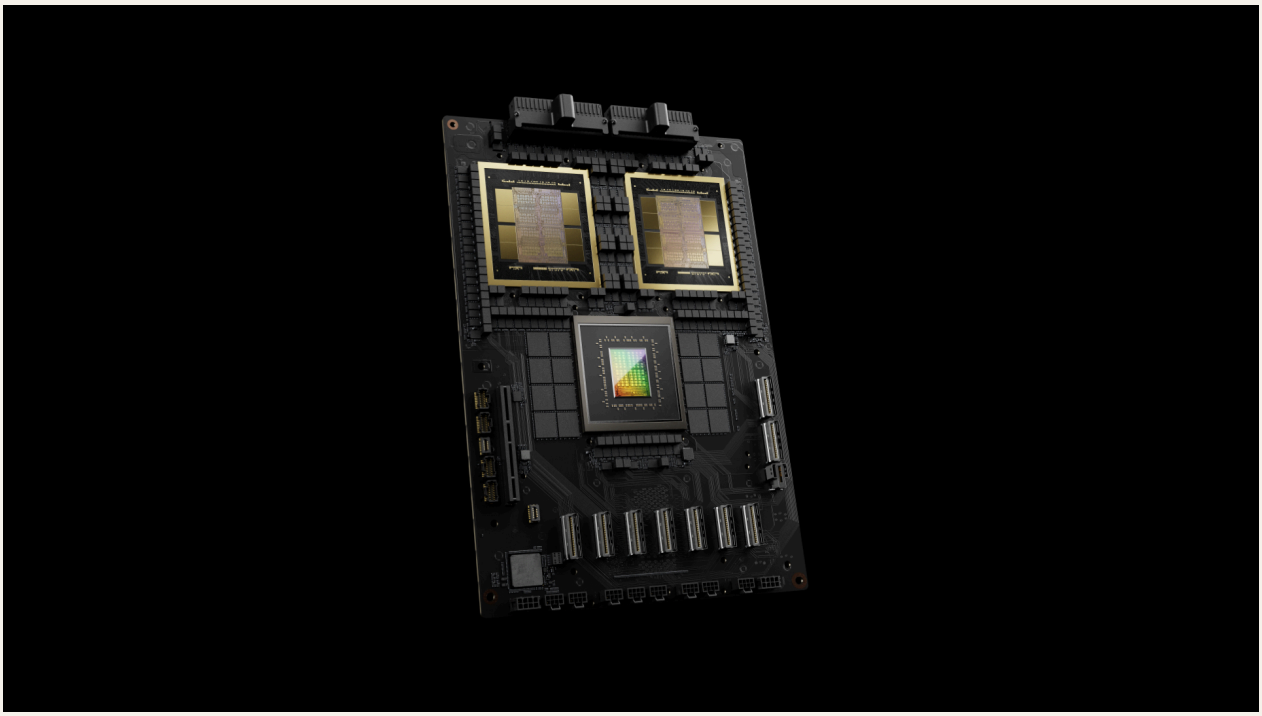

PREPARED BY

Philipp Eiselt
www.icelt.net

PREPARED FOR

Public Report

DATE

May 2026

CONTENTS

EXECUTIVE SUMMARY

A new operator role has emerged in the AI-compute stack. The **GPU lessee** takes possession of hardware they do not own, runs it in their own facility, and rents the compute hourly on decentralised marketplaces. The pitch is hard to argue with. No CapEx, and you can walk away clean if it fails. The reality is narrower.

We mapped the marketplaces an operator would list on, built the unit economics from current rental and power data, and stress-tested the "low-risk" claim. Put plainly, **the framing is honest on capital risk and materially misleading on operating risk**. Under a revenue-share contract both sides bleed together. Under a **fixed monthly lease** set in the 2024 pricing regime, the operator is now the shock absorber for every adverse scenario, and at mid-2026 H100 spot rates a $1,500/month fixed lease is **structurally insolvent**.

This paper is for operators considering the play, for owners structuring the deal, and for boards underwriting either side. Numbers are sourced; case stubs are flagged as illustrative composites.

$1.49
H100 spot
Floor per GPU-hour, mid-2026

46%
Break-even
Utilisation, 50/50, EU power

~15%
Real margin
Neocloud post-depr. (McKinsey)

18mo
Lifecycle
H100 active obsolescence

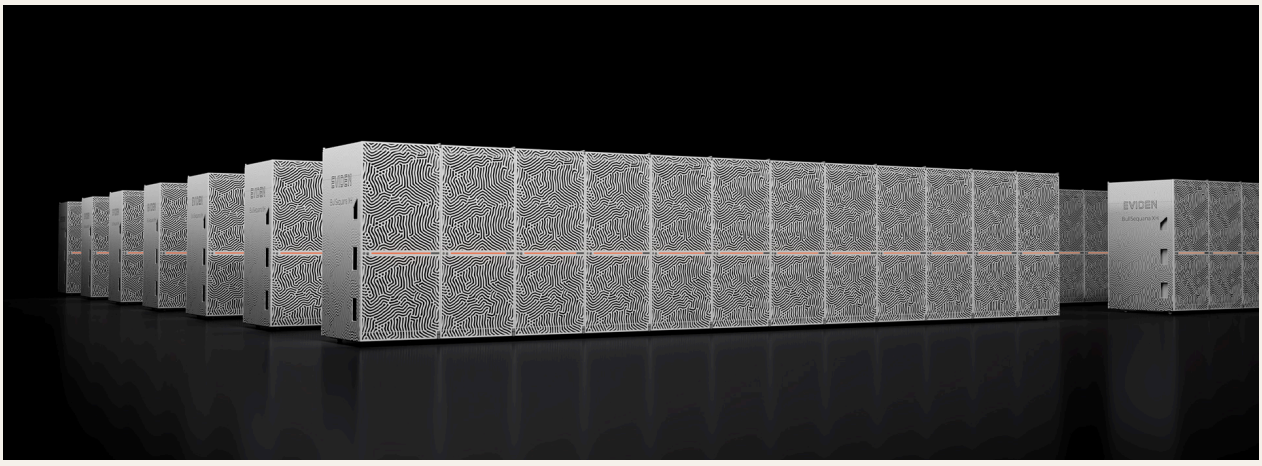

Eviden / NVIDIA "Jupiter". The reference scale this paper measures the small operator against.

01 What we mean by the operator role

This paper studies one specific actor in the AI-compute economy. The **operator who possesses GPUs they do not own** rents the compute out. The hardware owner (a financier, an SPV, a Tier-2 cloud, or a tokenisation platform) retains paper title. The operator contributes power contracts, real-estate, labour, and a customer book. The two sides split the revenue.

THE PITCH IN ONE LINE

"Run someone else's GPU, keep the spread, walk away clean if it doesn't work."

The pitch is structurally similar to **owner-operator trucking**, where the driver finances or leases a tractor, books loads, eats fuel and maintenance, and absorbs the residual-value risk indirectly through lease terms. Of the common analogs (Airbnb hosting, taxi medallions, mining colocation), trucking is the most honest. Airbnb fails because the asset belongs to the host. Medallions fail because the scarce thing is a regulatory permit, not a productive asset on an 18-to-30-month obsolescence curve.

What this paper does not cover

We exclude the **facility-scale** deals where ex-Bitcoin miners lease entire sites to hyperscalers. Riot's 10-year AMD deal at Corsicana, IREN's $9.7B Microsoft arrangement, Hut 8's $7B Google-backed deal, and Core Scientific's $10B CoreWeave hosting contract are institutional capacity transactions between regulated counterparties. They are not the same product as a small operator listing eight H100s on Vast.ai. Conflating them is the single most common error in the trade press.

The institutional view. NVIDIA Rubin platform racks, the rate-card the rest of the market discounts against.

02 The market today

Where an operator actually lists capacity

MARKETPLACE	MODEL & TERMS	OPERATOR ECONOMICS
Vast.ai	Self-priced; ISO-27001 + 5+ flagship GPUs for datacenter tier	Implicit 25–35% commission (third-party est.)
RunPod	Community Cloud, vetted third-party hosts	Host split not publicly disclosed
TensorDock	Stripe payouts, fixed split	Host keeps 75%, platform 25%
Aethir	DePIN; 91k+ Checker Nodes; 45-day payout hold for SLA	80% to host, 20% to foundation
io.net	DePIN; $IO staking required; USDC has 2% facilitation fee	Wide range; 327k verified GPUs, 5,350 cluster-ready (Mar 2025)
Akash	Bare-metal provider model; provider count 70→63 in Q3 2025	Variable; first contraction in years
Render Network	Render/compute-only; 8GB+ VRAM, 1TB NVMe minimum	Narrow vertical
Salad	Consumer GPUs (3090/4090) only; fiat & gift-card payouts	Up to $180/mo per top rig. Hobbyist scale.

The hardware-owner programs that matter

Four structures define the "owner-provides, operator-runs" market in 2026.

- **Compute Labs** tokenises GPUs as GNFTs on Solana, with NexGen Cloud as operating partner. It charges a flat **10%** across tokenisation, management, and yield, and advertises a **30% target yield** to GNFT holders. Thirty percent is a target, not a track record.
- **SPV-based GPU leasing** is the institutional version. NVIDIA placed roughly **$2B** into an xAI G-SPV that buys GPUs and leases them back, and completed a **$1.5B sale-leaseback** with Lambda in September 2025. GPU-backed loans run at 50–70% advance rates and 12–15% interest; ABS issuance across data centres and GPU assets is expected to grow from **$8B in 2025 to $25B by 2028**.
- **Public-market owner-operator vehicles**, exemplified by **Alpha Compute Corp.** (NASDAQ: ALP), which rebranded from AlphaTON Capital in April 2026 and now owns and operates dedicated Blackwell B200/B300 clusters under direct enterprise lease. Their financing template is the cleanest public

example of the **non-recourse GPU-secured loan**. The **$31.9M facility** drawn in April 2026 has collateral strictly limited to the hardware. The lender can repossess the GPUs but cannot pursue the corporate balance sheet. This is the institutional version of the "owner keeps the risk, operator keeps the spread" pitch that gets translated downward, on much worse terms, to retail operators.

- **Bitcoin-miner-turned-AI-host facility leases** covering Riot, IREN, Hut 8, and Core Scientific are institutional plays and not relevant to a small operator. They do, however, set the reference rate that the rest of the market prices against.

WHY THE ALPHA COMPUTE STRUCTURE MATTERS

Their first cluster is **504 B200 GPUs in a Canadian hydro datacenter.** It signed a single **$32.2M / 24-month fixed lease** with an unnamed frontier AI lab. That works out to **$16.1M ARR, or roughly $32k per GPU per year**, on a named-counterparty take-or-pay contract. It is a clean upper bound for the owner side. The retail operator listing on Vast.ai is competing for the **residual** of that same market with none of the credit profile, none of the take-or-pay protection, and a much shorter customer book.

Contract structures, the real range

01

Pure revenue-share

Marketplace-mediated. Operator bears 100% of power and colo costs; no floor. Both sides share price-compression risk. Example: Aethir 80/20, TensorDock 75/25.

02

Fixed monthly lease

Operator pays $/GPU/month regardless of utilisation. Owner is hedged; operator carries every downside. Was viable at $4–8/hr H100 spot. Not at $1.49 to $2.50.

03

Hybrid (floor + share)

Owner takes a minimum monthly draw, operator keeps the upside above a threshold. The honest structure for 2026; uncommon in current deal flow.

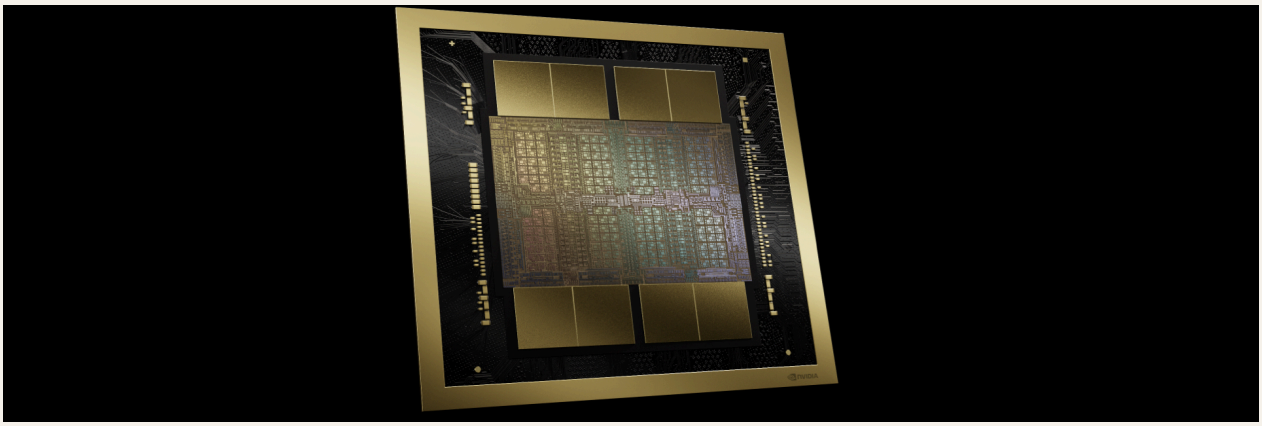

The asset being financed. NVIDIA Blackwell architecture die. The unit the SPV, the GNFT and the small operator are all priced against.

03 Unit economics.

The model assumes a 700 W GPU with a PUE of 1.4. Server overhead is roughly 200 W per unit, covering the CPU, NIC, and storage share of an 8-GPU node. The month has 730 hours at a realised rate of $2.00 per hour. Colo and bandwidth are allocated at $90 per month. Labour, monitoring, and the RMA pool at $80. All figures are **per GPU, per month**.

SCENARIO	POWER $/KWH	UTIL.	REVENUE	POWER	OTHER	LEASE	NET
A. 50/50 rev-share, Texas	0.06	60%	$876	$55	$170	$438	**+$213**
B. 50/50 rev-share, Texas	0.06	40%	$584	$55	$170	$292	**+$67**
C. 50/50 rev-share, Germany	0.18	60%	$876	$166	$170	$438	**+$102**
D. 50/50 rev-share, Germany	0.18	40%	$584	$166	$170	$292	**−$44**
E. $1,500 fixed lease, TX	0.06	80%	$1,168	$55	$170	$1,500	**−$557**
F. $1,500 fixed lease, TX	0.06	60%	$876	$55	$170	$1,500	**−$849**
G. $1,500 fixed lease, TX	0.06	40%	$584	$55	$170	$1,500	**−$1,141**
H. $1,500 fixed lease, DE	0.18	80%	$1,168	$166	$170	$1,500	**−$668**

Figure 1 · Per-GPU monthly P&L. "Other" rolls colo, bandwidth, labour and RMA pool. Power assumes PUE 1.4 and a 200 W server-overhead allocation per GPU. 730 hours/month at $2.00/hr realised.

WHY FIXED LEASES BREAK

A $1,500/month fixed H100 lease implies a **$2.05/hour** cost floor before any operating expense. With mid-2026 H100 spot at **$1.49–$2.50**, that floor is at or above realised revenue. The lease was priced against a market that no longer exists. A defensible 2026 fixed lease sits closer to **$700–900** per month. That sits in line with an asset depreciating 60% over 18 months on a $25k street price.

Break-even utilisation

Figure 2 · Required utilisation, % of hours, to reach cash-positive at $2.00/hr realised. A $1,500 fixed lease is mathematically infeasible at current 2026 spot prices.

Take-aways

- **The economics are thinner than the pitch suggests.** A 100-GPU fleet on the best available rev-share deal clears roughly **$21k/month** gross profit before the operator pays themselves. Net per GPU lands in the mid-single digits on a good day. This is an operations business, not an investment.
- **Where you plug in matters more than almost anything else.** The gap between Texas at $0.06/kWh and Germany at $0.18 runs to roughly $110 per GPU per month, which is more than the combined colo and labour bill. European operators running US-sourced hardware are being squeezed from both ends.
- **Utilisation is not yours to decide.** The marketplace queue and your reliability score set it for you. Budget for 30–40% in the first two months, not the 70% that featured in the underwriting model.
- **The depreciation clock was set by someone who does not have to keep the racks full.** Whoever booked a 4–6 year useful life in their accounts does not feel the same pressure. The hardware ages on their schedule; the revenue shortfall lands on yours.

04 Risk decomposition. Is the pitch honest?

RISK	PROB.	SEV.	SCORE	NOTE
Utilisation shortfall	High	High	**HIGH**	New listings ramp slowly; Blackwell is siphoning demand off H100.
Price compression (next-gen)	High	High	**HIGH**	H100 spot already halved from 2024; B200 ramp continues through 2026.
Repossession reality	Med	High	**HIGH**	Owner must physically extract racked GPUs from operator's colo across jurisdictions. Title is paper.
Power-cost spike	Med	Med	**MED**	EU exposed (gas-linked); US Texas hedgeable but constrained by ERCOT queue.
Hardware failure / RMA	Med	Med	**MED**	Title with owner ≠ liability with owner. Read the SLA section.
Export-control / regulatory	Low	High	**MED**	AI Diffusion Rule rescinded May 2025; successor pending. Cross-border sub-rental triggers tier risk.
FX (USD revenue, EUR costs)	Med	Med	**MED**	Hedgeable with quarterly forwards. Most operators don't.
Counterparty / renter non-pay	Low	Low	**LOW**	Vast.ai and RunPod are pre-pay; escrow handles it.

> " We don't own the GPUs and we don't really own the customers. What we own is a power contract, a permit, and 14 months of operational scar tissue. Everything else is rented. Including the optimism. "

OPERATOR, 80-GPU CLUSTER, WEST TEXAS (ILLUSTRATIVE COMPOSITE)

Verdict on “low risk”

The framing is **technically true for capital risk and materially misleading for operating risk**. The operator is short volatility on hourly rental prices, short power, and short their own labour. The hardware owner holds paper title to a depreciating asset they cannot easily repossess from a colo cage they don’t operate.

ONE-LINE VERDICT

“Low risk” is **honest on CapEx, marketing on operating risk.** It is also an outright trap under a fixed-lease structure at 2026 rental prices.

The repossession gap

A note for both sides of the table. In researching this paper we found no public, named operator default with a documented hardware repossession. The mechanics (UCC-1 filings to GPU serial numbers, geo-fencing, escrowed BMC/IPMI credentials, SPV ownership wrappers) are well documented in theory. The practice is silent. This is itself a signal. Either the market is too young for defaults to have surfaced, or losses are being absorbed quietly. Owners structuring a fixed-lease deal in 2026 should not assume the recovery path is liquid.

Where the contract actually lives. The cable plant and BMC layer are the real surface that determines uptime, SLA penalties, and whether a repossession route is liquid.

05 Three operator archetypes

ARCHETYPE A

The pivoting miner

Ex-Bitcoin operator with 14 MW in West Texas. Retrofitted two halls for air-cooled H100 hosting on behalf of a Tier-2 neocloud. Sells power, racks, hands, uptime, all marked up against a take-or-pay floor. Hall 4 still runs SHA-256 ASICs as a hedge.

ARCHETYPE B

The garage aggregator

Ex-ML engineer managing 38 RTX 4090s and 6000 Adas across three residential and one light-industrial site in Quebec. Doesn't own most cards. Manages for a syndicate of crypto-native investors. Edge is uptime discipline, custom containers, and a Discord reputation score.

ARCHETYPE C

The curious MSP

Munich engineering firm leasing two GB200 NVL36 racks from a financing vehicle, hosted in their existing colo suite. Resells fine-tuning and inference to regulated EU clients wanting non-hyperscaler residency. Not a cloud company. A margin arbitrageur with a compliance story.

Three case stubs

WINNER · COMPOSITE

Clean spread

A 22 MW former Ethereum-era site in Paraguay re-papered its hydro PPA at **$0.028/kWh**, took 1,400 H200s on sale-leaseback from a Singapore GPU fund, and signed a Korean foundation-model lab as anchor. **EBITDA clears 62%** on hardware it does not own. The real asset is the power contract.

STRUGGLING · COMPOSITE

The 4090 farm in suburbia

An aggregator running **240 consumer cards** across nine residential sites books **$0.41–$0.58/hr** at **47% utilisation**. Two landlords have served notice over breaker trips; insurance was non-renewed. Profitable on paper, insolvent on timing. Payouts settle in tokens. Costs settle in CAD.

CAUTIONARY · COMPOSITE

The reseller who got long

A UK MSP signed a **36-month take-or-pay** on 96 H100s in Q3 2024 at **$2.10/hr-equivalent** assuming a $3.50/hr resale. Spot slid; the anchor renegotiated, then defaulted. The lessor invoked a **personal guarantee on the principal's home**. The hardware still earns; the operator does not.

06 Outlook 2026. What is shifting under the operator

H100 rental pricing has re-tightened, not collapsed

After the late-2025 "compression" narrative, 1-year contract rates broke back above **$2/hr-equivalent** in Q1 2026 and rose 15–20% month-on-month into spring; the expected oversupply did not arrive on schedule. This is the single most important data point for any lessee underwriting a deal today.

Blackwell is sold out, but the rollout is uneven

GB200 and B200 backlogs sit near **3.6M units** and have been pre-allocated largely to hyperscalers. Sub-scale operators get crumbs and longer lead times, which **paradoxically extends the earnings life of installed H100/H200 fleets**. Plan for 2027 obsolescence, not 2026.

Bitcoin miners have become AI's landlords, at scale

IREN's $9.7B Microsoft deal, Hut 8's $7B Google-backed arrangement, and Core Scientific's $10B CoreWeave hosting contract signal that the "pivoting miner" archetype is now an institutional category. **AI/HPC is on track to deliver 70% of listed-miner revenue by year-end 2026.**

Power is the binding constraint, especially in Texas

ERCOT's large-load queue has roughly **quadrupled** in a year. About **360 GW** of prospective data-centre demand sits queued against an 85.5 GW historical peak. SB6 (June 2025) now forces collateral and cost-share before interconnection. The era of cheap, instant Texas siting is over.

DePIN marketplaces are maturing, not winning

io.net reports >1M aggregated GPUs and Akash quotes H100s at **$1.33/hr** versus AWS **$3.93/hr**, but settlement in tokens, utilisation economics, and enterprise-procurement friction keep these venues as a spot/overflow layer rather than primary revenue for serious operators.

07 Appendix

A. Definitions

- **Operator (lessee).** The party who takes possession of GPUs they do not own and rents the compute on the open market. They supply the power contract, colo space, labour, and customer relationships.
- **Hardware owner.** The party holding paper title, typically a financier, an SPV, a Tier-2 cloud, a tokenisation platform, or an NVIDIA partner.
- **Realised hourly rate.** What the operator actually receives per GPU-hour after marketplace commission and any discounting. Not the headline list price.
- **PUE (Power Usage Effectiveness).** Facility total power divided by IT load. We use 1.4 throughout, which is standard for air-cooled colo; liquid-cooled facilities now reach 1.1–1.2.

B. Method

This paper triangulates three independent passes. A market researcher covered marketplaces, contract structures, and programs. A unit-economics analyst built the P&L, break-even, and risk register. A strategy writer shaped personas, analogies, and case framing. Numerical claims are sourced. Case stubs are labelled **illustrative composites**; no individual operator or facility is identified. We found no public, named operator default with documented hardware repossession and flag that gap explicitly.

C. Primary sources

- McKinsey. The evolution of neoclouds and their next moves https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-evolution-of-neoclouds-and-their-next-moves
- Silicon Data. H100 Hyperscaler Index, April 2026 https://www.silicondata.com/blog/h100-hyperscaler-index-april-2026
- Silicon Data. B200 rental price, March 2026 update https://www.silicondata.com/blog/b200-rental-price-march-2026-update
- SemiAnalysis. The Great GPU Shortage, rental capacity index https://newsletter.semianalysis.com/p/the-great-gpu-shortage-rental-capacity
- Bird & Bird. GPU-based financing in the global data-center market https://www.twobirds.com/en/insights/2025/germany/gpubased-financing-in-the-global-data-center-market-a-new-standard-for-largescale-investment-structu
- Messari. State of Akash, Q3 2025 https://messari.io/report/state-of-akash-q3-2025
- Aterio. Former Bitcoin miners powering the AI boom https://www.aterio.io/blog/how-former-bitcoin-miners-are-powering-the-ai-boom
- CoinDesk. Compute Labs targets 30% yield on tokenised AI infrastructure https://www.coindesk.com/business/2025/06/18/a-startup-is-looking-to-pay-30-yield-by-tokenizing-ai-infrastructure

- Introl. GPU cloud price collapse, H100 market, December 2025 https://introl.com/blog/gpu-cloud-price-collapse-h100-market-december-2025
- Latitude Media. ERCOT's large-load queue has nearly quadrupled https://www.latitudemedia.com/news/ercots-large-load-queue-has-nearly-quadrupled-in-a-single-year/
- DCD. NVIDIA-Lambda $1.5B sale-leaseback https://www.datacenterdynamics.com/en/news/nvidia-signs-15bn-deal-to-lease-its-gpus-back-from-lambda-report/
- Greenberg Traurig. Navigating GPU export controls, December 2025 https://www.gtlaw.com/en/insights/2025/12/navigating-gpu-export-controls-and-ai-use-restrictions-in-data-center-operations
- GlobeNewswire. AlphaTON Capital rebrands as Alpha Compute Corp., April 2026 https://www.globenewswire.com/news-release/2026/04/20/3276894/0/en/alphaton-capital-rebrands-as-alpha-compute-corp-to-reflect-its-growing-ai-compute-business.html
- StockTitan. Alpha Compute executes binding $31.9M non-recourse GPU loan https://www.stocktitan.net/news/ALP/alpha-compute-executes-binding-31-9-million-non-recourse-gpu-z17rlmei9adp.html
- GlobeNewswire. Alpha Compute closes $32.2M revenue contract with AI lab customer, May 2026 https://www.globenewswire.com/news-release/2026/05/12/3293000/0/en/Alpha-Compute-Closes-32-2-Million-Revenue-Contract-with-AI-Lab-Customer.html

D. What we deliberately excluded

- Facility-scale leases between former Bitcoin miners and hyperscalers, which are a categorically different product.
- DGX Cloud reseller agreements, which are services contracts rather than GPU subleases.
- Crypto-native yield tokens that wrap GPU exposure without a physical operator in the chain.
- Sovereign computing programmes commissioned directly by national governments.

E. Image credits

- **Cover**. NVIDIA GB200 Grace Blackwell Superchip. NVIDIA Corporation, used under press / editorial allowance.
- **Page 3**. Eviden / NVIDIA "Jupiter" exascale supercomputer. NVIDIA Corporation / Eviden.
- **End of Section 1**. NVIDIA Rubin platform racks. NVIDIA Corporation.
- **Section 6**. Server backplane and cable plant. Stock editorial photograph.