

# Compute Real Estate

The New Asset Class Where Land, Power and GPUs Meet

**\$720B+**

Big-four hyperscaler  
capital spending tracking  
for 2026

**30 GW**

Frontier AI compute  
capacity, doubling  
roughly every 7 months

**>\$20B**

GPU-backed debt  
outstanding, up from  
~\$11B a year earlier

**\$514B**

Alphabet cloud backlog  
at Q2 2026 — demand  
already signed

For a century the safest rent came from owning a building and leasing it. In 2026 the most sought-after tenants in the world are AI labs looking for powered land with GPUs already inside it. This issue follows the money as real estate, energy and GPU-as-a-Service merge into a single investable asset — with long-term contracts, bank debt against the hardware, and rental income measured per GPU-hour instead of per square metre — and closes on Fluidstack, the asset-light operator turning landlords into AI infrastructure partners.

## IN THIS ISSUE

[01 Executive Summary](#)[02 The Market in Numbers](#)[03 Power — the binding constraint](#)[04 What One Gigawatt Costs](#)[05 The Neocloud Model](#)[06 League Table](#)[07 Public Neoclouds Scorecard](#)[08 Financialization of Compute](#)[09 Frontier Ventures](#)[10 Company Focus — Fluidstack](#)[11 Risks & Watch Items](#)

## 01 — EXECUTIVE SUMMARY

# Compute is revenue

**Core thesis:** the physical layer of AI — land, power, buildings and the GPUs inside them — has stopped being a cost centre and become an asset class in its own right. It now has everything a real-estate market has: tenants on multi-year leases, lenders who take the asset as security, and, from this autumn, a forward curve. We call it **compute real estate**. Its rent is measured per GPU-hour rather than per square metre, and its tenants are the best-funded companies on earth.

## FIVE HEADLINE FINDINGS

**01 The \$700 billion year**

Amazon, Alphabet, Microsoft and Meta are guiding to roughly \$720–745B of capital spending in 2026, the bulk of it data centres. Suppliers are reporting the other side of that cheque: NVIDIA's latest quarter delivered \$96.2B of revenue with \$108B guided next, and Oracle's contracted backlog reached \$638B.

**02 Power, not chips, is the constraint**

Data-centre electricity use is on course to roughly double from 485 TWh to ~950 TWh by 2030 (IEA). Operators have signed 13 nuclear deals for 9.8 GW, and local opposition now touches ~\$130B of proposed projects. Whoever controls powered land controls the schedule.

**03 One gigawatt costs ~\$38 billion**

Epoch AI's all-in estimate is ~\$38B per GW, of which roughly \$8.5B is the site — land, building, power and cooling — and the rest is chips and networking. At three-year depreciation that is ~\$12B a year to carry. The rent side: an H100 leases for ~\$4.00/hour at a neocloud, nearly double that at a hyperscaler.

**04 The lease has met the mortgage**

GPU-backed debt has passed \$20B, up from ~\$11B a year earlier, and CoreWeave closed the first investment-grade facility secured on GPUs — \$8.5B rated A3. When a bank lends against the chips the way it lends against the building, an asset class has been born.

**05 The Gulf already owns the ingredients**

Compute real estate rewards exactly three things: cheap power, land and patient capital. HUMAIN's plan targets **6.6 GW** of capacity behind ~\$77B of investment, and Saudi Arabia holds the world-record solar tariff of ~1 cent per kWh. This is the first AI asset class that pays a yield on all three Gulf advantages at once — the sovereign bridge made physical.

**What an allocator should take from this issue.** Treat compute the way you would treat a logistics park in 2015: a scarce, income-producing asset with a credit-worthy tenant, financed on the lease. The risks are real — chips depreciate faster than buildings, contracts run two to seven years, and grid connections decide who builds first — and each has a page here. But the direction is settled: the money in AI is flowing to whoever owns the ground the intelligence runs on.

02 — THE MARKET IN NUMBERS **SNAPSHOT · AUGUST**

# The \$700 billion year

Every number below is the same money seen from a different seat: what the buyers of compute say they will spend, and what the sellers of compute say they have already been promised. Together they describe a market that has moved from “will demand hold?” to “can it be built in time?”.

**\$96.2B**

NVIDIA quarterly revenue, \$89B of it data centre — with \$108B guided for next quarter and zero assumed from China

**\$638B**

Oracle remaining performance obligations, up \$85B in a quarter, on FY26 revenue of \$55.7B

**\$514B**

Alphabet cloud backlog at Q2 2026 — contracted work not yet delivered

**58%**

Synergy's compound growth rate for GPU cloud (“neocloud”) revenue: ~\$25B today toward ~\$400B

**2026 capital spending guidance, big four (US\$ billions, midpoint)**

|                  |                                                                                     |            |
|------------------|-------------------------------------------------------------------------------------|------------|
| Amazon           |    | ~\$220B    |
| Alphabet         |    | \$195–205B |
| Microsoft (FY26) |    | ~\$190B    |
| Meta             |  | \$130–145B |

Post-July guidance. Microsoft has guided FY27 (to June 2027) at \$255–260B. Combined big-four 2026 spend tracks ~\$720–745B, up from ~\$600B expected at the start of the year.

## What the numbers say in plain language

Four companies will spend more on data centres this year than the entire US commercial real-estate market changes hands in a normal year. Their suppliers are not guessing at demand: Oracle and Alphabet together hold more than \$1.1 trillion of signed-but-undelivered contracts. The bottleneck has moved from the customer to the construction site — which is why the rest of this issue is about land, power and the cost of a gigawatt.

## Supply is sold out

TSMC's CoWoS advanced-packaging capacity — the step that turns a chip into an AI accelerator — is sold out through 2026 and tight into 2027. NVIDIA is guiding ~70% growth for its fiscal 2028, and describes itself as supply-constrained.

## Rents are holding

The GetDeploying GPU price index was up 1.9% year on year at 31 August. Two years into a build-out of this size, the fact that rental prices have not fallen is the single clearest sign that demand is running ahead of delivered capacity.

## The reset that wasn't a retreat

OpenAI trimmed its headline compute commitments from ~\$1.4 trillion to ~\$600 billion. Even the smaller figure is more than the market capitalisation of all but a dozen companies — and it is contracted, not aspirational.

## 03 — POWER

# The binding constraint is a grid connection

A GPU can be bought in months. A substation takes years. That gap is the reason “compute real estate” is a property story and not a semiconductor story: the scarce input is land that already has power attached, and the people who own it now set the pace of the whole industry.

**485 → 950**

TWh of global data-centre electricity, 2024 to 2030 (IEA) — roughly a doubling in six years

**9.8 GW**

Nuclear capacity contracted by AI and cloud operators across 13 announced deals

**~0.3 GW**

Stargate capacity actually live at 31 August, against 9+ GW announced

**~\$130B**

Proposed data-centre projects facing local opposition in Q1 2026 alone (75 projects; Data Center Watch)

## Why power and not chips

Epoch AI counts about 30 GW of frontier AI compute today, doubling roughly every seven months. US electricity generation is growing by around 120 TWh a year (2.8% in 2025) — healthy by historical standards, but a fraction of what that doubling implies. The result is a queue: projects wait not for GPUs or money but for the utility to say yes. Any site with a signed interconnection agreement is worth more than the building on it.

## Buying the power plant

Operators have stopped waiting for the grid. Amazon has contracted ~2.9 GW of nuclear output (1.92 GW from Talen's Susquehanna plant and 960 MW of X-energy small modular reactors). Meta's 1 GW-plus Prometheus campus in Ohio is slated to come online in 2026 with on-site gas. Battery storage is being added at 20–25 GW a year globally (IEA) to smooth the load. In each case the operator is behaving like a landlord securing utilities for a tenant, not like a software company.

## The opposition is organised

Data Center Watch tracked \$162B across 36 projects blocked or delayed by local opposition by mid-2025; by Q1 2026 the tally was ~\$130B across 75 projects in a single quarter. Water use, noise and electricity bills are the objections. Jurisdictions with land, water and spare generation — and public support for using them — become the scarce resource.

## The Gulf angle

Every constraint on this page is an advantage for the region. Cheap, abundant generation (Saudi Arabia's record solar tariff is ~\$0.0104/kWh), land without neighbours to object, and permitting under a single authority. HUMAIN's targeted 6.6 GW would be, on its own, more than a fifth of today's global frontier fleet. The question is no longer whether the Gulf can host compute, but on what terms it leases it.

**Plain-language summary.** Think of AI capacity as a factory that needs a river. There are plenty of factories for sale; there are very few rivers. Owning the river — powered land with a grid connection — is what this asset class is really about.

04 — WHAT ONE GIGAWATT COSTS **SNAPSHOT · AUGUST**

# The bill, the rent and the payback

Real estate is valued on three numbers: what it costs to build, what it earns in rent, and how long the tenant stays. Compute real estate can now be valued the same way — and the numbers are large enough to explain every financing structure in the rest of this issue.

**~\$38B**

All-in cost of one gigawatt of frontier AI capacity (Epoch AI) — chips, networking, building and power

**~\$8.5B**

Of that, the “real estate” — land, shell, power and cooling per gigawatt


**~\$12B/yr**

Annual cost to carry one gigawatt at a three-year depreciation life (Epoch AI)

**\$4.00/hr**

H100 rental at a neocloud, 31 August (GetDeploying); hyperscalers charge ~93–97% more

What an hour of an H100 costs, by channel (US\$, August 2026)

|                          |                                                                                   |              |
|--------------------------|-----------------------------------------------------------------------------------|--------------|
| Hyperscaler on-demand    |  | ~\$7.75–7.90 |
| Neocloud on-demand       |  | \$4.00       |
| Spot / marketplace floor |  | ~\$1–2       |

Hyperscaler figure derived from the ~93–97% premium in the GetDeploying index. Spot floor is the lowest observed marketplace price, not a contracted rate. Index +1.9% YoY.

## The rent, in plain language

An H100 renting at \$4.00 an hour, busy 85% of the time, earns roughly \$30,000 a year (our arithmetic). A tenant paying hyperscaler rates pays nearly double for the same chip — the premium buys integration, security and a single invoice. Marketplaces where idle capacity is sold by the hour set the floor at \$1–2; nobody builds a gigawatt for the floor price, which is why long contracts matter more than list prices.

## The payback model

Mercatus AI's model (May 2026) frames the economics the way a landlord would: the building lasts 15–20 years, the chips inside it three to six, and the lease that pays for both typically runs **two to seven years**. The operator's return therefore depends on three things — utilisation, the contracted rate, and whether the tenant renews before the chips are obsolete. The bulk of the capital is debt-financed against that contract, so a strong tenant lowers the cost of every gigawatt.

## Why the shell is the cheap part

Only about \$8.5B of the \$38B gigawatt is bricks, power and cooling. That inversion — the fit-out worth four times the building — is what makes this asset class new. It rewards owners who can refill the building with the next generation of chips, and it punishes anyone who financed the chips over the life of the building. Every structure in Section 05 is an attempt to match the money to the right lifespan.

**Rule of thumb for the reader:** a gigawatt is roughly a \$40 billion building with a \$12 billion annual carrying cost, let to a handful of tenants who each need it for two to seven years. Price the tenant, not the chip.

## 05 — THE NEOCLOUD MODEL

# Contract → debt → deploy → amortize

A “neocloud” is a company that does one thing: it rents GPUs. Strip away the branding and the business is a property developer's — find a tenant, borrow against the lease, build fast, and earn back the capital before the asset ages. The four steps below are the whole industry.

1

**Contract**

An anchor tenant — a hyperscaler or AI lab — signs a take-or-pay agreement for capacity, typically two to seven years. This is the lease, and it is signed before the building exists.

2

**Debt**

Lenders advance most of the capital against the contract and the GPUs themselves. Rates have run from ~10% on early facilities to investment-grade pricing on the best-secured deals.

3

**Deploy**

Prefabricated halls and pre-negotiated power turn a signed site into a live cluster in months rather than years. Speed to revenue is the operator's main skill.

4

**Amortize**

The chips earn rent for three to six years. The contract must repay the debt before the hardware needs replacing — or the operator must refill the building with the next generation.

## TWO WAYS TO OWN THE SAME ASSET

**Asset-heavy: own the whole stack**

**Who:** CoreWeave, Nebius, Crusoe. **How:** the operator controls the sites (owned or on long leases), buys the chips, raises the debt and carries the depreciation. **Reward:** keeps the entire margin and the terminal value of the land and power. **Cost:** balance sheets that now rival utilities — CoreWeave carries ~\$35B of debt — and full exposure if a tenant does not renew. Capital spending guidance for the two listed leaders alone is \$55–64B for 2026.

**Asset-light: bring the tenant, not the building**

**Who:** Fluidstack and a growing set of operators built on the same idea. **How:** a landlord — often a former bitcoin miner with powered land — owns the site and builds the hall; a hyperscaler or the tenant backstops the landlord's lease; the operator designs, fills and runs the cluster and holds the customer relationship. **Reward:** hyperscale capacity without hyperscale debt. **Cost:** the operator shares the economics with the landlord and the guarantor.

**The landlord**

Owns land, power and the building. Increasingly a listed “digital infrastructure” company converting from crypto mining, earning a long-dated lease guaranteed by a large tenant. Behaves like a REIT.

**The operator**

Designs the cluster, procures the GPUs, runs the software and holds the customer. Earns the spread between the contracted rate and its costs. Behaves like an operating company.

**The guarantor**

A hyperscaler or AI lab that backstops the landlord's lease so the building can be financed cheaply — often in exchange for warrants or priority access. Behaves like an anchor tenant in a mall.

**The plain-language version.** Somebody owns the shopping centre, somebody runs the shops, and a big anchor tenant guarantees the rent so the bank will lend. The heavy model does all three jobs at once; the light model splits them among specialists. Section 10 shows the light model in practice.

06 — LEAGUE TABLE **SNAPSHOT · AUGUST**

# Who owns the ground

Six operators now define the independent compute market. Two are listed, three are private and pricing for the public markets, and one — Fluidstack — has chosen the asset-light path described in Section 05. Figures are the latest public disclosures at 31 August 2026.

| OPERATOR                            | MODEL                                                                           | SCALE SIGNAL                                                                                 | ANCHOR CUSTOMERS (PUBLIC)                              | CAPITAL POSITION                                                                   |
|-------------------------------------|---------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|--------------------------------------------------------|------------------------------------------------------------------------------------|
| <b>CoreWeave</b><br>(NYSE:<br>CRWV) | Asset-heavy; owns and long-leases sites across the US and Europe                | H1 2026 revenue \$4.65B (+112%); \$104B contracted backlog                                   | Microsoft, OpenAI, Meta                                | Debt \$35.6B; first A3-rated \$8.5B GPU-backed facility; 2026 capex guide \$35–39B |
| <b>Nebius</b><br>(Nasdaq:<br>NBIS)  | Asset-heavy; European-rooted, building in the US and Nordics                    | H1 2026 revenue \$981M (+529%); ARR \$3.0B; \$40B of committed backlog                       | Microsoft, Meta                                        | Debt \$8.5B vs cash \$8.0B; 2026 capex guide \$20–25B                              |
| <b>Crusoe</b><br>(private)          | Asset-heavy; energy-first developer (Abilene, Texas campus for Stargate)        | Multi-gigawatt pipeline; a ~\$30B-valuation round was reported but unconfirmed at press time | Oracle / OpenAI (Stargate)                             | Project-finance heavy; no public filing as of 31 August                            |
| <b>Lambda</b><br>(private)          | Asset-heavy; developer-facing cloud plus large hyperscaler contracts            | Last priced at \$5.4B (November 2025)                                                        | Microsoft                                              | Added ~\$1B of GPU-backed debt on 28 August; IPO preparation reported              |
| <b>Together AI</b><br>(private)     | Inference-first cloud; rents capacity rather than owning sites                  | Valued at \$8.3B; strong open-model developer base                                           | Developers and open-model labs                         | Equity-funded; lighter balance sheet by design                                     |
| <b>Fluidstack</b><br>(private)      | Asset-light; landlords own sites, operator brings tenant, design and operations | 100,000+ GPUs under management; 2,500 GPUs deployed in 48 hours; ~6-month prefab builds      | Anthropic, Mistral, Meta and other publicly named labs | Terms not disclosed in this publication — see Section 10                           |

## What the table shows

Contracted backlog is the number that matters, and the two listed leaders alone hold ~\$144B of it. Backlog is the lease roll of this asset class: it tells you how much rent is already promised before a single new building is financed. Private operators are now pricing off the same metric, which is why the public scorecard in Section 07 sets the valuation benchmark for the whole group.

## Customer concentration is the quiet risk

Every operator on this list depends on a handful of tenants, and several on one. That is normal for a new property market — the first office towers in any district are anchored by a single large lessee — but it means a tenant's decision to build its own capacity, or to renegotiate, moves an operator's value more than any technology change. Section 11 returns to this.

07 — PUBLIC NEOCLOUDS SCORECARD **SNAPSHOT · 31 AUGUST**

# What the stock market pays for compute real estate

**~\$6.3B**

Combined H1 2026 revenue of the four listed neoclods

**~\$125B**

Combined market capitalisation at 31 August

**\$55–64B**

2026 capital-spending guidance, CoreWeave and Nebius alone

**~\$56B**

Combined debt across the four

| METRIC                           | COREWEAVE (CRWV)                                                       | NEBIUS (NBIS)                                     | IREN                                                       | APPLIED DIGITAL (APLD)                                      |
|----------------------------------|------------------------------------------------------------------------|---------------------------------------------------|------------------------------------------------------------|-------------------------------------------------------------|
| <b>H1 2026 REVENUE</b>           | \$4.65B (+112%)                                                        | \$981M (+529%); ARR \$3.0B                        | \$282M (AI cloud \$104M)                                   | \$385M, Dec–May half (~+270%)                               |
| <b>Q2 ADJ. EBITDA MARGIN</b>     | 59%                                                                    | 41%                                               | 14% (fiscal Q4)                                            | 16% (fiscal Q4)                                             |
| <b>Q2 GAAP NET INCOME</b>        | –\$626M (H1 –\$1.37B)                                                  | –\$190M (H1 +\$431M on a \$781M revaluation gain) | –\$684M incl. ~\$550M mining impairments                   | –\$111M                                                     |
| <b>FY26 GUIDANCE &amp; CAPEX</b> | Rev \$12.4–13.2B; adj. OI \$0.96–1.15B; capex \$35–39B (raised 11 Aug) | Rev \$3.0–3.4B; YE ARR \$7–9B; capex \$20–25B     | Contracted ARR \$1B → >\$4B by Dec-26; FY27 capex \$25–30B | \$1B NOI run-rate targeted within a year                    |
| <b>BACKLOG &amp; ANCHORS</b>     | \$104B backlog                                                         | >\$40B commitments                                | Microsoft \$9.7B; NVIDIA \$3.4B                            | \$36B of leases / 1.41 GW; one hyperscaler >80% (CoreWeave) |
| <b>DEBT / CASH</b>               | \$35.1B / \$5.5B                                                       | \$8.5B / \$8.0B                                   | \$7.6B / \$7.6B                                            | \$5.0B / \$4.2B                                             |
| <b>SHARE PRICE, 31 AUG</b>       | \$84.89 (+19% YTD; +112% vs \$40 IPO)                                  | \$206.32 (+147% YTD)                              | \$37.12 (–2% YTD)                                          | \$25.41 (+4% YTD)                                           |
| <b>MARKET CAP</b>                | ~\$47B                                                                 | ~\$56B                                            | ~\$15B                                                     | ~\$7.4B                                                     |

## Adjusted profit, GAAP loss — both are true

Every name on this page is profitable before depreciation and loss-making after it. That is not an accounting trick; it is the asset class. A landlord who has just built a tower shows the same pattern in year one. The question investors are really asking is whether the chips will still be earning rent when the depreciation stops — which is a question about contract length and renewal, not about margins.

## What the share prices say

The market pays for contracted backlog and punishes concentration. Nebius, with balanced debt and cash, has more than doubled; CoreWeave trades at roughly twice its IPO price despite \$35B of debt because \$104B of rent is signed; IREN and Applied Digital, still carrying mining legacies and single-tenant exposure, have gone sideways. Compute real estate is being priced like real estate: on the quality of the lease.

Data as of 31 August 2026 from company filings and results releases. Market capitalisation computed as closing price × shares outstanding. IREN and Applied Digital report on June and May fiscal years respectively; calendar-half figures are derived. Fluidstack's landlords are deliberately excluded. Lambda and Crusoe remain private with no public filing at the snapshot date.



# Rent, mortgage, forward curve

A property market is complete when three things exist: a lease, a lender who will take the property as security, and a forward market where the price of future space can be traded. Compute acquired the third of these this summer. In the space of one quarter it stopped being an IT budget line and became a capital market.

**>\$20B**

GPU-backed debt outstanding, up from ~\$11B a year earlier

**A3**

Rating on CoreWeave's \$8.5B facility — the first investment-grade loan secured on GPUs

**\$114.8B**

Analyst tally (Electron Economics) of NVIDIA-linked backstops, guarantees and commitments across the ecosystem

**Oct 2026**

CME's targeted launch of exchange-traded compute futures; CFTC consultation completed August

## The mortgage

GPU-backed lending has moved from specialist credit funds to the mainstream. Early facilities were priced like risky project finance — one CoreWeave loan carried a 10.44% coupon. Today the lender list includes Blackstone, Magnetar, BlackRock, PIMCO, JPMorgan and Macquarie, and the best-secured facility is rated A3. Lambda added ~\$1B of GPU-backed debt on 28 August. The security package is the same as a commercial mortgage: the asset plus the lease.

## The guarantee

Behind many leases stands a bigger name. NVIDIA has guaranteed ~\$105B of capacity at the Stargate campus in Ohio; hyperscalers backstop landlords' leases in exchange for warrants; AI labs sign take-or-pay contracts that let a small operator borrow at a large operator's rate. An analyst tally puts NVIDIA-linked support at \$114.8B. This is the credit umbrella that makes the asset-light model possible — and the exposure regulators will study first.

## The forward curve

ICE and Ornn have launched compute futures; the CME targets October for its contract, after a CFTC consultation in August. A forward price for a GPU-hour lets an operator hedge the rent on capacity it has not yet built, and lets a lender mark its collateral. Every mature property market has a forward curve in some form; compute now has one in the most literal sense.

## The bear case, stated fairly

Michael Burry's critique — that the largest operators are understating depreciation by an estimated \$176B over 2026–2028 by extending chip lives — is the correct question to ask of any leveraged property market: is the rent roll longer than the asset life? OpenAI's decision to scale its stated commitments from ~\$1.4T to ~\$600B is the other side of the same coin. Neither changes the direction of the market; both change the price at which it should be financed.

## Why this matters for sovereign capital

Gulf institutions are natural lenders and guarantors in this structure. They already own the power; providing the **credit umbrella** — lease guarantees, investment-grade facilities, anchor commitments — converts that advantage into a seat at the top of the capital stack, with real assets as security. The financing layer, not the chip, is where sovereign capital is best placed.

## 09 — FRONTIER VENTURES

# Where the next gigawatts come from

If powered land is the constraint, the most interesting ventures are the ones that find power where nobody else is looking — in orbit, in idle machines, in sovereign campuses built from scratch. Four categories, one Singular Link portfolio company.

## Orbital compute — Starcloud SL portfolio

Starcloud is building data centres in space, where solar power is continuous and cooling is free by radiation. It has raised **\$450M** to date and has flown its first hardware. The thesis is the purest form of this issue's argument: if the scarce input is powered land, the answer may be to remove the land. It is early, and the launch-cost curve decides the timing — but it is the only venture that escapes the grid queue entirely.

## Sovereign campuses — HUMAIN and Stargate UAE

HUMAIN, backed by Saudi Arabia's Public Investment Fund, targets **6.6 GW** of capacity behind ~\$77B of planned investment, with ~\$23B in its first phase; its Riyadh and Dammam campuses are targeted for 2026, with commissioning not yet publicly confirmed at the snapshot date. Abu Dhabi's Stargate UAE is building a 1 GW campus with OpenAI and partners. Both sit on the cheapest utility-scale power in the world — the Kingdom's record solar tariff is ~\$0.0104/kWh.

## Marketplaces — the spot market for compute

SF Compute (\$40M raised) runs an exchange where idle clusters are sold by the hour, setting the ~\$1–2 spot floor referenced in Section 04. Prime Intellect (\$130M) goes further, pooling distributed machines to train models across the internet. Neither owns a building; both make the asset class more liquid — the equivalent of short-let platforms for property.

## Inference clouds — renting the finished product

Together AI (\$8.3B valuation) and Groq (a \$1.5B Saudi commitment, announced February 2025) sell model output rather than raw GPU time. They are the tenants of compute real estate as much as its operators — and their growth is the demand signal that keeps the rent roll full. Issue 02 covered why inference volume is compounding faster than prices fall.

**\$450M**

Starcloud total raised — orbital data centres

**6.6 GW**

HUMAIN capacity target — the largest sovereign compute plan announced

**1 GW**

Stargate UAE campus under construction in Abu Dhabi

**\$0.0104**

Per kWh — Saudi Arabia's world-record solar tariff

**The Singular Link view.** The venture opportunity in compute is not another neocloud; the listed leaders already have the balance sheets for that. It is in the layers that make the asset class work — new sources of power, liquidity for the spot market, and the financing and operating structures that let sovereign capital own compute without operating it. That is the lens for Section 10.

## 10 — COMPANY FOCUS

# Fluidstack: the landlord's partner

**Fluidstack in one paragraph.** Founded in 2017 by Oxford students who began by aggregating idle GPU capacity, London-headquartered Fluidstack is now one of the largest independent AI cloud operators. Its distinctive choice is to own almost none of the ground it runs on: landlords own the sites and power, large tenants and hyperscalers guarantee the leases, and Fluidstack brings the customer, the design and the operations.

**100k+**

GPUs under management for AI labs and enterprises

**48 hrs**

Time to deploy 2,500 GPUs for a customer — a public marketing benchmark

**~6 mo**

Typical build time for a prefabricated, purpose-built AI hall

## THREE MOVES THAT DEFINE THE MODEL

**I From aggregator to operator**

Fluidstack's origin was software that pooled under-used GPUs from data centres and sold them by the hour. That taught it two things the asset-heavy operators learned later and more expensively: how to fill capacity fast, and how to run other people's hardware reliably. The company's name for its business — "AI cloud platform" — is really a description of an operating layer that sits on top of any site.

**II The landlord partnership**

Rather than buying land, Fluidstack partners with owners of powered sites — frequently listed infrastructure companies that converted from bitcoin mining and hold grid connections that would take years to replicate. The landlord builds the hall to Fluidstack's specification and earns a long-dated lease; Fluidstack fills it. Speed comes from the fact that the hardest asset — the interconnection — already exists.

**III The credit umbrella**

A landlord cannot borrow cheaply against a lease from a private operator. The model solves this by bringing in a guarantor — a hyperscaler or the end tenant — that backstops the lease, so the building is financed at an investment-grade cost of capital. The guarantor gains priority capacity; the landlord gains a bankable tenant; Fluidstack gains hyperscale capacity without hyperscale debt.

**Who it builds for**

Publicly named customers include Anthropic — which named Fluidstack as a partner in its US data-centre programme — Mistral, Meta and a set of leading open-model labs. The customer list is the clearest evidence for the model: labs that could rent from anyone choose an operator with no balance sheet of its own, because what they are buying is speed and operating competence, not real estate.

**Discipline in what it walks away from**

In March 2026 the company stepped back from a large European campus project it had announced a year earlier, choosing to concentrate on North American sites where power and guarantors were in place. For an asset-light operator that is the right reflex: it should never be the party carrying a gigawatt of unfunded land.

Fluidstack's commercial terms, valuation, contract values, site details and financing arrangements are not disclosed in this publication. Facts above are drawn from the company's public materials and named-customer announcements.

## 10 — COMPANY FOCUS (CONTINUED)

# Why Fluidstack matters for Gulf capital

The asset-light model is, in effect, a template for how a region with power, land and capital can own compute real estate without becoming a technology operator. Each role in the structure maps onto something the Gulf already has.

## The landlord role

Sovereign and semi-sovereign entities own the two scarcest inputs — powered land and permitting authority. In the Fluidstack structure the landlord earns a long-dated, guaranteed lease and keeps the terminal value of the site. That is a real-estate return on an energy advantage, with an operator handling everything inside the building.

## The guarantor role

The credit umbrella needs an investment-grade name willing to backstop leases in exchange for priority capacity or equity-like upside. Gulf institutions are among the few balance sheets in the world that can play that role at gigawatt scale — and it places them at the top of the capital stack with real assets as security.

## Speed as the differentiator

A six-month prefabricated build on a site with existing power is the fastest path from decision to revenue in the industry. For national programmes that have announced capacity targets on public timelines, an operator that can fill a hall in months — rather than a developer that will still be pouring concrete in two years — is the difference between a headline and a rent roll.

## Insulation from the chip cycle

Because Fluidstack does not carry the hardware on its own balance sheet, it is less exposed to the depreciation question raised in Section 08. When the next generation of chips arrives, the operator refills the building; the landlord's lease continues; the guarantor's capacity is upgraded. The party that owns the ground is the party that benefits from every refresh.

## THE SINGULAR LINK VIEW

The largest neoclouds have proven that compute is an asset class; Fluidstack is proving that it can be owned the way real estate is owned — by separating the land, the operations and the credit into roles that different specialists play best. For Gulf capital, which holds the land and the credit and does not want to run the operations, that separation is the whole point. We regard the asset-light model as the most direct bridge between the region's structural advantages and the frontier of AI infrastructure — and we have positioned accordingly.

## What to watch

Renewal terms on the first long-dated leases; whether guarantors extend the umbrella to new regions; and the pace of North American builds through 2026.

## What could go wrong

Guarantor concentration — the model depends on a few very large names; a change in their appetite changes the operator's cost of capital overnight.

## Our stance

Constructive. The model is the right shape for sovereign participation; the execution record on speed is public and verifiable.

Fluidstack's commercial terms are not disclosed in this publication. Company names shown, including current portfolio positions, are illustrative of Singular Link's thematic focus and not a recommendation.

## 11 — RISKS &amp; WATCH ITEMS

# What could change the picture

## Asset life versus lease length

Chips are depreciated over three to six years; leases run two to seven. If a tenant does not renew, the operator is left with ageing hardware and a building to refill. Burry's estimate of \$176B of understated depreciation over 2026–2028 is the sharpest statement of this risk.

## Customer concentration

Several operators depend on one tenant for most of their revenue — Applied Digital's largest hyperscaler accounts for more than 80% of leases. A tenant building its own capacity, or renegotiating, matters more than any technology change.

## The grid and the neighbours

Interconnection queues stretch years and local opposition touched ~\$130B of projects in a single quarter. Sites already announced may slip: Stargate has ~0.3 GW live against 9+ GW announced, and HUMAIN's 2026 campuses were not yet publicly commissioned at the snapshot date.

## Commitments can be re-sized

OpenAI's stated commitments moved from ~\$1.4T to ~\$600B. The reset did not reduce demand for delivered capacity, but it shows that headline numbers in this market are ambitions until they are contracts — and contracts until they are backed by an investment-grade name.

## Rent could fall

The GPU price index is only +1.9% year on year and the spot floor is \$1–2/hour. When CoWoS constraints ease after 2027 and the next chip generation lands, older hardware reprices first. Operators on short contracts feel it before landlords on long ones.

## Policy and geopolitics

NVIDIA now assumes zero revenue from China. Export rules, tariffs and sovereign-AI procurement policies can redirect where gigawatts are built — a risk for some regions and an opening for those, like the Gulf, positioned as trusted hosts.

### WATCH ITEMS FOR ISSUE 04

- ✓ **October:** CME compute futures launch — the first exchange-traded forward price for a GPU-hour.
- ✓ **Q3 results (November):** CoreWeave and Nebius backlog conversion; whether adjusted margins hold as depreciation rises.
- ✓ **Private-market pricing:** Crusoe's reported round and Lambda's IPO preparation set the private benchmark against the public scorecard.
- ✓ **Gulf commissioning:** first public confirmation of HUMAIN campus energisation and Stargate UAE's first 200 MW.

**Sources.** Company results releases and filings (NVIDIA, Oracle, Alphabet, Amazon, Microsoft, Meta, CoreWeave, Nebius, IREN, Applied Digital) to 31 August 2026; Epoch AI frontier compute and gigawatt cost estimates; IEA *Energy and AI* and battery-storage data; GetDeploying GPU Price Index (31 August 2026 snapshot); Synergy Research Group; Data Center Watch (local opposition tallies); Mercatus AI payback model (May 2026); Electron Economics backstop tally; CME Group, ICE and Ornn product announcements; CFTC consultation (August 2026); HUMAIN and Stargate UAE public statements; Fluidstack and customer public announcements. Analysis, the compute-real-estate framing and the Gulf-capital lens are Singular Link's own. Crusoe's reported round was unconfirmed at the snapshot date.

This material is provided by Singular Link for informational purposes only. It is not investment advice, nor an offer to sell or a solicitation of an offer to buy any security or interest in any fund. Company names shown, including current portfolio positions, are illustrative of Singular Link's thematic focus and should not be taken as a recommendation, a complete list of holdings, or an indication of future portfolio composition. Terms of private portfolio positions are not disclosed in this publication. Past performance is not indicative of future results.