

Scoring the Stack: Who Owns What in AI Infrastructure

14 October 2026 · Tugay Dilikoğlu · Information cutoff: 14 October 2026

Thirteen companies scored on eight layers, from power to inference APIs, and what their position means for customer concentration, hardware life and valuation.

By **Tugay Dilikoğlu** · Follow-up to From Megawatts to Equity Value (2 October 2026)

Investment takeaways

- **Business model and physical depth are different questions.** CoreWeave and Nebius are AI cloud platforms that own their GPUs and software but lease most of their floor space (2.5 of 4 on physical integration). Crusoe runs a similar platform on land and power it controls (3.5).
- **A full software stack has not diversified revenue.** The disclosed examples show substantial counterparty exposure, measured against different revenue, capacity and contract bases. Nebius's estimate is an upper bound, not a minimum concentration level.
- **Hardware life matters most where price per megawatt is lowest.** At the economics of IREN's Microsoft contract, moving from a 6-year to a 4-year GPU life cuts the margin before operating costs from 47% to 22% (derived, illustrative). At the \$20–25 million of annual contract value per MW that Nebius reports, the same change matters far less. We could not reconcile the two bases.
- **Controlling power is not delivering it.** IREN had about 40 MW of AI cloud capacity operating against about 5 GW of grid agreements at 30 June 2026. CoreWeave had 1.5 GW active against 3.7 GW contracted.


From classification to measurement

In From Megawatts to Equity Value we classified AI infrastructure companies by what they own and sell: data-center landlords, hosting providers, GPU infrastructure and AI cloud platforms. We also noted that classification should follow the project, because one company can span several models.⁵⁶ This note adds a second dimension: how much of the stack each company controls, from the power connection to the inference API.

We scored thirteen companies on eight layers. Four are physical: power, land, the data-center shell, and the GPU fleet. Four are software: cluster orchestration, a managed platform, inference APIs, and developer tooling. Each layer scores 1 if the company owns or runs it, 0.5 if partly, and 0 if it relies on others. A company's horizontal position is the sum of its physical scores; its vertical position is the sum of its software scores. Color shows the company's primary business model under our existing classification. The full score table is in Exhibit 2.

EXHIBIT 1

Integration map: position is integration depth, color is business model



Derived: $x = P1+P2+P3+P4$, $y = S1+S2+S3+S4$, scores from Exhibit 2. **Business-model colors:** Dilikoglu Research classification;⁵⁶ each company's primary model, an analyst judgement. Points that share a position are drawn side by side. IREN's pre-Mirantis position is shown hollow.

Same business model, different depth

The largest AI cloud platforms lease their buildings. CoreWeave reports 3.7 GW of contracted power "via a diversified provider portfolio".⁵ It owns one site, in Kenilworth, New Jersey,⁹ and its Lancaster, Pennsylvania campus is a third-party build with CoreWeave as tenant.¹⁰ Its attempt to buy its landlord Core Scientific ended in October 2025 when Core Scientific's shareholders voted against it.⁸ Nebius owns its Finnish site and land in Kansas City but leases its New Jersey site and others; management describes the model as asset-light.^{12,13} Both have full software stacks: CoreWeave through Weights & Biases and a serverless inference API,^{6,11} Nebius through Token Factory.¹⁴ Crusoe sells the same kind of platform from a different base: about 360 MW of on-site gas turbines at Abilene and its own construction capability.^{18,20}

Landlords are converging on one point. Applied Digital, Galaxy, TeraWulf, Cipher and Hut 8 control power, land and halls, and leave software to tenants. Hut 8, for example, has 949 MW of contracted IT capacity across Beacon Point and River Bend after its July 2026 lease expansion. Contracted capacity is distinct from energized or revenue-generating capacity.^{54,55} Applied Digital and Hut 8 score half a point for GPUs through small GPU-as-a-service units.^{29,54}

GPU infrastructure is the model in motion. IREN bought Mirantis, a Kubernetes and AI-orchestration company, in a deal that closed on 4 August 2026; its software score rose from 0.5 to 1.5.² Lambda moved the other way and shut down its inference API in September 2025.²² It now sits beside Fluidstack, which we classify as a hosting provider: Fluidstack operates clusters in buildings it leases or develops, and "rents the space to host chips rather than buying them directly".²⁸

EXHIBIT 2 Layer scores behind every point


COMPANY	MODEL	P1	P2	P3	P4	S1	S2	S3	S4	PHYS.	SOFT.	SOURCES
CoreWeave	AI cloud platform	0.5	0.5	0.5	1	1	1	1	1	2.5	4.0	5, 6, 9, 10, 11
Nebius	AI cloud platform	0.5	0.5	0.5	1	1	1	1	1	2.5	4.0	12, 13, 14, 48
Crusoe	AI cloud platform	1*	0.5	1	1*	1*	0.5	1	0.5*	3.5	3.0	18, 19, 20, 51
IREN	GPU infrastructure	1	1	1	1	1	0.5	0	0	4.0	1.5	1, 2, 50
Together AI	AI cloud platform	0*	0.5*	0.5*	0.5	1*	0.5*	1	1	1.5	3.5	25, 26
Vultr	AI cloud platform	0	0	0	1	1*	0.5*	1	0.5*	1.0	3.0	27
Lambda	GPU infrastructure	0	0.5*	0.5*	1	1	0.5	0	0	2.0	1.5	22, 23, 24, 49
Fluidstack	Hosting provider	0.5*	0.5	0.5	0.5*	1	0.5	0	0	2.0	1.5	28, 52
Applied Digital	Data-center landlord	1*	1	1	0.5*	0*	0	0	0	3.5	0.0	29, 30
Hut 8	Data-center landlord	1*	1*	1	0.5*	0	0	0	0	3.5	0.0	54, 55
Galaxy	Data-center landlord	1	1*	1	0*	0	0	0	0	3.0	0.0	31, 32
TeraWulf	Data-center landlord	1	1*	1	0	0	0	0	0	3.0	0.0	33, 34
Cipher Digital	Data-center landlord	1*	1	1	0	0	0	0	0	3.0	0.0	35, 36, 37

Rubric: P1 Power, owns generation or holds grid interconnection in its own name. P2 Land, owns or controls the site. P3 Shell, designs and builds its own halls. P4 GPU fleet, owns or finances the accelerators. S1 Orchestration, its own cluster provisioning, managed Kubernetes or Slurm. S2 Platform, managed storage, observability and job tooling. S3 Inference, serverless or per-token endpoints. S4 Developer tools, fine-tuning, experiment tracking, model catalog. * marks 26 of 104 cells that rest on secondary sources or incomplete disclosure.

Software capability is not revenue diversification

The vertical axis measures what a company can sell, not what it does sell. Most revenue in this group still comes from multi-year reserved-capacity contracts or long leases with a small number of hyperscalers and AI labs. Exhibit 3 shows the largest disclosed counterparty for each listed company. The bases differ, so the comparison is about order of magnitude, not rank.

EXHIBIT 3 Largest single counterparty, share of disclosed base



Bases: Galaxy, CoreWeave as sole Helios tenant³² (reported). Hut 8, unnamed investment-grade tenant, 704 of 949 leased IT MW⁵⁴ (derived). Applied Digital, unnamed hyperscaler, 810 of 1,410 leased MW²⁹ (derived). CIPHER, AWS, 300 of 600 gross MW³⁵ (derived; a third, undisclosed lease would lower it). IREN, Microsoft \$9.7bn ÷ 5 years = \$1.94bn a year against \$4bn contracted ARR^{3,4} (derived; assumes even recognition). TeraWulf, Anthropic about 401 of about 839 leased MW^{33,34} (derived). CoreWeave, "Customer A", 45% of Q1 2026 revenue⁷ (reported). Nebius, Microsoft \$17.4bn against backlog above \$40bn^{15,17} (derived upper bound; backlog from a secondary summary).


Concentration is still high. CoreWeave's largest customer fell from 72% to 45% of quarterly revenue in a year, but its two largest still made up 65% in Q1 2026.⁷ Risk also runs down a chain. Galaxy's only AI tenant is CoreWeave, and CoreWeave leases a further 400 MW at Applied Digital's Polaris Forge.²⁹ A landlord's 15-year lease is only as strong as the tenant's ability to keep re-pricing GPUs, so for these landlords hardware risk is one step removed, not absent.

Hardware life is where the models diverge

CoreWeave depreciates technology equipment over six years.⁴³ Nebius is reported to use four (secondary source; not confirmed in its annual report).⁴⁴ Among hyperscalers, Microsoft depreciates servers over two to six years and Alphabet over six; Meta moved to 5.5 years in 2025 and Amazon shortened part of its fleet from six to five.^{53,45,46} We did not find IREN's GPU useful life in its FY2026 10-K.

Exhibit 4 shows illustrative economics per megawatt-year for three models, anchored on two real contracts. It isolates depreciation: power, staff, overheads, interest and tax are excluded.

EXHIBIT 4 Margin before operating costs vs. GPU useful life, per MW-year



INPUT	VALUE	TYPE	BASIS
GPU-cloud revenue	\$9.70m / MW-yr	DERIVED	IREN–Microsoft: \$9.7bn ÷ 5 yrs ÷ 200 MW ³
GPU and network capex	\$29.0m / MW	DERIVED	\$5.8bn Dell GB300 order ÷ 200 MW ³ ; rack-price estimates imply \$36–43m ⁴²
Shell capex, owner-built	\$6.5m / MW	REPORTED	IREN Horizon 1 estimate, \$6–7m per MW (May 2025) ⁴⁰
Shell capex, landlord	\$12.0m / MW	ESTIMATE	B. Riley estimate for Applied Digital (secondary) ⁴¹
Shell useful life	20 yrs	ASSUMPTION	Blended; Equinix uses 12–60 yrs for buildings, 3–40 for core systems ⁴⁷
Lease revenue (= tenant lease cost)	\$1.71m / MW-yr	DERIVED	Applied Digital: \$36.2bn ÷ 15 yrs ÷ 1,410 MW ³⁰

Formulas (L = GPU life in years): Landlord = $(1.71 - 12.0/20) \div 1.71$. Integrated = $(9.70 - 6.5/20 - 29.0/L) \div 9.70$. GPU owner leasing the shell = $(9.70 - 1.71 - 29.0/L) \div 9.70$. Whether the Applied Digital lease includes power is not disclosed. Illustrative sensitivity, not a forecast for any company. Depreciation life affects EBIT and net income, not EBITDA directly. Inputs span different projects and dates, so the comparison does not isolate ownership structure as the cause of the margin difference.

The landlord's margin does not move with GPU life because it owns no GPUs. The two GPU-owning models fall steeply as the assumed life shortens, and both are below zero before operating costs at three years. The integrated model holds up better than the tenant because it builds its shell at about half the landlord's estimated cost and does not pay a landlord's return.

On a five-year contract, a six-year GPU life is a bet that five-year-old chips will still earn money.

With a five-year contract and six-year depreciation, one sixth of the GPU cost, about \$4.8m per MW (derived: $\$29.0\text{m} \div 6$), remains on the balance sheet when the IREN–Microsoft contract ends. Recovering it depends on re-contracting end-of-term hardware at whatever the next generation leaves it worth. This is the asset-life and financing-horizon mismatch we described on 2 October, at the level of a single megawatt.⁵⁶

Price per megawatt changes the answer. Our 2 October note cited Nebius's Q2 2026 shareholder letter, which reported annual contract value of \$20–25 million per MW, and a CoreWeave disclosure of about \$40 million of annualized revenue per MW on contracts of three to six months.⁵⁶ Exhibit 5 reruns the integrated case at those contract values, holding cost per MW constant. We could not reconcile why IREN's contract implies less than half the Nebius figure. The MW definitions, the scope of services or the GPU generation may differ, so treat the higher rows as an upper bound, not a peer comparison. The CoreWeave figure covers short-term pricing and is not shown.

EXHIBIT 5

Integrated margin before operating costs by revenue per MW and GPU life

REVENUE PER MW-YEAR	3-YR LIFE	4-YR LIFE	5-YR LIFE	6-YR LIFE	7-YR LIFE	BREAK-EVEN LIFE
\$9.7m IREN–Microsoft (derived)	-3%	22%	37%	47%	54%	3.1 yrs
\$20.0m Nebius reported, low	50%	62%	69%	74%	78%	1.5 yrs
\$25.0m Nebius reported, high	60%	70%	76%	79%	82%	1.2 yrs


Derived: $\text{margin} = (R - 6.5/20 - 29.0/L) \div R$. Break-even life = $29.0 \div (R - 0.325)$. **Revenue rows:** \$9.7m, IREN–Microsoft (derived)³; \$20m and \$25m, the range of annual contract value per MW reported by Nebius (as cited in our 2 October note).⁵⁶ Holding GPU cost at \$29.0m per MW across rows is an assumption; if higher-priced contracts use denser or newer hardware, their margins are overstated.

Controlling power is not delivering it

The map scores control, not scale. IREN gets full marks for power while running a small fraction of what it controls. Exhibit 6 puts operating capacity next to contracted or controlled capacity on a log scale.

EXHIBIT 6 Operating vs. contracted capacity (log scale)

■ Operating / energized □ Contracted or controlled



Bases differ: CoreWeave, active vs contracted power, Q2 2026.⁵ IREN, AI cloud capacity operating at 30 Jun 2026 vs grid connection agreements, letters of agreement or equivalents;¹ its first 50 MW Horizon was delivered to Microsoft in August. Applied Digital, operating vs leased critical IT.³⁰ TeraWulf, revenue-generating critical IT vs leased (derived sum, 438 + 401 MW).^{33,38} Galaxy, delivered vs leased critical IT.^{32,38} Cipher, energized vs "operating and contracted", which may mix gross and IT MW.³⁹ Not shown: Hut 8 (597 MW leased, no AI capacity yet energized)⁵⁴ and Nebius (5 GW contracted-power target for year-end 2026, no current operating figure).¹³ Several operating figures come from secondary summaries of company releases.

On these bases, CoreWeave has delivered about 41% of its contracted power. Galaxy and Cipher are at roughly a quarter, TeraWulf at 12%, Applied Digital at 7% and IREN under 1% (all derived). The physical score tells you who holds the scarce input. The gap illustrates delivery exposure on each company's own basis; the percentages are not comparable completion rates. For companies on the right of the map, construction timing is a valuation input.

Valuation should follow the model and the position

Each business model calls for a different primary method, and the map adds a second question: how much of the stack the buyer is paying for. We do not quote trading multiples in this note; a sector average would blend the models this note separates.

EXHIBIT 7 Valuation approach by business model

BUSINESS MODEL	COMPANIES IN THIS NOTE	PRIMARY METHOD	INPUT THAT MATTERS MOST	MAIN WAY IT GOES WRONG
Data-center landlord	Applied Digital, Galaxy, TeraWulf, Cipher, Hut 8	DCF of contracted lease cash flows; EV per contracted MW as a check	Tenant credit; cost to complete	The tenant is itself a leveraged GPU owner
Hosting provider	Fluidstack (private)	DCF of service fees net of pass-through costs	Occupancy; cost pass-through	Service obligations without asset backing
GPU infrastructure	IREN, Lambda (private)	Contract-by-contract DCF with an explicit residual value for the GPUs; sum-of-parts for undelivered power	GPU life and residual pricing; delivery dates	Pipeline MW valued as if delivered
AI cloud platform	CoreWeave, Nebius, Crusoe, Together AI, Vultr	DCF on backlog with leases treated as debt-like; EV/EBITDA supplemented by replacement-capex and residual-value analysis	GPU life; customer concentration; price per MW	EBITDA excludes depreciation and can obscure replacement costs

Analytical framework: Dilikoglu Research; not a sourced industry standard.

What would change our view?

- **Older GPUs keep earning.** If previous-generation fleets re-contract at reasonable prices as the first 2023–24 contracts roll off, six-year lives are justified and the GPU-owning models look much better than Exhibit 4. This is the most important signal to watch.
- **The revenue-per-MW gap closes.** If IREN-style contracts reprice toward the \$20–25m per MW Nebius reports, hardware-life risk falls sharply (Exhibit 5). If the gap reflects different MW definitions, the higher figures overstate it.
- **Power stops being scarce.** Faster grid build-outs and on-site generation would weaken the premium for controlling power. Landlords on long leases are protected for the term, not beyond it.
- **The scores are judgement calls.** Moving one cell by 0.5 moves a company half a unit. Lambda and Fluidstack sit on the dividing line, and 26 of 104 cells rest on secondary sources or incomplete disclosure.
- **Platform revenue becomes measurable.** None of the listed companies reports platform or inference revenue separately. If they start to, the vertical axis can be weighted by revenue instead of capability.

Sources and evidence notes

Company statements are attributed disclosures, not independent verification. Sources marked "secondary" are trade press or third-party summaries and carry less weight. Figures labelled derived show their formula in the exhibit notes.

1. IREN Ltd, Form 10-K, fiscal year ended 30 Jun 2026
2. IREN, "IREN Completes Acquisition of Mirantis", 4 Aug 2026
3. IREN, "IREN Secures \$9.7bn AI Cloud Contract with Microsoft", 3 Nov 2025

4. IREN, FY26 results release (8-K exhibit), 27 Aug 2026
5. CoreWeave, Q2 2026 results release, 11 Aug 2026
6. CoreWeave, Form 10-K, FY2025, filed 2 Mar 2026
7. CoreWeave, Form 10-Q, Q1 2026
8. DatacenterDynamics, CoreWeave–Core Scientific deal ended (secondary)
9. Commercial Observer, CoreWeave Kenilworth purchase, 5 Aug 2025 (secondary)
10. DatacenterDynamics, CoreWeave Lancaster plans, 15 Jul 2025 (secondary)
11. CoreWeave docs, Serverless Inference
12. DatacenterDynamics, Nebius Kansas City campus, 12 Jan 2026 (secondary)
13. Blockspace, Nebius Q2 2026, 12 Aug 2026 (secondary)
14. Nebius, Token Factory launch, 5 Nov 2025
15. Nebius/City of Vineland, Microsoft \$17.4bn agreement, 8 Sep 2025
16. Nebius, Q3 2025 shareholder letter (6-K), Meta agreement
17. MarketBeat, Nebius Q2 2026 call highlights, 12 Aug 2026 (secondary)
18. Cleanview, Stargate Abilene behind-the-meter report, c. Apr 2026 (secondary)
19. Futuriom, Crusoe: cloud company or construction kingpin, Jul 2025 (secondary)
20. Crusoe, Managed Inference launch, 20 Nov 2025
21. Bloomberg Government, Crusoe raises over \$3bn at \$30bn valuation, 3 Sep 2026 (secondary)
22. Lambda community forum, sunsetting Chat and Inference, Sep 2025
23. RuntimeWire, Lambda pre-IPO round, 24 Aug 2026 (secondary)
24. Missouri Partnership, Lambda Kansas City AI factory, 28 Oct 2025
25. TechCrunch, Together AI raises \$800m at \$8.3bn, 1 Jul 2026 (secondary)
26. Sacra, Together AI company profile, accessed Oct 2026 (secondary)
27. Data Center Knowledge, Vultr Ohio cluster, 2 Dec 2025 (secondary)
28. Forbes, Fluidstack valued at \$18bn, 3 Sep 2026 (secondary)
29. Applied Digital, FY2026 Q4 results release (8-K), 27 Jul 2026
30. Applied Digital, Form 10-K, fiscal year ended 31 May 2026
31. Galaxy, ERCOT interconnection approval, 15 Jan 2026
32. The Energy Mag, Galaxy Q2 2026 and 526 MW CoreWeave lease, 5 Aug 2026 (secondary)
33. TeraWulf, Form 10-K, FY2025
34. TeraWulf, Q2 2026 results release (8-K), 5 Aug 2026
35. Cipher, Q4/FY2025 results release (8-K), 24 Feb 2026
36. Cipher Digital, Barber Lake lease amendment (8-K), 25 Sep 2026
37. DatacenterDynamics, AWS 300 MW lease at Cipher, 3 Nov 2025 (secondary)
38. Mitrade, TeraWulf and Galaxy operating capacity, 5 Aug 2026 (secondary)
39. Trefis, Cipher Digital capacity, 10 Aug 2026 (secondary)
40. IREN, Q3 FY25 results release, 14 May 2025

41. Blockspace, B. Riley on Applied Digital Polaris Forge, 21 Aug 2026 (secondary)
42. Wing VC, How much does an NVIDIA NVL72 cost, Aug 2026 (secondary)
43. CoreWeave FY2025 10-K property and equipment note, via edgar.tools
44. SiliconANGLE, GPU depreciation and useful-life assumptions, 22 Nov 2025 (secondary)
45. Hudson Labs, server useful lives at Amazon and Meta (secondary)
46. Global Data Center Hub, How long does an AI chip last, 11 Sep 2026 (secondary)
47. Equinix, FY2025 10-K, property and equipment useful lives
48. Nebius docs, services overview
49. Lambda, product overview (llms.txt)
50. IREN, "IREN Announces Acquisition of Mirantis", 5 May 2026
51. Crusoe, AutoClusters and platform services announcement, 13 Mar 2025
52. Fluidstack docs, getting started
53. Hudson Labs, Microsoft server useful life, 30 Jul 2026 (secondary)
54. Hut 8, second 352 MW Beacon Point IT lease; total contracted IT capacity 949 MW, 20 July 2026 (company announcement)
55. DatacenterDynamics, Hut 8 \$3.25bn bonds for River Bend, 29 Apr 2026 (secondary)
56. Dilikoglu Research, From Megawatts to Equity Value: Who Captures the AI Infrastructure Buildout?, 2 Oct 2026

Research disclosure

- Tugay Dilikoğlu holds no securities in the companies discussed, has no material interests in them and received no issuer or third-party compensation for this commentary.
- Produced independently by Tugay Dilikoğlu, separately from his professional responsibilities. Views are the author's own and do not represent an employer or affiliated institution.
- For informational purposes only; not personalized investment advice or an offer to transact.
- Opinions reflect Dilikoglu Research's assessment at the stated cutoff and may change. This sector commentary does not initiate or change any company rating or target; refer to the separately published company reports for their dated recommendations, assumptions and report-specific conflict disclosures.
- Public disclosures do not provide complete project-level economics; no assurance of future financing, performance or investment returns is implied.