

EXECUTIVE REFERENCE REPORT · Q2 2026

AI Infrastructure Power Rankings

A ranking of the companies shaping the AI economy through compute, energy, distribution, capital, R&D, adoption, and strategic dealmaking.

The AI economy is no longer just a model race. It is a contest over chips, power, land, networks, capital, and distribution.

Audience	Edition	Format
Policymakers, CEOs, operators, VCs, founders	Q2 2026	Print-ready executive HTML

Executive summary

The AI infrastructure race is now a stack contest. The winners are not simply the companies with the best models. They are the companies that control scarce accelerators, advanced packaging, memory, networking, cloud platforms, data center real estate, power procurement, and enterprise distribution.

NVIDIA ranks first because it controls the most valuable near-term choke point: the accelerator, networking, and software bundle around which the AI economy still plans capacity. Microsoft, Alphabet, and Amazon follow because they combine cloud reach, custom silicon, global infrastructure, massive balance sheets, and enterprise distribution. Meta and OpenAI sit just behind because they are among the strongest demand creators in the market.

The insurgent infrastructure tier matters. Oracle has become a serious AI-capacity broker. CoreWeave has become the defining AI-native cloud. xAI is assembling frontier-scale compute with unusual speed. Crusoe shows the new energy-first logic: move compute to power, not merely power to compute.

The practical takeaway is blunt: AI infrastructure power now comes from portfolio construction. The winners combine at least four of these assets at once: compute, power, distribution, capital, R&D, and adoption feedback loops.

Contents

- | | |
|--------------------------|-----------------------------|
| 1. Methodology | 5. Top-tier profiles |
| 2. Top ten comparison | 6. Strategic spotlights |
| 3. Full ranking table | 7. Implications by audience |
| 4. What the rankings say | |

Methodology

This ranking scores companies on seven dimensions, each on a 0–10 scale, then converts the weighted result to a 0–100 overall score. It measures power over the AI economy, not stock-market performance, not brand awareness, and not model quality alone.

Dimension	Weight	What it captures
Compute	24%	Accelerators, CPUs, custom silicon, packaging, networking, owned or contracted compute, data center footprint.
Energy	10%	Power availability, efficiency, renewable or carbon-free sourcing, water/carbon intensity, energy strategy.
Distribution	16%	Cloud/platform reach, enterprise installed base, APIs, developer leverage, ecosystem lock-in.
Capital	15%	Revenue scale, profitability, cash flow, funding access, backlog and contract depth.
R&D	15%	Labs, model development, patents, publications, talent density, research velocity.
Adoption	15%	Customer base, production deployment, sector breadth, deployment scale.
M&A	5%	Recent strategic acquisitions, JVs, inorganic expansion that changes stack control.

A score of 10 means a global control point. A score of 8 means strong and hard to displace. A score of 6 means important but substitutable. A score below 5 means niche, regional, or still emerging. The ranking is intentionally debatable. That is part of its value.

Top ten comparison

Rank	Company	Compute	Energy	Distribution	Capital	R&D	Adoption	M&A	Overall
1	NVIDIA	10.0	7.8	9.6	10.0	9.5	9.7	8.5	95.2
2	Microsoft	9.5	9.0	10.0	10.0	9.2	9.8	6.8	94.7
3	Alphabet	9.6	8.8	9.6	10.0	9.8	9.0	6.8	93.8
4	Amazon	9.5	8.7	9.8	10.0	8.8	9.5	6.5	92.9
5	Meta	9.3	8.0	9.2	9.9	9.5	9.2	7.5	91.7
6	OpenAI	9.0	7.0	9.5	8.8	10.0	9.7	6.0	89.6
7	TSMC	9.4	7.5	7.8	9.8	8.8	9.2	5.5	87.0
8	Broadcom	8.7	7.3	8.4	9.6	8.2	8.7	9.0	85.9
9	AMD	8.8	7.8	8.0	9.2	8.7	8.3	8.8	85.4

Rank	Company	Compute	Energy	Distribution	Capital	R&D	Adoption	M&A	Overall
10	Oracle	9.0	7.8	8.4	9.2	7.8	8.7	8.0	85.4

Ranked table of ninety-nine companies

The original research brief requested 100. This version contains 99 scored companies because the source ranking had an empty placeholder at rank 100. That should be fixed before public release.

Rank	Company	Compute	Energy	Distribution	Capital	R&D	Adoption	M&A	Overall
1	NVIDIA	10.0	7.8	9.6	10.0	9.5	9.7	8.5	95.2
2	Microsoft	9.5	9.0	10.0	10.0	9.2	9.8	6.8	94.7
3	Alphabet	9.6	8.8	9.6	10.0	9.8	9.0	6.8	93.8
4	Amazon	9.5	8.7	9.8	10.0	8.8	9.5	6.5	92.9
5	Meta	9.3	8.0	9.2	9.9	9.5	9.2	7.5	91.7
6	OpenAI	9.0	7.0	9.5	8.8	10.0	9.7	6.0	89.6
7	TSMC	9.4	7.5	7.8	9.8	8.8	9.2	5.5	87.0
8	Broadcom	8.7	7.3	8.4	9.6	8.2	8.7	9.0	85.9
9	AMD	8.8	7.8	8.0	9.2	8.7	8.3	8.8	85.4
10	Oracle	9.0	7.8	8.4	9.2	7.8	8.7	8.0	85.4
11	Samsung Electronics	8.6	7.8	8.2	9.7	8.3	8.8	5.5	84.5
12	Apple	7.2	8.5	8.7	9.8	8.4	8.8	5.0	82.7
13	Huawei	8.5	7.2	8.0	8.4	8.8	8.6	6.5	82.3
14	Alibaba Cloud	8.3	7.0	8.4	9.1	8.0	8.7	6.0	82.1
15	ASML	8.8	7.0	7.0	9.0	8.8	8.9	5.0	81.9
16	Anthropic	8.2	7.0	8.0	8.2	9.4	8.5	5.5	81.4
17	CoreWeave	9.2	7.2	7.8	8.1	7.5	8.4	7.0	81.3
18	Arm	7.8	8.0	8.7	8.5	8.5	8.0	5.5	80.9
19	xAI	9.1	6.2	7.8	8.5	8.8	7.6	6.0	80.9
20	Equinix	7.6	9.0	8.5	8.8	7.0	8.7	6.0	80.6
21	SK hynix	8.4	7.5	7.5	8.8	8.0	8.8	5.0	80.6
22	IBM	7.4	7.8	8.0	9.1	8.5	8.0	6.5	80.0
23	Cisco	7.5	7.3	8.4	9.3	7.6	8.0	7.5	79.8
24	Dell Technologies	7.8	7.0	8.5	8.8	7.5	8.5	6.5	79.8

Rank	Company	Compute	Energy	Distribution	Capital	R&D	Adoption	M&A	Overall
25	Applied Materials	8.3	7.2	7.0	8.8	8.2	8.6	5.5	79.5
26	Tencent Cloud	7.6	7.1	8.2	8.9	7.8	8.2	5.5	78.6
27	Qualcomm	7.2	8.0	8.5	8.6	8.0	8.0	5.5	78.5
28	HPE	7.7	7.2	8.1	8.0	7.4	8.1	8.8	78.3
29	Digital Realty	7.4	8.7	8.1	8.5	6.8	8.2	6.5	77.9
30	Micron	8.0	7.6	7.2	8.5	7.8	8.4	5.0	77.9
31	Arista Networks	7.7	7.5	7.8	8.6	7.5	8.3	5.5	77.8
32	Schneider Electric	6.8	9.0	7.6	9.0	7.2	8.4	6.5	77.6
33	Marvell	8.0	7.3	7.6	8.2	7.8	8.0	5.8	77.6
34	Lam Research	8.1	7.1	6.8	8.5	8.0	8.2	5.2	77.1
35	Intel	7.6	7.0	8.0	8.2	8.2	7.5	6.0	76.9
36	Baidu AI Cloud	7.7	7.0	7.8	8.2	8.0	7.8	5.5	76.7
37	Databricks	6.6	7.0	8.0	8.4	8.0	8.6	6.2	76.2
38	Vertiv	7.0	8.0	7.8	8.4	6.8	8.5	6.0	75.8
39	NTT DATA	7.1	7.8	7.8	8.6	7.2	7.8	5.8	75.6
40	KLA	7.8	7.0	6.6	8.5	8.0	8.0	5.1	75.6
41	Supermicro	8.0	6.8	7.5	7.6	6.9	8.4	5.5	75.1
42	Snowflake	6.4	7.2	7.8	8.4	7.6	8.4	5.8	74.5
43	G42	7.5	7.2	7.4	8.0	7.0	7.8	6.5	74.5
44	Crusoe	8.0	9.0	6.8	7.0	6.8	7.5	6.8	74.4
45	Palantir	6.2	7.0	7.6	8.7	7.8	8.2	6.0	74.1
46	Lenovo	7.2	7.0	8.0	8.0	6.8	8.0	5.5	74.0
47	Quanta Computer	7.6	7.0	7.1	8.0	6.8	8.2	5.2	73.7
48	MediaTek	6.5	8.0	8.0	8.0	7.4	7.6	4.8	73.3
49	GlobalFoundries	7.4	7.0	6.8	8.0	7.4	7.8	5.3	73.1
50	Foxconn	7.0	7.0	7.2	8.4	6.5	8.0	5.5	72.4
51	Cerebras	7.8	7.4	6.0	6.5	7.8	6.4	5.0	72.2
52	AirTrunk	6.8	8.4	6.5	7.6	5.4	7.6	5.2	71.9
53	Wistron	7.0	6.8	6.8	7.7	6.3	7.8	5.0	71.4
54	Inventec	6.9	6.8	6.8	7.5	6.3	7.7	5.0	70.9

Rank	Company	Compute	Energy	Distribution	Capital	R&D	Adoption	M&A	Overall
55	Celestica	6.8	7.0	6.8	7.8	6.4	7.6	5.0	70.9
56	QTS Realty Trust	7.0	8.2	6.9	8.0	5.8	7.8	5.8	70.8
57	NextEra Energy	5.6	9.6	6.5	9.2	5.8	7.6	5.7	70.8
58	Groq	7.6	8.4	6.2	6.8	7.2	6.8	5.2	70.8
59	Sanmina	6.7	7.0	6.6	7.6	6.2	7.4	6.0	70.0
60	Applied Digital	7.4	8.0	6.2	6.9	5.8	7.6	6.5	69.9
61	Vantage Data Centers	7.2	8.0	6.8	7.8	5.8	7.8	5.5	69.9
62	Keppel	6.6	8.0	6.4	8.1	5.8	7.2	5.5	69.2
63	Constellation Energy	5.8	8.8	6.2	8.7	5.6	7.4	5.8	69.1
64	Goodman Group	6.6	7.8	6.3	8.3	5.5	7.3	5.2	68.9
65	GE Vernova	6.2	8.2	6.4	8.5	7.0	7.5	6.0	68.8
66	Together AI	7.0	7.0	6.6	6.4	7.0	7.2	5.0	68.7
67	Siemens Energy	6.1	8.0	6.2	8.0	6.8	7.3	5.5	68.1
68	GDS Holdings	7.0	7.8	6.6	7.0	5.8	7.6	5.5	68.0
69	Ciena	6.0	7.4	6.8	7.8	6.6	7.2	5.0	67.8
70	ABB	5.9	8.3	6.7	8.3	6.6	7.6	5.3	67.7
71	Iron Mountain Data Centers	6.8	8.0	6.8	7.8	5.6	7.4	5.5	67.7
72	Eaton	6.0	8.3	6.8	8.4	6.7	7.8	5.5	67.7
73	Chindata Group	6.9	7.8	6.5	7.2	5.6	7.4	5.3	67.4
74	Stack Infrastructure	7.0	7.8	6.6	7.4	5.5	7.4	5.3	67.4
75	Fujitsu	6.0	7.6	7.0	7.6	7.0	7.0	5.0	67.4
76	SoftBank	5.8	6.8	7.4	8.8	6.2	7.4	7.0	67.4
77	UMC	6.8	7.0	6.5	7.8	6.8	7.5	4.8	67.4
78	ASE Technology	6.9	7.0	6.6	7.7	6.8	7.7	4.8	67.4
79	CyrusOne	6.9	8.0	6.7	7.7	5.6	7.5	5.5	67.4
80	Nebius	7.3	7.2	6.4	6.5	6.6	6.8	5.0	67.3
81	NEXTDC	6.7	8.5	6.2	7.4	5.4	7.4	5.0	67.3
82	Compass Datacenters	6.9	7.7	6.4	7.3	5.3	7.3	5.2	67.0
83	Bloom Energy	6.0	8.4	6.0	6.9	6.8	7.0	5.6	66.2
84	Trane Technologies	5.6	8.0	6.2	8.1	6.2	7.2	5.2	66.2

Rank	Company	Compute	Energy	Distribution	Capital	R&D	Adoption	M&A	Overall
85	Corning	5.8	7.2	6.6	8.0	6.2	7.2	4.8	66.2
86	Quanta Services	5.7	7.0	6.7	8.0	5.9	7.3	5.8	66.1
87	Cummins	5.8	7.2	6.2	8.1	6.4	7.0	5.6	66.1
88	Switch	6.6	8.2	6.4	7.2	5.2	7.0	4.8	65.8
89	SambaNova	7.2	7.4	5.8	6.3	7.0	6.2	4.7	65.6
90	Coherent	6.2	7.0	6.2	7.5	6.4	7.0	4.8	65.5
91	Carrier Global	5.5	7.8	6.2	8.0	6.1	7.0	5.0	65.1
92	Johnson Controls	5.4	7.8	6.4	8.0	6.0	7.0	5.0	65.0
93	Credo Technology	6.3	7.2	6.0	7.0	6.4	6.8	4.8	64.6
94	Legrand	5.5	7.8	6.2	7.8	6.0	7.0	5.0	64.6
95	Astera Labs	6.2	7.0	5.8	7.0	6.4	6.8	4.8	63.9
96	Modine	5.6	7.6	5.8	7.2	5.8	6.8	4.8	62.4
97	Lumentum	5.8	7.0	5.8	7.0	6.0	6.8	4.7	62.2
98	NRG Energy	5.2	7.0	5.9	7.8	5.2	6.8	5.0	61.1
99	Vistra	5.2	6.8	5.8	7.8	5.1	6.7	5.0	60.5

What the rankings say

PATTERN 1

Concentration with selective leakage

NVIDIA remains central, but hyperscalers are pushing custom silicon. The market is not de-NVIDIA-ing. It is becoming more multi-architecture, especially for inference.

PATTERN 2

Distribution beats raw brilliance

Models and chips need procurement channels, identity, security, compliance, APIs, and developer tooling. The best standalone technology can still lose without distribution.

PATTERN 3

Energy is the governing variable

The frontier has moved into grid strategy, generation, interconnection, land, water, and cooling. AI capacity now depends on power strategy as much as chip supply.

PATTERN 4

Upstream bottlenecks are underpriced

TSMC, HBM suppliers, equipment vendors, networking firms, cooling companies, and power-equipment players quietly govern how much AI can physically exist.

Top-tier profiles

NVIDIA

NVIDIA is the central control point of the AI infrastructure economy because it sells the scarce asset every serious AI player needs: high-performance accelerated compute that works at scale. Its advantage is no longer just the GPU. Blackwell, NVLink, Spectrum-X, CUDA, enterprise software, reference systems, and ecosystem gravity together make NVIDIA an AI-factory architecture company. It does not merely sell components into the stack; it defines the stack's upgrade rhythm. The mental model is simple: NVIDIA is the operating system of the physical AI economy. Its biggest risk is not immediate technical collapse, but customer diversification as hyperscalers push custom silicon for inference economics. Even there, challengers must prove they can match the whole system, not only the chip.

Microsoft

Microsoft is the strongest distribution engine in the ranking. Azure gives it compute reach, OpenAI gives it model-market pull, GitHub gives it developer gravity, and Microsoft 365 gives it enterprise workflow access. The Maia silicon effort matters because it signals that Microsoft does not want to remain purely dependent on external accelerators. Its moat is procurement, identity, security, compliance, and existing enterprise budget lines. The mental model is AI as a feature of the enterprise operating system. Microsoft does not need to win every benchmark; it needs to make AI adoption feel safe, bundled, and inevitable.

Alphabet

Alphabet is the most technically complete AI institution outside NVIDIA. It combines DeepMind research, custom TPU hardware, global cloud infrastructure, search distribution, ads cash flow, Android, YouTube, and energy-efficient data center operations. TPU strategy gives Alphabet a serious non-NVIDIA compute path, especially as training and inference architectures diverge. Its weakness has historically been commercial packaging and enterprise conversion compared with Microsoft and AWS. The mental model is the vertically integrated AI lab-cloud-advertising machine. If custom silicon economics matter more in the next cycle, Alphabet's position strengthens.

Amazon

Amazon remains the broadest cloud landlord in the AI economy. AWS has immense distribution, customer trust, developer reach, and balance-sheet power. Trainium is strategically important because it gives Amazon a path to control inference and training cost curves without depending fully on NVIDIA. Bedrock, SageMaker, global regions, enterprise accounts, and public-sector contracts make Amazon a default infrastructure provider even when it lacks the cultural AI dominance of OpenAI or Google. The mental model is AI tenancy at planetary scale: Amazon sells the rails, the rooms, and increasingly the custom engines inside them.

Meta

Meta is one of the largest demand creators in AI infrastructure. Its capex appetite, social distribution, ad systems, messaging platforms, open-weight strategy, and custom MTIA roadmap make it a structurally important force. Meta's difference is that it can turn inference improvements directly into ads, feeds, agents, recommendation systems, and consumer engagement. It is less enterprise-native than Microsoft or AWS, but it has enormous product-surface leverage. The mental model is consumer-scale inference as an industrial weapon. Meta treats infrastructure as the product-development bottleneck itself.

OpenAI

OpenAI is the clearest example of demand pull becoming infrastructure power. It does not fully own the physical substrate, yet it reshapes the behavior of clouds, chip vendors, capital providers, and governments because demand for its models and products is so intense. Stargate-style capacity planning moves OpenAI from lab to infrastructure orchestrator. Its weakness is dependency on partners for compute, capital, and energy buildout. Its strength is that enterprises, developers, and policymakers increasingly plan around OpenAI availability. The mental model is model demand as a gravitational field.

TSMC

TSMC is the bottleneck behind the bottleneck. NVIDIA, AMD, Apple, hyperscaler silicon teams, and many AI-chip challengers depend on its advanced manufacturing, packaging, yield discipline, and trusted execution. AI infrastructure is not just GPUs; it is wafers, CoWoS, advanced nodes, testing, supply coordination, and expansion timing. TSMC has less distribution power than hyperscalers but more physical leverage over the future supply of AI compute. The mental model is civilizational fabrication capacity. A frontier lab cannot scale if TSMC and its packaging ecosystem clog.

Broadcom

Broadcom is an infrastructure toll collector. It benefits from AI networking, switching, custom silicon exposure, VMware software control, and enterprise infrastructure consolidation. It is not as visible as NVIDIA or OpenAI, but it is deeply embedded in the systems that make large-scale AI work. VMware gives Broadcom a hybrid-cloud and private-cloud control plane, which matters as enterprises bring AI into existing infrastructure. The mental model is stack compression through silicon and software. Broadcom profits from the AI buildout whether workloads centralize in hyperscale clouds or spread into enterprise estates.

AMD

AMD is the most credible large-scale challenger to NVIDIA, but only because it has begun building a system strategy rather than only a GPU strategy. Its acquisitions around AI talent, rack-scale design, and systems enablement show a move toward full-stack infrastructure credibility. AMD still trails NVIDIA in software ecosystem gravity, but customers want a second source. That desire alone gives AMD strategic power. The mental model is insurance against one-vendor dependency. AMD does not need to beat NVIDIA everywhere; it needs to become the trusted alternative where cost, supply diversity, and portability matter.

Oracle

Oracle has become the unexpected AI infrastructure broker. Its cloud business now matters because it is willing to build, finance, customize, and operate capacity at a scale that large model companies need. Oracle's strength is not developer love; it is operational willingness and enterprise infrastructure discipline. Its role in major AI capacity programs shows that the market rewards companies willing to be unfashionably physical. The mental model is the comeback of the boring asset owner. When demand exceeds supply, a company with land, contracts, cloud infrastructure, and buildout appetite becomes strategically important again.

Strategic spotlights

xAI

xAI is a frontier demand magnet with extreme compute ambition. Its strategic power comes from speed, capital access, and the ability to assemble massive clusters quickly. Its weakness is narrower distribution and more exposed energy positioning compared with hyperscalers. It rises if it turns compute scale into durable enterprise and developer adoption.

CoreWeave

CoreWeave is the defining AI-native cloud specialist. Its advantage is speed to capacity and focus on GPU-heavy workloads. Its weakness is concentration risk, financing intensity, and narrower platform distribution than AWS, Azure, or Google Cloud. It is the most important neocloud because it proved specialist clouds can become strategic infrastructure, not just arbitrage machines.

Crusoe

Crusoe is the cleanest expression of the energy-first AI infrastructure thesis. Instead of treating power as a support function, it treats power access and siting as the center of the product. This matters because AI buildout is increasingly constrained by interconnection, generation, land, and cooling rather than model ideas alone. Crusoe's mental model is compute moving to power.

Micron

Micron is a memory bottleneck company in an AI market that keeps underestimating memory. HBM, server DRAM, and data center storage become more important as models grow and inference scales. Micron's power is upstream and less glamorous, but it directly affects how much useful AI capacity can be deployed.

SK hynix

SK hynix is one of the critical HBM leaders. Its influence comes from supplying the memory layer that allows accelerators to perform at frontier scale. Policymakers and investors often obsess over GPUs while ignoring the memory architecture required to make those GPUs useful.

Schneider Electric

Schneider Electric matters because AI infrastructure has become an electrical engineering problem. Power distribution, cooling, grid interfaces, and data center energy systems are no longer back-office concerns. Schneider's role is to translate compute ambition into deployable facilities.

Strategic implications

FOR POLICYMAKERS

AI infrastructure is now industrial policy, grid policy, water policy, competition policy, and national-security policy at the same time. Sovereign AI requires compute availability, effectiveness, and resilience, not slogans.

FOR VCS

The durable opportunities are not lazy model wrappers. They are in interconnect, power, cooling, inference cost accounting, orchestration, memory, packaging, and regional capacity.

FOR FOUNDERS

Build for portability and margin control. Hard-wiring to one cloud, one model, or one accelerator without a cost reason is a strategic mistake.

FOR EXECUTIVES

AI infrastructure strategy is portfolio management: training versus inference, owned versus leased capacity, GPU versus custom silicon, and renewable versus dispatchable power.

Editorial note: This executive HTML version is formatted for printing to PDF. It is designed as a shareable reference, not a final audited investment product. Before publishing, add a proper source appendix and fill the missing rank 100.