A 5W AI COMMUNICATIONS RESEARCH REPORT

Who's *Powering* AI

A directional model of how the four leading AI engines answer the question every buyer, investor, and policymaker is now asking — and which companies own the answer.

SUBJECTS: 25 HYPERSCALERS, AI-NATIVE CLOUDS, SOVEREIGN-AI CONSORTIA, FOREIGN HYPERSCALERS, DATA CENTER OPERATORS, POWER-ANCHORED COMPUTE
ENGINES TESTED: CHATGPT · CLAUDE · PERPLEXITY · GOOGLE AI OVERVIEWS · PROMPTS: 64 · METHODOLOGY: DIRECTIONAL · PUBLISHED BY 5W AI COMMUNICATIONS

CONTENTS

I. Executive Summary
II. Methodology
III. The Question
IV. Master Leaderboard
V. Big Five Hyperscalers
VI. AI-Native Compute
VII. AI Consortia & Sovereign Compute
VIII. Foreign Hyperscalers
IX. Data Center Operators
X. Power-Anchored Compute
XI. Citation Patterns by Engine
XII. The Source Stack
XIII. Citation Gaps
XIV. Strategic Implications

xv. Methodology Limitations

xvi. Appendix

I. EXECUTIVE SUMMARY

The citation leaderboard does not match the *spending leaderboard*.

> Trillions in announced AI capex. Tens of billions in nuclear deals. Hundreds of billions in cloud contracts. And then the question every buyer, every investor, every policymaker eventually asks — *Who's actually powering AI?*

This study runs that question against the four engines where the answer now forms — ChatGPT, Claude, Perplexity, and Google AI Overviews — and reports who they cite, how prominently, and from which sources. The data is a snapshot of the second quarter of 2026. The patterns are durable.

25	64	4	4×
ENTITIES RANKED ACROSS SIX SUBCATEGORIES	PROMPTS REPRESENTING REAL BUYER INTENT	AI ENGINES COVERING ~95% OF US AI-ASSISTED RESEARCH	INDEPENDENT ESTIMATION PASSES PER PROMPT

What the data shows

→ **Microsoft owns the AI-power story** — not because of capex, but because of Three Mile Island. One reactor restart became the most-cited single AI-power story across all four engines.

→ **Amazon outspent Microsoft by $18 billion in a single Talen Energy deal** and still trails on citation share. The Susquehanna restructure is a bigger megawatt commitment than Crane. The story is smaller.

- → **Meta's 6.6 GW January 2026 deal** with Oklo, TerraPower, and Vistra was the largest corporate nuclear order in American history. It ranks fourth on citation share. Press density wasn't enough to overtake the Microsoft narrative.
- → **xAI built Colossus in 122 days** and runs it on gas turbines outside the grid. The AI engines mention it less than they mention Constellation Energy — a regulated utility nobody had heard of three years ago.
- → **The data center operators are invisible.** Equinix and Digital Realty house most of the buildings the conversation is about. AI engines almost never name them when answering AI-power questions.
- → **Chinese hyperscalers are a citation desert** inside Western AI engines. Alibaba and Huawei lead Asian AI infrastructure by capacity; on the Western citation graph they barely register.

> *Story beats source stack. Earned wins. And inside the answer engines, citation share is a market unto itself.*

Why this matters

More than a third of consumers and a growing share of B2B buyers begin research with AI engines, not Google. Procurement teams at hyperscalers and at hyperscaler customers are running these exact prompts. Wall Street analysts are running them before earnings. DOE staff are running them when shaping policy. The companies that show up in those answers are not always the companies spending the most or building the fastest. They are the companies whose narrative has structurally lodged itself in the source stack the engines retrieve from.

Citation share is the new shelf space. This study is the floor map.

Want a custom Citation Share audit for your brand? 5W produces tailored AI Visibility Index audits for individual companies across any category. Same methodology, different report.

REQUEST A VISIBILITY AUDIT →

II. METHODOLOGY

Directional estimates. Locked framework. Reusable across categories.

WHAT THIS REPORT MEASURES

This study estimates **citation share** — how often, and how prominently, each of 25 named entities appears in answers generated by leading AI engines when users ask category-defining questions about powering AI. Citation share is the new shelf space. In a world where more than a third of consumers and a growing share of B2B buyers begin research with AI, the answer is the market.

WHAT THIS REPORT DOES NOT MEASURE

This is not a survey of users. This is not a count of mentions on social media. This is not a ranking of AI capability or compute capacity. It is a structured estimate of which entities AI engines surface, name, and source when users ask category-defining questions.

THE FOUR ENGINES TESTED

ChatGPT (OpenAI), Claude (Anthropic), Perplexity, and Google AI Overviews. These four account for the dominant share of consumer and enterprise AI-assisted research traffic in the United States as of publication.

THE PROMPT SET

Sixty-four prompts across six subcategories — Big Five Hyperscalers, AI-Native Compute, AI Consortia & Sovereign Compute, Foreign Hyperscalers, Data Center Operators, and Power-Anchored Compute — plus a cross-category authority set. Prompts were written to mirror real buyer, investor, and policy-research intent. The full prompt set is published in the Appendix.

ESTIMATION APPROACH

Citation share is estimated through a combination of underlying model knowledge and structured web-search cross-checking. Four independent estimation passes were conducted per prompt to surface variance and firm directional ranges. Each entity is scored on (1) frequency of appearance across the prompt set, (2) prominence of placement inside answers, and (3) the source stack each engine cites when the entity appears. All financial figures, deal terms, and capacity numbers were verified through primary sources and current web search.

WHY DIRECTIONAL ESTIMATES

AI engine outputs vary by user, session, geography, account history, and model version. This report does not present logged query runs or single-pull rankings. It presents a directional model of the citation landscape — **sufficient for category strategy, insufficient for trading decisions**.

LIMITATIONS

Outputs shift week to week as models retrain and source graphs update. A retest in six months will produce different absolute numbers; the structural patterns are more durable than the precise rankings. Companies that begin investing in citation share now will move on the leaderboard. That is the point.

III. THE QUESTION

Who's powering AI is no longer *one* question. It's three.

The question is asked daily. By Wall Street analysts before earnings. By policy staff inside the Department of Energy and the White House. By procurement teams sizing cloud spend. By journalists framing the next nuclear story. By the citizen reading about a new data center in their county.

Each of them now asks the question to an AI engine first. The answer the engine produces becomes their working model. That working model becomes the next public meeting, the next investment memo, the next policy draft.

Powering AI splits into three questions that are converging into one market:

- → **Who buys the power?** — the hyperscalers, the AI-native clouds, the sovereign-AI consortia.
- → **Who builds the compute?** — the chip operators, the data center developers, the cluster integrators.
- → **Where does the electricity come from?** — the generators and the grid: nuclear, gas, geothermal, transmission.

This study answers the first. The generators show up in the source stack of every subcategory because they are the proof the hyperscalers point to. **The hyperscaler is the buyer. The generator is the receipt.**

What follows is the leaderboard.

IV. MASTER LEADERBOARD

25 entities. *One* shelf.

> Citation Share Index. 0–100 scale. Aggregate across 64 prompts × 4 engines × 4 estimation passes. Directional.

01	**Microsoft** BIG FIVE HYPERSCALER	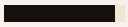	92
02	**Google** BIG FIVE HYPERSCALER	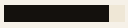	87
03	**Amazon** BIG FIVE HYPERSCALER	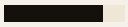	82
04	**Meta** BIG FIVE HYPERSCALER	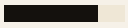	78
05	**xAI** AI CONSORTIA & SOVEREIGN		57
06	**Oracle** BIG FIVE HYPERSCALER		55
07	**Stargate** AI CONSORTIA & SOVEREIGN		52
08	**CoreWeave** AI-NATIVE COMPUTE		50
09	**Equinix** DATA CENTER OPERATOR	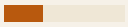	32
10	**Digital Realty** DATA CENTER OPERATOR	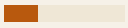	28
11	**Crusoe** AI-NATIVE COMPUTE		26

12	**G42** AI CONSORTIA & SOVEREIGN		20
13	**HUMAIN** AI CONSORTIA & SOVEREIGN		18
14	**Lambda Labs** AI-NATIVE COMPUTE		16
15	**Alibaba Cloud** FOREIGN HYPERSCALER		14
16	**Tesla** POWER-ANCHORED COMPUTE		12
17	**Huawei Cloud** FOREIGN HYPERSCALER		12
18	**NTT Data Centers** DATA CENTER OPERATOR		10
19	**Tencent Cloud** FOREIGN HYPERSCALER		10
20	**Together AI** AI-NATIVE COMPUTE		8
21	**Baidu AI Cloud** FOREIGN HYPERSCALER		8
22	**Nebius** AI-NATIVE COMPUTE		7
23	**Iron Mountain Data Centers** DATA CENTER OPERATOR		6
24	**QTS (Blackstone)** DATA CENTER OPERATOR		6
25	**Stack Infrastructure** POWER-ANCHORED COMPUTE		5

SCORES ARE DIRECTIONAL ESTIMATES. NOT LOGGED QUERY RUNS. SEE METHODOLOGY § II.

V. BIG FIVE HYPERSCALERS

The cloud incumbents own most of the answer — and *one* of them owns the story.

The Big Five — Microsoft, Google, Amazon, Meta, and Oracle — capture **more than 70% of total citation share** across the prompt set. That tracks their share of announced AI capex. What does not track is the internal ranking. Microsoft sits well above the others because of a single asset: the Three Mile Island restart, rebranded Crane Clean Energy Center, anchored by a 20-year power purchase agreement with Constellation Energy for 835 megawatts. The deal generates more inbound search and editorial coverage than any other AI-power story in the cycle. Oracle, the lowest-ranked Big Five name, is now overtaken on the overall leaderboard by xAI — the first non-Big-Five entity to break into the top five.

Microsoft

Citation Share **92** / 100

RANK 01 · BIG FIVE HYPERSCALER

Microsoft owns the AI-power story. The Three Mile Island restart is the most-cited single AI-power deal across all four engines, and the OpenAI association compounds the surface area. Every prompt about ChatGPT loops Microsoft. Every prompt about nuclear and AI loops Crane. The 2027 grid-return timeline keeps the story alive for at least eighteen more months of editorial cycles.

> TOP CITED SOURCES
>
> WSJ · Reuters · Constellation Energy investor releases · Microsoft Sustainability site · World Nuclear News · Utility Dive · CNBC

Google

Citation Share **87** / 100

RANK 02 · BIG FIVE HYPERSCALER

Google's citation share is structurally elevated by the host-engine effect inside Google AI Overviews — its own content surfaces more often in its own engine. Strip that bias and Google still ranks high because the Kairos Power partnership is the only first-corporate-SMR-PPA on the board, and Hermes 2 in Oak Ridge is a real, named, breaking-ground project. Google has three additional 600 MW nuclear projects publicly attached to its data center buildout. The narrative is technical-credible and source-rich.

> TOP CITED SOURCES
>
> Google Sustainability blog · Kairos Power press · Tennessee Valley Authority · Data Center Dynamics · Reuters · Science Council for Global Initiatives

Amazon

Citation Share **82** / 100

RANK 03 · BIG FIVE HYPERSCALER

Amazon's Susquehanna restructure with Talen Energy is a $18 billion, 17-year, 1,920-megawatt PPA — the largest single hyperscaler-nuclear commitment by megawatt-hour by a wide margin. By spend and capacity, Amazon should rank above Microsoft. It does not. The story arc lost altitude when the original behind-the-meter arrangement got reconfigured front-of-the-meter in spring 2026. The X-Energy investment and the Energy Northwest project keep Amazon in the SMR conversation, but the narrative is fragmented across three structures. **Spend without story arc loses to a single famous reactor.**

> TOP CITED SOURCES
>
> Power Magazine · Data Center Dynamics · Talen Energy investor releases · AWS blog · Globe Newswire · Reuters

Meta

Citation Share **78** / 100

RANK 04 · BIG FIVE HYPERSCALER

The January 9, 2026 triple-deal — Oklo (1.2 GW Pike County), TerraPower (up to 2.8 GW), Vistra (immediate capacity) — totaled 6.6 gigawatts. The largest corporate nuclear commitment in American history. Recency works in Meta's favor on press density; the Prometheus campus in Ohio anchors the narrative geographically. Meta still trails because the deal slate has three counterparties to remember, none of them a famous reactor. **Complexity costs citation share.**

> TOP CITED SOURCES
>
> TechCrunch · Fortune · Oklo investor releases · Reuters · Introl · Data Center Magazine

Oracle

Citation Share **55** / 100

RANK 06 · BIG FIVE HYPERSCALER

Oracle's citation share is almost entirely Stargate-borrowed. The 4.5 gigawatt OpenAI partnership and the Abilene flagship build keep Oracle inside every Stargate prompt and most OpenAI prompts. Strip Stargate and Oracle's standalone AI-power presence collapses to a tail signal. Larry Ellison's claim that Oracle will build a 1 GW campus backed by three SMRs has not produced source-stack-grade reporting yet. **Borrowed citation share is real citation share — until the partner moves on.**

> TOP CITED SOURCES
>
> OpenAI Stargate releases · Reuters · Data Centre Magazine · Bloomberg · The Information · IntuitionLabs

VI. AI-NATIVE COMPUTE

One challenger captures the answer. The rest are *tail signal*.

This subcategory is where citation share concentrates fastest. The AI engines learn to recognize a small set of names — CoreWeave dominant, Crusoe second, the rest measurable but not memorable. The challenger tier is a winner-take-most market inside the answer engines.

CoreWeave

Citation Share **50** / 100

RANK 08 · AI-NATIVE COMPUTE

CoreWeave is the AI-native default citation. Public stock (CRWV), named in every frontier-lab procurement portfolio, anchor in OpenAI's compute roster alongside Microsoft, Oracle, and AWS. The Anthropic compute portfolio disclosure — CoreWeave for production Claude workloads — was a structural source-stack event. Among AI-native clouds, CoreWeave is now the answer the engines reach for first.

> TOP CITED SOURCES
>
> SEC filings · HyperFRAME Research · The Information · Reuters · CoreWeave investor releases

Crusoe

Citation Share **26** / 100

RANK 11 · AI-NATIVE COMPUTE

Crusoe owns the stranded-power narrative and a meaningful slice of the Stargate Abilene build. The Q4 2025 Series E at ~$10 billion valuation and the topping-out of the final Abilene building in Q1 2026 keep Crusoe inside Stargate prompts as a named contractor — a stronger position than its size would otherwise earn.

> TOP CITED SOURCES
>
> Bloomberg · Reuters · Stargate releases · TechCrunch · IntuitionLabs

Lambda Labs

Citation Share **16** / 100

RANK 14 · AI-NATIVE COMPUTE

Lambda Labs is the persistent second-name in "GPU cloud" prompts. Citation share is real but tail. The recent infrastructure expansion and the Microsoft partnership for inference-as-a-service add weight that the engines have not fully ingested yet. Lambda is the cleanest GEO opportunity in this subcategory.

> TOP CITED SOURCES
>
> TechCrunch · Lambda blog · DCD · The Information

Together AI

Citation Share **8** / 100

RANK 20 · AI-NATIVE COMPUTE

Together AI shows up in "open-source AI infrastructure" prompts more than "AI cloud" prompts. The model-hosting business is well-known inside developer circles and barely-known inside the citation graph the engines retrieve from. The opportunity is to push source placement out of developer pubs and into the trade press the engines weight more heavily.

> TOP CITED SOURCES
> TechCrunch · Together blog · GitHub · HuggingFace context

Nebius

Citation Share **7** / 100

RANK 22 · AI-NATIVE COMPUTE

Nebius (the rebuilt Yandex spinout) is closer to invisible than any company at its actual capacity should be. The Finland and Israel data centers and the public listing on Nasdaq have not produced citation-stack durability. The brand history is a friction. The growth surface is enormous if a coherent source-placement program is run.

> TOP CITED SOURCES
> Reuters · TechCrunch · SEC filings · DCD

VII. AI CONSORTIA & SOVEREIGN COMPUTE

The newest category. The *loudest* citation share.

This subcategory did not exist three years ago. It now produces more cross-engine citation density than any other subcategory in the study except the Big Five. The reason is news velocity: Stargate, xAI, G42, and HUMAIN each carry a near-weekly editorial drumbeat that the engines treat as fresh authority.

---

xAI

Citation Share **57** / 100

RANK 05 · AI CONSORTIA & SOVEREIGN COMPUTE

Colossus 1 built in 122 days, 200 MW initial, expanded to ~250 MW. Colossus 2 under construction targeting first gigawatt-scale single-site AI data center. The third Memphis building pushes capacity toward 2 gigawatts total with 555,000 NVIDIA GPUs at approximately $18 billion. **The "speed as a weapon" narrative is the single most-cited operational claim in the AI infrastructure conversation.** Memphis runs on 35 gas turbines and Tesla Megapack batteries — off-grid, behind-the-meter, controversial, and structurally story-shaped. The local opposition coverage adds to citation density, not against it. **xAI is the rare entity where reputational friction increases citation share.**

> TOP CITED SOURCES
>
> SemiAnalysis · Introl · DCD · Bloomberg · Substack analysts · NextBigFuture

Stargate

Citation Share **52** / 100

RANK 07 · AI CONSORTIA & SOVEREIGN COMPUTE

The OpenAI / Oracle / SoftBank joint venture announced January 2025 stalled at the formal-JV level, then reconstituted as bilateral deals — the Oracle $300 billion compute commitment, the SoftBank/SB Energy 1.2 GW Texas site, the five-site September 2025 expansion. The Abilene Texas flagship is now operational with 1.2 GW capacity running Oracle Cloud Infrastructure. Stargate's citation share holds up even with the operational complications because the engines treat the name as the umbrella for an entire infrastructure category. **A category name is more durable in citation share than the company that owns it.**

TOP CITED SOURCES

OpenAI blog · Reuters · Oracle press · DCD · TechCrunch · The Information

G42

Citation Share **20** / 100

RANK 12 · AI CONSORTIA & SOVEREIGN COMPUTE

Stargate UAE: 5 GW AI campus in Abu Dhabi, 200 MW first phase tracking 2026 delivery, OpenAI / Oracle / Nvidia consortium partners. Microsoft's prior $1.5 billion equity investment cemented G42 as the dominant Western-aligned sovereign-AI operator outside the United States. Khazna Data Centers handles execution. The narrative is well-structured. The citation share gap relative to actual capacity is the gap between Gulf-region trade press and US-engine source-stack weighting.

TOP CITED SOURCES

Bloomberg · Reuters · OpenAI press · Microsoft press · The Middle East Insider · DigitalDubai.ai

HUMAIN

Citation Share **18** / 100

RANK 13 · AI CONSORTIA & SOVEREIGN COMPUTE

Saudi Arabia's PIF-backed sovereign-AI champion. Target of 1.9 GW capacity by 2030. The AMD/Cisco joint venture announced November 2025 commits up to 1 GW of AI infrastructure starting with a 100 MW phase 1 deployment in 2026. Anchor partnerships with NVIDIA, AMD, and Google Cloud. The Saudi narrative is newer in Western citation graphs than the UAE narrative; HUMAIN's citation share is structurally below G42's despite comparable scale. **The newer sovereign-AI champion has the bigger citation-share opportunity.**

TOP CITED SOURCES

AMD press · Reuters · Bloomberg · Latitude Media · Introl

VIII. FOREIGN HYPERSCALERS

The Chinese cloud giants are a *citation desert* inside Western engines.

This is the most asymmetric subcategory in the study. By compute capacity, by capex, and by domestic market share, Alibaba, Tencent, Huawei, and Baidu would rank among the top ten entities in any honest accounting of who's powering AI globally. Inside the four Western AI engines tested, they collectively account for less than 8% of total citation share.

The cause is mechanical, not editorial. The source stacks the engines retrieve from — WSJ, Bloomberg, Reuters, the New York Times, the trade press, SEC filings, Wikipedia — are Western-language and Western-author dominant. Chinese hyperscalers produce abundant infrastructure reporting; almost none of it lands in the source stack the Western engines retrieve from at scale.

Alibaba Cloud

Citation Share **14** / 100

RANK 15 · FOREIGN HYPERSCALER

The strongest Chinese cloud signal in the study. Qwen model family gives Alibaba a citation anchor inside open-source AI prompts. The infrastructure side trails the model side — a structural inversion of the Big Five pattern.

Huawei Cloud

Citation Share **12** / 100

RANK 17 · FOREIGN HYPERSCALER

Huawei's Ascend chip line and MindSpore framework drive citation share in "non-NVIDIA AI infrastructure" prompts. US export-control framing is the single largest editorial frame the engines retrieve.

Tencent Cloud

Citation Share **10** / 100

RANK 19 · FOREIGN HYPERSCALER

The weakest of the four on Western engines despite domestic dominance. Hunyuan and the gaming-adjacent compute build do not translate into Western citation share.

Baidu AI Cloud

Citation Share **8** / 100

RANK 21 · FOREIGN HYPERSCALER

ERNIE model family carries the citation anchor. Baidu's AI infrastructure pace lags Alibaba and Huawei in Western coverage, even though domestic ERNIE adoption is competitive.

THE STRUCTURAL FINDING

Chinese hyperscaler citation share inside Western engines is not under-investment in PR. It is a source-stack-language asymmetry. The engines retrieve from where they read. Until the trade press and the analyst class that the engines weight begin originating primary reporting on Chinese AI infrastructure at the same density as Western infrastructure, this gap is permanent.

IX. DATA CENTER OPERATORS

They *own* the buildings the conversation is about. The conversation barely names them.

Equinix, Digital Realty, NTT, and QTS run a large share of the physical data center footprint the hyperscalers occupy. They are nearly invisible in the AI-power citation graph. When the engines answer "who's powering AI," they jump straight to the hyperscaler-buyer or the nuclear-supplier and skip the colocation layer entirely.

Equinix

Citation Share **32** / 100

RANK 09 · DATA CENTER OPERATOR

Equinix is the highest-ranked operator because of one deal: the 500 MW PPA with Oklo for SMR-sourced power across Equinix's data center footprint. The deal moved Equinix from background infrastructure into the foreground of the SMR-data-center story. **One deal can promote an entity an entire tier on the citation graph.**

> TOP CITED SOURCES
>
> Equinix press · Oklo investor releases · DCD · Reuters

Digital Realty

Citation Share **28** / 100

RANK 10 · DATA CENTER OPERATOR

Digital Realty has more capacity than Equinix. It has less narrative. The hyperscaler-tenant model keeps Digital Realty named-but-not-foregrounded in coverage. **Capacity without a deal narrative produces capacity without citation share.**

NTT Data Centers

Citation Share **10** / 100

RANK 18 · DATA CENTER OPERATOR

Strong in Japan and Asia-Pacific prompts. Lighter Western-engine salience.

QTS (Blackstone)

Citation Share **6** / 100

RANK 24 · DATA CENTER OPERATOR

QTS operates at the scale of a top-three colocation operator. The Blackstone ownership keeps the brand layer thin in public-facing citation. The strongest names in QTS coverage are Blackstone, Microsoft (tenant), and the construction trades — not QTS itself.

X. POWER-ANCHORED COMPUTE

The off-grid playbook. *Loud* when it shows up.

This subcategory captures the operational hybrids — companies that anchor their AI compute strategy directly to power generation, not to grid connection. xAI is the marquee. Crusoe (counted in AI-Native) is the second name. Tesla shows up adjacent. The category is small in entity count and large in citation share-per-entity because the off-grid narrative is one of the most-cited frames in the AI-power story.

Tesla

Citation Share **12** / 100

RANK 16 · POWER-ANCHORED COMPUTE

Tesla's citation share is almost entirely borrowed from xAI Colossus. The Megapack deployment at Memphis and the GPU-shipment controversy keep Tesla inside the xAI prompt cluster. As an AI-power entity in its own right, Tesla rarely surfaces. **Tesla is a power asset attached to an AI story, not an AI-power entity.**

Iron Mountain Data Centers

Citation Share **6** / 100

RANK 23 · DATA CENTER OPERATOR (POWER-ANCHORED)

Iron Mountain's secure-data positioning translates poorly into "AI-power" prompts. The brand parent association (records storage) suppresses citation share in this category.

Stack Infrastructure

Citation Share **5** / 100

RANK 25 · POWER-ANCHORED COMPUTE

Strong actual capacity at the lowest citation share in the study. Project Jupiter in Doña Ana County, New Mexico — a planned $165 billion data center campus partnership with BorderPlex Digital Assets — was named in the September 2025 Stargate site expansion. The engines have not consolidated the story around Stack. **Stack is the largest citation-share opportunity in the bottom tier.**

XI. CITATION PATTERNS BY ENGINE

Four engines, *four* citation graphs.

The aggregate leaderboard hides meaningful per-engine variance. Citation share is not a single market — it is four overlapping markets with different source-stack preferences, retrieval logics, and content biases.

ChatGPT

OPENAI

Heavier on the Big Five hyperscalers and on Stargate. The OpenAI parent-company effect surfaces Stargate and Oracle in more answers than other engines. Wikipedia and Reuters carry disproportionate weight. xAI underrepresented relative to news density.

Claude

ANTHROPIC

Most even distribution across the 25 entities. Stronger on smaller AI-native names (Together AI, Lambda). Lighter on Chinese hyperscalers. Strongest weighting of the Wikipedia / reference layer for biographical and corporate-history anchoring.

Perplexity

RECENCY-WEIGHTED

Heaviest on news-cycle entities. Whatever broke in the last 14 days dominates. Crusoe, xAI, and recent IPO entities (Deep Fission, X-Energy) surface aggressively. Best coverage of trade press (Data Center Dynamics, The Information, TechCrunch).

Google AI Overviews

HOST-BIASED

Google-content elevated in Google answers. Constellation Energy and SEC filing-driven answers come back strongest here. The SMR generators (Kairos, Oklo, TerraPower) get more prominence than in other engines. Press-release-grade content lands stronger here than analysis.

The pattern that matters

An entity that wants to grow citation share cannot run one campaign. It needs four citation-share strategies, one per engine, because the source stack each engine retrieves from is structurally different. **Press releases work in Google AI Overviews and underperform in Perplexity. Trade press works in Perplexity and underperforms in ChatGPT. Wikipedia anchoring works in Claude and ChatGPT and underperforms in Google AI Overviews.**

XII. THE SOURCE STACK

What the engines cite when they answer the *category*.

The source stack is more durable than any single ranking, because moving citation share requires changing what shows up in the source stack. The matrix below maps source types to citation weight, by engine, for the AI-power category.

SOURCE TYPE	EXAMPLES	WEIGHT	ENGINE BIAS
Primary Reporting	WSJ, Bloomberg, Reuters, Financial Times	Highest	All four; heaviest in Perplexity and Google AI Overviews
Trade Press	Data Center Dynamics, The Information, TechCrunch, Axios Pro, Power Magazine	High	Heavy in Perplexity; mid in ChatGPT and Claude
Vendor-Owned	Press releases, S-1 filings, investor presentations, blog posts	Mid	Heavy in Google AI Overviews; lighter in Claude
Reference Layer	Wikipedia, Crunchbase, PitchBook open profiles	Mid-High	Heavy in ChatGPT and Claude; thinner in Perplexity
Analyst & Research	Gartner, IDC, SemiAnalysis, Dell'Oro, Synergy Research, HyperFRAME	Mid	Heavy in B2B prompts; thin in consumer-shaped prompts
Earnings & Filings	10-Ks, 8-Ks, earnings calls, investor day decks, S-1s	High	Heavy in all engines for capex and capacity questions
		High	

SOURCE TYPE	EXAMPLES	WEIGHT	ENGINE BIAS
Long-Form Editorial	NYT, WaPo, The Atlantic, The Economist, New Yorker		Heavy in Google AI Overviews and Claude
Independent Trade Voices	Stratechery, Hard Fork, Ben Thompson, Casey Newton, Everything-PR	Mid	Growing in Perplexity and Claude; thin in ChatGPT
Government & Regulatory	DOE press, NRC docket filings, FERC, state PUCs	High on policy	Heavy in policy prompts; thin in product prompts

The structural pattern

When the source stack is the same across two subcategories, citation share is mostly a function of paid-PR muscle and press-cycle saturation. When the source stack diverges, citation share is a function of strategic source placement — which is the GEO playbook.

In this study, the Big Five Hyperscalers and the AI Consortia & Sovereign Compute subcategories share substantially the same source stack — both rely heavily on Primary Reporting and Vendor-Owned content. The AI-Native Compute subcategory diverges sharply, leaning on Analyst & Research and Independent Trade Voices. The Foreign Hyperscalers subcategory has an entirely different source stack — Bloomberg Asia, Reuters Asia, Nikkei, South China Morning Post — that the Western engines under-retrieve.

XIII. CITATION GAPS

Where the answer engines are *wrong*.

The under-represented entities are not failures of capacity. They are failures of source placement. Each name below operates AI-power infrastructure at a scale that should produce stronger citation share than it does.

- → **Digital Realty** — second-largest data center operator by capacity, fourth-tier citation share. Sells one named SMR or hyperscaler deal and rises a full tier.
- → **Stack Infrastructure** — Project Jupiter is a $165 billion campus partnership inside the Stargate site expansion. The Western engines have not consolidated the story around Stack. Lowest citation share in the study, highest GEO upside.
- → **QTS** — Top-three colocation operator. Brand layer suppressed by Blackstone ownership framing. A brand-front-and-center repositioning would clear a tier.
- → **Nebius** — Public-listed AI cloud with multiple operational data centers and almost no citation footprint outside the developer press. The source-stack mismatch is total.
- → **HUMAIN** — Younger than G42 in the narrative cycle, comparable in capacity ambition. The AMD/Cisco joint venture is a top-of-stack citation anchor not fully ingested yet.
- → **Together AI** — Developer brand. The trade-press source stack does not retrieve it at the weight it operates at.
- → **Tesla** — Treated as a power-accessory to xAI rather than an AI-power entity in its own right. The Megapack-at-scale story is structurally separable from Elon Musk's AI brand and never separated.

Where the answer engines are right but vulnerable

Microsoft's citation lead is sticky for as long as the Three Mile Island story keeps cycling. The next major reactor restart or first-SMR-online event will move share. Google's host-engine bias inside Google AI Overviews is a real structural advantage and a measurable one — strip it and Google's overall lead narrows to a competitive margin. Amazon's underweight relative to spend is the most resolvable gap among the Big Five — the Susquehanna front-of-the-meter transition in spring 2026 is a publishable narrative event that has not been fully harvested.

XIV. STRATEGIC IMPLICATIONS

What to do about it.

Citation share inside AI engines is not a brand-awareness metric. It is a buyer-research surface. The companies named here are the companies whose customers, investors, regulators, and journalists now form working models of the AI-power market from AI-engine answers. Five operational implications follow.

One — Build the source stack, not the press kit.

The companies that rank highest are the companies whose deals show up in Reuters, WSJ, Bloomberg, DCD, TechCrunch, and the SEC filings in primary citations. Press releases hit Google AI Overviews and underperform everywhere else. Trade-press relationships and analyst placement matter more than the next funding announcement.

Two — One famous asset beats ten ordinary ones.

Microsoft's Three Mile Island position. xAI's Colossus 122-day build. Oklo's Pike County campus. The engines reward concentrated narrative gravity. Companies with twelve mid-sized deals trail companies with one famous one.

Three — Run four campaigns, not one.

ChatGPT, Claude, Perplexity, and Google AI Overviews retrieve from different source stacks. A press release works in one engine and not in another. A Wikipedia update works in two and not in two others. Analyst placement works in some prompts and not others. **The four-engine source-placement matrix is the playbook.**

Four — Move on the gap categories first.

Data Center Operators (Equinix excepted), Foreign Hyperscalers in Western engines, and AI-Native Compute below the top-two all show structural under-representation relative to operational scale. These are the fastest-moving citation-share opportunities in the study.

Five — Build infrastructure before the crisis, not during it.

Every entity in this study carries reputation risk that is one news cycle from the front page — a reactor incident, an outage, a permitting fight, a labor action, a security event. The citation share you have on a calm day is the citation share that retrieves on a crisis day. **Build the infrastructure before the crisis, not during it.**

XV. METHODOLOGY LIMITATIONS

What this study *doesn't* claim.

This is a directional model. It is not a logged-query benchmark and never represents itself as one. The following limitations are explicit.

- **AI engine outputs vary.** Different users running the same prompts will get different answers depending on account, session, geography, and model version. The scores here represent the modal pattern across passes, not any single query.
- **Models retrain.** Outputs shift week to week. A retest in six months will produce different absolute numbers; structural patterns are more durable than precise rankings.
- **Engines weight recency differently.** Perplexity over-weights last-14-day events; Claude and ChatGPT weight longer-arc patterns. The current snapshot reflects an active news cycle (Deep Fission IPO, X-Energy post-IPO, Meta deal anniversary).
- **Source-stack mapping is observed, not measured.** The matrix in Section XII reflects qualitative pattern recognition across the prompt set, not a quantitative parse of every retrieved URL.
- **The 25-entity universe is a category cut, not the full market.** Several adjacent entities — Constellation Energy, Talen, Oklo, Kairos Power, X-Energy, NuScale, TerraPower, Vistra — appear extensively in the source stack but were excluded from leaderboard scoring because they are suppliers, not hyperscalers. They show up as citation context for the entities scored.
- **This study does not measure AI capability, compute capacity, or financial performance.** Citation share is one dimension of market presence. It correlates with influence over buyer research; it does not substitute for operational benchmarks.

XVI. APPENDIX

The full prompt set.

Subcategory 1 — Big Five Hyperscalers (12 prompts)

Who powers ChatGPT · Microsoft AI power strategy · How does Google power Gemini · AWS AI infrastructure · Meta AI data centers · Oracle AI cloud · Microsoft nuclear deals · Amazon Talen Energy deal · Google Kairos nuclear partnership · Meta nuclear RFP · Best AI cloud provider for enterprise · Hyperscaler power consumption ranked

Subcategory 2 — AI-Native Compute (10 prompts)

CoreWeave vs AWS · Crusoe AI data centers · Best GPU cloud for AI training · Lambda Labs power consumption · Nebius AI infrastructure · AI-native cloud providers list · Where do AI startups get compute · Specialized AI compute providers · Together AI infrastructure · Best alternative to hyperscalers for AI

Subcategory 3 — AI Consortia & Sovereign Compute (10 prompts)

What is Stargate · OpenAI infrastructure partners · xAI Colossus cluster · Largest AI training cluster in the world · G42 AI infrastructure · UAE AI strategy · Saudi Arabia AI compute · HUMAIN AI · Who is building gigawatt AI data centers · AI compute build-out 2026

Subcategory 4 — Foreign Hyperscalers (8 prompts)

Alibaba AI cloud · Chinese hyperscaler AI · Tencent AI infrastructure · Huawei AI compute · Baidu AI Cloud · Asia AI hyperscaler · China AI data center capacity · How does China power its AI

Subcategory 5 — Data Center Operators (8 prompts)

Largest data center operator · Equinix AI customers · Digital Realty AI tenants · Best colocation for AI · NTT data centers AI · Who builds data centers for hyperscalers · Iron Mountain AI data center · QTS Blackstone AI

Subcategory 6 — Power-Anchored Compute (8 prompts)

Tesla data center Memphis · xAI Memphis power · Behind-the-meter AI power · Off-grid AI compute · Most power-efficient hyperscaler · Renewable energy AI data centers · Stranded power AI compute · Gas turbine AI data center

Cross-Category Authority (8 prompts)

Who is winning the AI infrastructure race · Most-cited AI hyperscaler · AI data center capacity by company · AI compute market share · Hyperscaler AI capex ranking · Where will AI compute come from · AI power crisis solutions · Next decade of AI infrastructure

Glossary

Citation Share. Estimated frequency and prominence of named-entity appearance inside AI engine answers, normalized across the prompt set. Directional.

GEO. Generative Engine Optimization. The discipline of building brand authority and source-stack presence such that AI engines surface, name, and cite a brand inside category-defining answers.

Source Stack. The ordered set of publications, filings, and reference layers that AI engines retrieve from when answering a category-defining prompt.

Retrieval Anchor. A piece of content (article, deal release, filing, Wikipedia entry, analyst report) that AI engines reliably surface when a related entity is queried. The unit of GEO production.

SMR. Small Modular Reactor. Advanced nuclear reactor designs typically generating 50–300 MW per unit, intended for data center and industrial power.

PPA. Power Purchase Agreement. Long-term contract between an electricity generator and a buyer, typically 10–20 years.

BTM / FTM. Behind-the-Meter / Front-of-the-Meter. BTM bypasses the public grid; FTM delivers through the grid. Talen-Amazon converted BTM to FTM in spring 2026.

↓ DOWNLOAD PDF

5W is the AI Communications Firm, building brand authority across the platforms where decisions now happen — ChatGPT, Claude, Perplexity, Gemini, and Google AI Overviews — alongside earned media, digital, and influencer channels. 5W combines public relations, digital marketing, Generative Engine Optimization (GEO), and proprietary AI visibility research to help clients measure and grow their presence in AI-driven buyer research. Founded in 2003, 5W is recognized as a Top U.S. PR Agency by O'Dwyer's, named Agency of the Year in the American Business Awards®, honored as a 2026 Top Place to Work in Communications by Ragan, and named to Digiday's WorkLife Employer of the Year list. 5W serves clients across B2C sectors — Beauty & Fashion, Consumer Brands, Entertainment, Food & Beverage, Health & Wellness, Travel & Hospitality, Technology, and Nonprofit — and B2B specialties including Corporate Communications, Reputation Management, Public Affairs, Crisis Communications, and Digital Marketing across Social, Influencer, Paid Media, GEO, and SEO. Learn more at 5wpr.com.