

INFORMATION MATTERS

CATEGORY REPORT // FOUNDATION MODEL PROVIDERS

Foundation Model Providers: The 2026 Landscape

A nine-vendor scoring through the IM Framework - two structural tiers, compute supply as the binding constraint, and the European story consolidated to two entities.

May 2026**Martin De Saulles | Principal Analyst****Fiona | AI Research Analyst**

Information Matters | May 2026

Report #IM109**Headline findings**

- * **The cohort is two-tiered, not nine-way comparable.**
OpenAI, Anthropic, Google DeepMind and Meta operate at a structural scale - combined disclosed run-rate revenue well above \$50bn, combined committed forward compute spend in excess of \$1 trillion through 2032 - that the other five entities are not.
- * **Defensibility concentrates with audited filers; Disruption with model-and-product alignment cases.**
Google DeepMind sits at the cohort's structural Defensibility peak; OpenAI and Anthropic carry the highest Disruption Potential. Robustness ranges from 8.03 (Google DeepMind) to 3.69 (Aleph Alpha).
- * **Anthropic's 28 May \$65bn raise at \$965bn leapfrogged OpenAI's \$852bn.**
Anthropic is now the cohort's most valuable pure-play AI lab - a structural shift within the briefing window.
- * **The Cohere-Aleph Alpha merger reshapes the European story but has not yet closed.**
24 April 2026 all-stock acquisition at approximately \$20bn combined entity, plus \$600m Schwarz Series E. Awaits German FDI and Canadian ICA review.
- * **Compute supply, not capability, is the binding constraint at the frontier.**
Three top-tier vendors publicly signalled compute-constrained operations in Q1 2026. The next eighteen months read more as a supply story than a capability story.
- * **Frontier pricing has stopped declining; openness is doing the price-decline work.**
GPT-5.5, Claude Opus 4.8 and Gemini 3 Pro converge in a \$2.50-5 input / \$15-30 output band; DeepSeek V3.2 at \$0.28/\$0.42 anchors the open-weight track two orders of magnitude below.

1. EXECUTIVE SUMMARY

This report scores nine foundation-model providers active in agentic AI as of May 2026 - OpenAI, Anthropic, Google DeepMind, Meta, Mistral AI, Cohere, DeepSeek, xAI (now a SpaceX subsidiary), and Aleph Alpha - through the full IM Framework. The cohort's structural shape, not any single benchmark, is the central editorial finding.

Key findings

- **Two tiers, not nine-way comparable.** OpenAI, Anthropic, Google DeepMind and Meta sit at a structural scale - combined disclosed run-rate revenue approaching \$100bn (Anthropic alone crossed \$47bn in May 2026), forward compute commitments above \$1 trillion through 2032 - that the other five entities are not.
- **Anthropic leapfrogged OpenAI on valuation.** The 28 May 2026 \$65bn raise at \$965bn post-money (FT) overtook OpenAI's \$852bn. Anthropic stated run-rate revenue crossed \$47bn in May - more than a fivefold increase across the year.
- **Compute supply is the binding constraint, not capability.** Three top-tier vendors signalled compute-constrained operations in 2026. NVIDIA's Q1 FY27 print confirmed the supply-led reading: Data Center \$75.2bn (+92% YoY), Q2 FY27 guide \$91bn.
- **Robustness is the buyer's optimisation target.** Eight-scenario robustness leads at Google DeepMind (8.03) and trails to Aleph Alpha (3.69). For multi-year integration buyers, the Trajectory shape matters more than the macro score.
- **Disclosure quality varies by an order of magnitude.** Audited 10-Q filers (Alphabet, Meta) sit at one end; press-leaked-and-now-S-1 (xAI through SpaceX); undisclosed at the other (DeepSeek, Aleph Alpha). The Strategic Resilience score weights this variance explicitly.

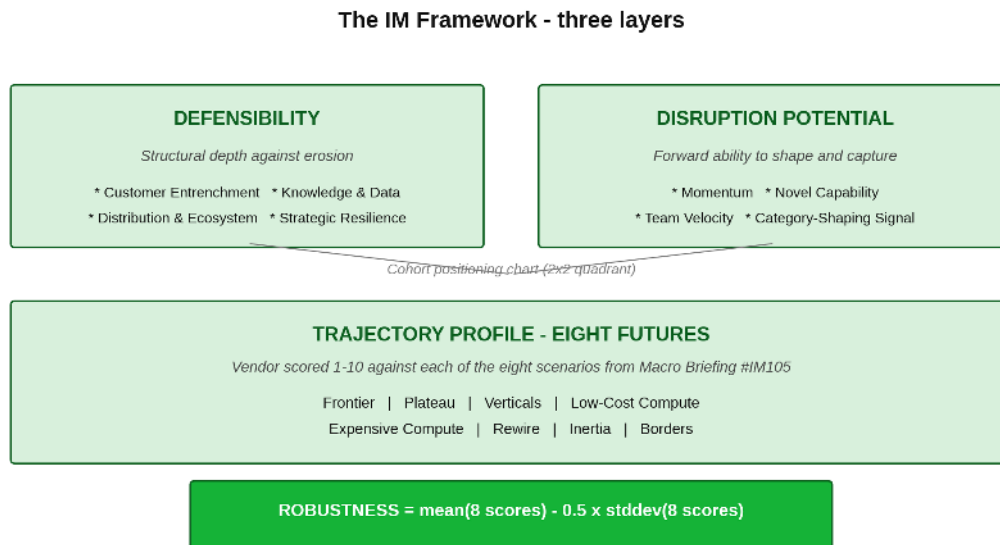


Figure 1.1 - The IM Framework: three layers, with Robustness as the buyer's optimisation target. Detailed methodology in section 2.

2. THE IM FRAMEWORK - HOW THIS REPORT SCORES

Each vendor receives two macro-construct scores. Defensibility measures the structural depth of a vendor's position against erosion - across Customer Entrenchment (depth of workflow embedment), Knowledge & Data Advantage (proprietary corpora and captive evaluation infrastructure), Distribution & Ecosystem (channel reach across hyperscalers, consumer surfaces, IDE / agent / browser integrations), and Strategic Resilience (capital base, parent-company strength, disclosure-quality variance). Disruption Potential measures the forward-looking ability to shape and capture category direction - across Momentum (revenue or adoption trajectory), Novel Capability (whether the product does something category-defining), Team Velocity (engineering and execution quality), and Category-Shaping Signal (the degree to which the vendor's moves define what the category becomes).

Each vendor also carries a Trajectory Profile - a vector of eight scores, one per scenario from the IM Macro Briefing Eight Futures for Agentic AI. The eight scenarios are Frontier (foundation-model capability keeps leaping), Plateau (open-source closes on closed-source, inference price collapses), Verticals (sector specialists prove their data advantage), Low-Cost Compute (on-device and specialty hardware scale), Expensive Compute (inference economics tighten), Rewire (enterprise restructures around agentic workflows), Inertia (deployment lags the technology), and Borders (regulatory divergence fragments the market). The robustness statistic - the mean of the eight scores minus half the standard deviation - is the single number that captures how a vendor performs across the combined future rather than in any one scenario.

Scoring note on disclosure-quality variance

The foundation-model-providers cohort splits sharply on disclosure quality, and the Strategic Resilience sub-dimension of Defensibility weights that variance explicitly. Alphabet and Meta file audited 10-Qs and 10-Ks under SEC discipline. OpenAI, Anthropic and Cohere have company-stated revenue without audited equivalent - Cohere's \$240m FY25 ARR via the CNBC-cited investor memo is the cleanest of these but is not an audited filing. xAI's \$3.2bn reported revenue is, since the February 2026 SpaceX acquisition, disclosed through the SpaceX S-1 filing. DeepSeek revenue is effectively undisclosed; Aleph Alpha revenue is not disclosed at all. The framework's scoring penalty for variance is real and is one of the load-bearing methodology decisions of this report.

3. THE FOUNDATION-MODEL-PROVIDERS CATEGORY IN 2026

The foundation-model-providers segment is the capital-intensity and run-rate-scale top of the agentic AI market. Hyperscaler capex committed to AI infrastructure across the five named providers (Alphabet, Microsoft, Amazon, Meta, Oracle) reconciles to approximately \$745-775bn of FY2026 spend, of which the genuinely AI-specific share is explicitly named only by Microsoft (approximately two-thirds short-lived assets primarily for AI on a \$190bn base) and partially by Meta. The infrastructure layer below the cohort - NVIDIA Data Center segment at approximately \$301bn annualised on the Q1 FY27 print (\$75.2bn in Q1, +92% YoY, on \$81.6bn total company revenue); AMD Data Center at approximately \$23bn - is concentrated, with hyperscalers accounting for just over 50% of NVIDIA's Data Center revenue. The category cannot be read separately from the compute supply chain that feeds it.

Three structural facts shape the cohort.

- **The cohort is two-tiered, not nine-way comparable.** OpenAI's stated \$24bn run-rate, Anthropic's \$14-47bn range over February-May 2026 (Series G plus the 28 May \$65bn follow-on round at \$965bn post-money; FT), Alphabet's \$109.9bn Q1 2026 revenue with Google Cloud at a \$70bn-plus annual run-rate, and Meta's \$56.31bn Q1 2026 revenue together sit at a structural scale that the cohort's remaining five - xAI (\$3.2bn 2025 revenue per the SpaceX S-1), DeepSeek (effectively undisclosed), Mistral (Bloomberg-stated approximately \$400m), Cohere (\$240m FY25 ARR), and Aleph Alpha (undisclosed) - cannot match by orders of magnitude.
- **The openness axis is doing real economic work.** Epoch AI's 2025 reconciliation of OpenAI's 2024 compute spend puts R&D compute at approximately \$5bn against final-training-run cost below \$1bn - approximately 80% of frontier R&D compute goes to experimentation rather than to released-model training. Pricing reflects the same axis: Western frontier-tier per-million-token pricing has converged in a \$2.50-5 input / \$15-30 output band that has stopped declining, while DeepSeek V3.2 at \$0.28 / \$0.42 anchors the open-weight track at roughly two orders of magnitude below frontier.
- **The compute supply story is the binding constraint.** Three top-tier cohort vendors have publicly signalled compute-constrained operations in Q1 2026: Anthropic explicitly on the Claude Code rate-limit tightening; Alphabet through Sundar Pichai's "we are compute constrained in the near term" statement on the Q1 2026 call; implicitly OpenAI through the stacked Azure and Stargate forward commitments. Alphabet's FY2026 capex guide stands at \$180-190bn against \$403bn FY2025 revenue. Meta's \$237.67bn of non-cancellable contractual commitments at the 31 March 2026 10-Q is up \$107bn in the quarter.

A note on the attribution rate

This report's per-vendor scoring rests on the same 25-35% foundation-model-revenue attribution rate IM's framework applies generally - the share of foundation-model API and product revenue judged genuinely agentic versus generative or chat. The four cross-checks we run on this rate land in different places depending on the question. On a narrow "share of named-agentic-SKU revenue" reading, the disclosed numbers sit well below 25-35%: Anthropic's Claude Code at \$2.5bn is approximately 8% of Anthropic's overall ARR; Salesforce Agentforce's \$800m ARR (reported with the Q4 FY2026 print on 25 February 2026) is approximately 2% of Salesforce's total annualised revenue.

On a broader "share of value-generating agentic runtime" reading that captures genuinely agentic API usage - Codex, Operator, computer-use, MCP-driven tool use, multi-step orchestration inside the foundation-model APIs - the 25-35% range is defensible. The a16z 100-CIO survey wave for 2025 captures the substance of the shift: 54% of CIOs say reasoning models accelerated LLM adoption, with these models enabling entirely new agentic workflows; "as agentic workflows have risen, switching costs have increased." The IM framework's 25-35% is a value-generating-runtime measure rather than a named-SKU-revenue measure, and the distinction is load-bearing for any reader reconciling the framework against disclosed agentic-SKU numbers.

4. COHORT POSITIONING

The 2x2 quadrant chart on the report cover plots each vendor's Defensibility (x-axis) against its Disruption Potential (y-axis), with dot size scaled to cumulative capital raised. Where capital is undisclosed (DeepSeek) the SemiAnalysis cumulative-infrastructure estimate is used as a proxy; where the vendor is internally funded by a parent (Google DeepMind, Meta), the dot size reflects the parent's FY2026 capex guide. The four named quadrants - Dominant Innovators, Disruptive Challengers, Established Incumbents, Emerging Players - frame the read.

- Dominant Innovators (top-right).** OpenAI, Anthropic, Google DeepMind, Meta. All four score above 7 on both axes. Within the quadrant Google DeepMind sits at the Defensibility extreme; OpenAI at the Disruption Potential extreme; Anthropic ties OpenAI and Meta at Defensibility 8.5, with the highest Knowledge & Data Advantage score in the cohort thanks to the Claude Code model-and-product alignment plus the 28 May \$65bn raise at \$965bn post-money that overtook OpenAI's \$852bn; Meta sits at the lower end on Disruption but the highest Strategic Resilience thanks to the ad-revenue-funded AI-capex model.
- Disruptive Challengers (top-left).** DeepSeek, xAI. Both vendors score in the top half on Disruption Potential and the bottom half on Defensibility. DeepSeek's Disruption sits on capability and cost (MIT-licensed open weights, the Nature-published R1 paper, the V3.2 launch with DeepSeek Sparse Attention), while its Defensibility is bounded by single-funder dependency (High-Flyer), undisclosed revenue, and US-states / EU / Korea regulatory restrictions. xAI's Disruption sits on the largest single named training cluster (Colossus 2 targeting 555,000 GPUs) and Grok 4.x's HLE scores; xAI became a SpaceX subsidiary on 2 February 2026 at \$250bn within the \$1.25 trillion combined entity.
- Established Incumbents (bottom-right).** Cohere, Mistral AI. Cohere's 6.2 / 6.5 profile is the textbook Established Incumbent for the enterprise-only segment - 70% gross margin, \$240m FY25 ARR, named Canadian-federal-government and Royal-Bank-of-Canada deployments, the explicit "ROI not AGI" positioning. Mistral's 6.0 / 6.75 profile carries the additional sovereignty dimension - French and German government partnerships, the SAP-led sovereign-AI stack, ASML at 11% on the cap table.
- Emerging Players (bottom-left).** Aleph Alpha. A single cohort vendor sits below the cohort median on both axes. Aleph Alpha's standalone position scores at the cohort floor on Defensibility (4.2; reflecting the early-2026 layoff of approximately 50 staff, the October 2025 leadership change with Andrulis stepping down, and the undisclosed revenue base) and the cohort floor on Disruption (4.0; reflecting the explicit decision since 2024 to stop competing on frontier-model performance). The framework treats the position honestly: the pending Cohere acquisition is the structural resolution.

The Cohere-Aleph Alpha merger as a load-bearing quadrant event

The pending merger is the most important structural event in the European segment of the cohort. The deal terms - all-stock combination at a \$20bn combined-entity valuation, Cohere shareholders to own approximately 90% post-close, Aleph Alpha shareholders approximately 10%, concurrent \$600m Series E commitment from Schwarz Group as the anchor investor - collapse the European frontier-lab cohort to Mistral plus a transatlantic combined entity.

The framework's read of the post-close combined entity, were it scored as a single cohort member at close, would inherit Cohere's commercial profile (the \$240m ARR, 70% gross margin, enterprise-only positioning) plus Aleph Alpha's sovereign-deployment capabilities (Bundeswehr clearance, the STACKIT runtime, on-premises and air-gapped delivery). The Borders sub-dimension of the Trajectory Profile would score materially higher for the combined entity than for either standalone vendor - the combined sovereign footprint across Canada, Germany and the broader EU is the only one in the cohort that sits credibly in two major jurisdictions at once. Until close, we hold the scoring at the standalone-vendor level and flag the deal in both vendor profiles, in this section, and in the cross-cohort patterns section.

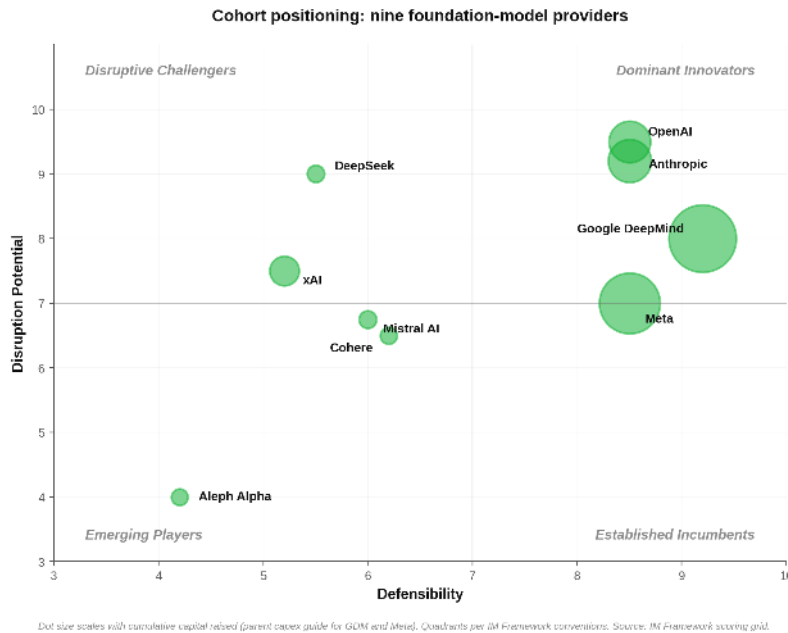



Figure 1. Cohort positioning: Defensibility (x-axis) against Disruption Potential (y-axis); dot size scaled to cumulative capital raised. Named quadrants per IM Framework conventions.

5. TRAJECTORY PROFILES - THE CROSS-COHORT VIEW

The Trajectory Profile heat-map below shows each vendor's eight-scenario score vector, sorted top-to-bottom by robustness. Cell shading uses the Blues sequential palette - light blue for low scores, dark blue for high - per the framework conventions; the palette deliberately carries no value judgement, in contrast to the green-gradient bar chart on each per-vendor profile page.



Eight-scenario Trajectory Profile from Macro Briefing #IM105 "Eight Futures for Agentic AI". Sequential Blues palette deliberately carries no value judgement. Source: IM Framework scoring.

Figure 2. Cross-cohort Trajectory Profile heat-map. 9 vendors x 8 scenarios; sorted by robustness.

The eight scenarios:

Scenario	What it describes
Frontier	Foundation-model capability keeps leaping; benchmarks stay the load-bearing axis.
Plateau	Open-source closes on closed-source; inference price collapses; output commoditises.
Verticals	Sector specialists with proprietary domain data outperform horizontal generalists.
Low-Cost Compute	On-device and specialty hardware scale; inference economics improve sharply.
Expensive Compute	Inference economics tighten; capability-cost trade-offs harden.
Rewire	Enterprises restructure their workflows around agentic operation.
Inertia	Deployment lags the technology; existing buyers do not move at the pace of capability.
Borders	Regulatory divergence fragments the market into sovereign-deployment regimes.

The cohort splits into three groups

Top tier - robust everywhere (Google DeepMind alone)

Google DeepMind's robustness of 8.03 is more than 1.4 points above the next vendor (Meta at 6.63) and reflects scores of 7 or higher in every single scenario. The Trajectory shape is the cleanest cohort-internal case for what robust-everywhere actually means: capability frontier through Gemini 3 Deep Think, the Plateau scenario protected by the Vertex AI distribution of third-party Claude alongside first-party Gemini, the Verticals scenario covered by AlphaFold / AlphaProteo / GNoME, and the Borders scenario covered by Google Cloud's audited cross-jurisdiction footprint.

Middle tier - capability-frontier and rewire-leading (OpenAI, Anthropic, Meta)

Three vendors in the 6.16-6.63 robustness band, all with the highest Frontier and Rewire scores in the cohort, but with one or two specific scenarios where the floor falls dangerously. OpenAI scores 9 on Frontier and 9 on Rewire (Operator, Codex, AgentKit) but only 5 on Low-Cost Compute. Anthropic scores 9 on Frontier and the cohort's only 10 on Rewire (Claude Code is the leading evidence point for agentic-native run-rate) but only 5 on Plateau and 5 on Low-Cost Compute. Meta scores 9 on Inertia and 9 on Plateau but 4 on Borders - the only cohort top-tier vendor with a Trajectory score below the cohort median on any single scenario, and the consequence of the Llama 4 license's EU exclusion combined with Meta's refusal to sign the EU GPAI Code of Practice.

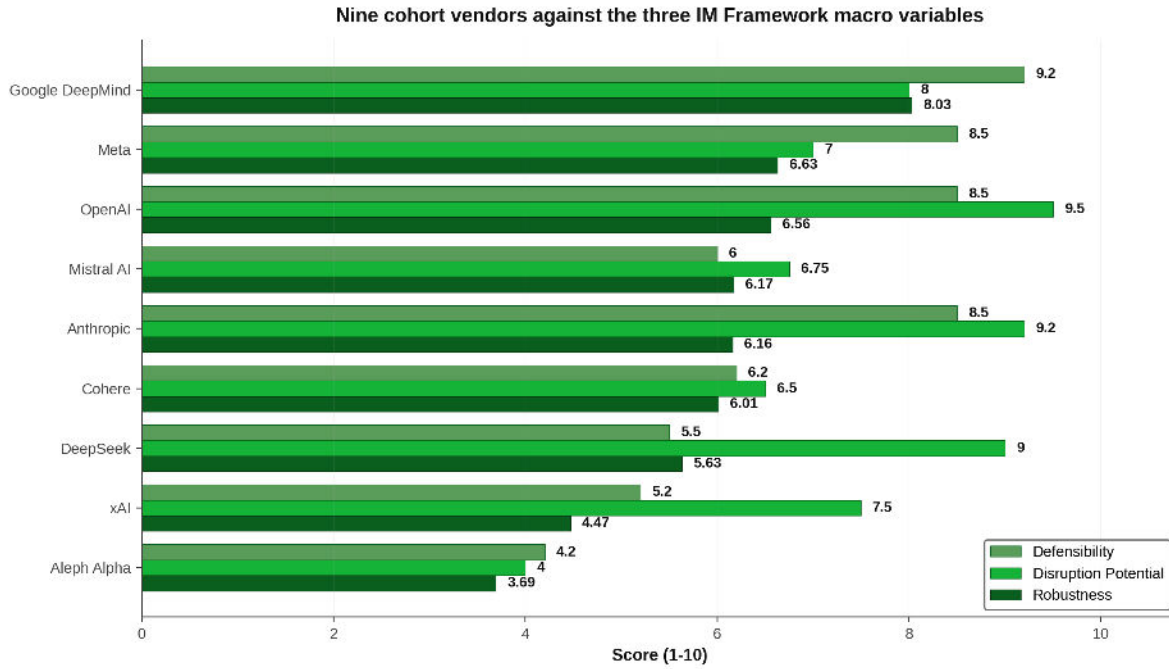
Bottom tier - two distinct modes (Mistral, Cohere, DeepSeek, xAI, Aleph Alpha)

Five vendors in the 3.69-6.17 robustness band, splitting into two distinct shapes. The European mode (Mistral 6.17, Cohere 6.01, Aleph Alpha 3.69) is anchored by high Borders scores - the sovereign-deployment thesis is the single shared structural feature across the three. The asymmetric-bet mode (DeepSeek 5.63, xAI 4.47) carries opposite shapes: DeepSeek scores 9 on Plateau, 9 on Low-Cost Compute but 3 on Borders; xAI scores 7 on Frontier and 6 on Rewire on the back of Colossus capability bets but 4 on Verticals, Low-Cost Compute and Borders.

The shape of the heat-map captures the report's central editorial finding: there is no scenario in which a single cohort vendor beats Google DeepMind on its own ground, but there are specific scenarios (Rewire for Anthropic, Plateau and Low-Cost Compute for DeepSeek, Borders for Mistral and the combined Cohere+Aleph Alpha entity) where a non-top-tier vendor outperforms the cohort average on the scenario the vendor's strategy is most aligned with. For buyers whose own scenario weights are not the framework's defaults, the per-scenario reads on the heat-map are the more useful artefact.

6. VENDOR COHORT - NINE PROFILES

Nine per-vendor profiles follow, ordered by Trajectory Profile robustness. Each carries the macro-construct scores in narrative form, the per-vendor Trajectory bar chart on the page (single-hue green gradient), and the Trajectory commentary paragraph naming best- and worst-positioned scenarios. The chart below summarises all nine vendors against the three Framework macro variables on a single page before the profiles unpack the detail.



Vendors ordered top-to-bottom by Robustness. Defensibility and Disruption Potential are macro-composite scores; Robustness is $\text{mean}(9 \text{ Trajectory scores}) - 0.5 \times \text{stddev}$. Source: IM Framework scoring grid.

Figure 6.0 - Cohort summary: Defensibility, Disruption Potential, and Robustness across the nine vendors, sorted by Robustness.

6.1 Google DeepMind (Alphabet)

Defensibility 9.2 / Disruption Potential 8.0 / Robustness 8.03.

Defensibility.

The strongest sub-dimensions are Distribution & Ecosystem and Strategic Resilience, both at 10. Google Cloud ended 2025 at over \$70bn annual run-rate and Q1 2026 cloud backlog stood at over \$460bn; the Vertex AI distribution channel carries first-party Gemini alongside Anthropic Claude, the only major distribution surface in the cohort that does. Strategic Resilience is parent-balance-sheet sized - Alphabet's FY2025 revenue was \$402.836bn (+15% YoY) and Q1 2026 revenue was \$109.9bn (+22% YoY) with audited 10-K and 10-Q filings. Customer Entrenchment is reinforced by the Workspace footprint and the AI Overviews / AI Mode roll-through to Search. Knowledge & Data Advantage is the cohort's strongest in scientific-AI assets (AlphaFold 3, AlphaProteo, AlphaProof, GNoME) and is reinforced by Pichai's Q1 2026 disclosure that AI Overviews / AI Mode cost-per-core-response fell more than 30% after the Gemini 3 upgrade.

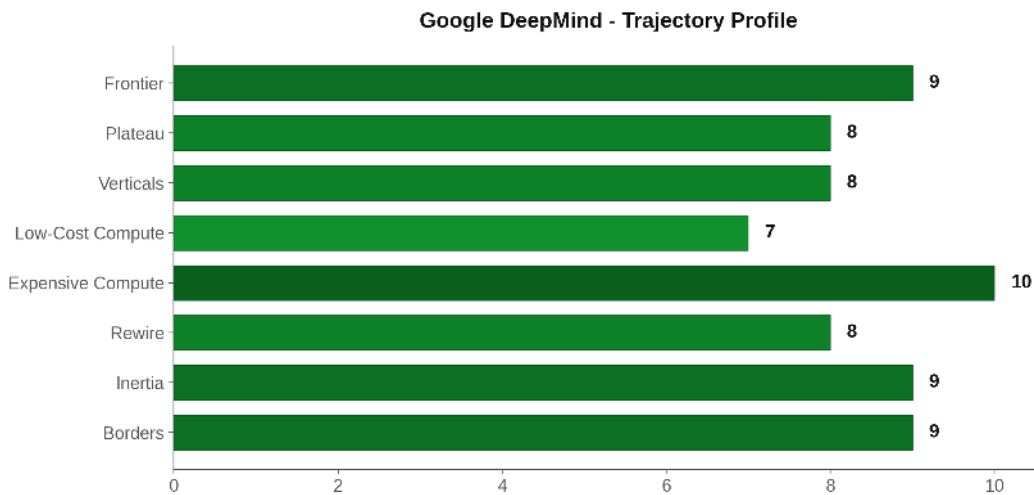
Disruption Potential.

Momentum is high - Gemini Enterprise paid monthly active users grew 40% QoQ in Q1 2026 and Pichai stated "revenue from products built on Alphabet's gen AI models grew nearly 800% year-over-year in Q1." Novel Capability is meaningful through Gemini 3 Deep Think (reported 45.1% on ARC-AGI-2 versus 4.9% for Gemini 2.5 Pro, 40%-plus on Humanity's Last Exam) and through the generative-media stack (Veo 3, Imagen 4, Lyria 2) plus the Genie world-models programme. Team Velocity is parent-organisation-sized but the November / December 2025 cadence (Gemini 3 Pro, Flash, Deep Think within four weeks) shows the org is shipping at frontier pace. Category-Shaping Signal is moderate - shaping through capability releases and the AlphaFold-Nobel halo, but does not set the run-rate cadence the way OpenAI and Anthropic do.

Trajectory commentary.

Best-positioned for Inertia (deep Google Cloud customer integration plus the Workspace footprint survives any slow-adoption future), for Borders (the cross-jurisdiction Google Cloud audited footprint is incumbent-grade across every major region), and for Expensive Compute (the captive TPU stack and the Pichai-disclosed 30%-plus unit-economics improvement on Gemini 3 mean inference-cost inflation hurts less than at the API-tier-only competitors). Worst-positioned for Low-Cost Compute - a future in which on-device and consumer-grade inference dominates is broadly neutral but not actively favourable; Gemini Nano addresses it, but the value proposition is not optimised for the scenario.

6.1 Google DeepMind (Alphabet) (continued)



Eight scenario Trajectory per IM Framework Macro Briefing #IM105 Eight Futures for Agentic AI. Darker green = higher score. Source: IM scoring grid.

Figure 3.1. Google DeepMind - Trajectory Profile across eight scenarios (green gradient).

6.2 Meta (Llama / Meta AI / Superintelligence Labs)

Defensibility 8.5 / Disruption Potential 7.0 / Robustness 6.63.

Defensibility.

The strongest sub-dimension is Strategic Resilience at 10 - Meta's Q1 2026 revenue was \$56.31bn (+33% YoY) with operating income of \$22.9bn (+30%) and net income of \$26.77bn (+61%), and the AI capex is funded out of advertising cash flow rather than out of foundation-model API revenue. Distribution & Ecosystem is similarly first-tier through Facebook, Instagram, WhatsApp, Messenger, the standalone Meta AI app, and Llama distribution through Together AI, Fireworks, Groq, AWS Bedrock and Azure. Customer Entrenchment is reinforced by the 8 million advertisers using Meta's AI ad-creative tools (up from 4 million at end-2024), the Value Optimization suite at over \$20bn annual run-rate, and Partnership Ads at \$10bn Q1 2026 run-rate. Knowledge & Data Advantage is moderate - the Llama corpus is real, but the family is now competing on the open-weight benchmark frontier against Chinese-origin models (Qwen, DeepSeek, GLM).

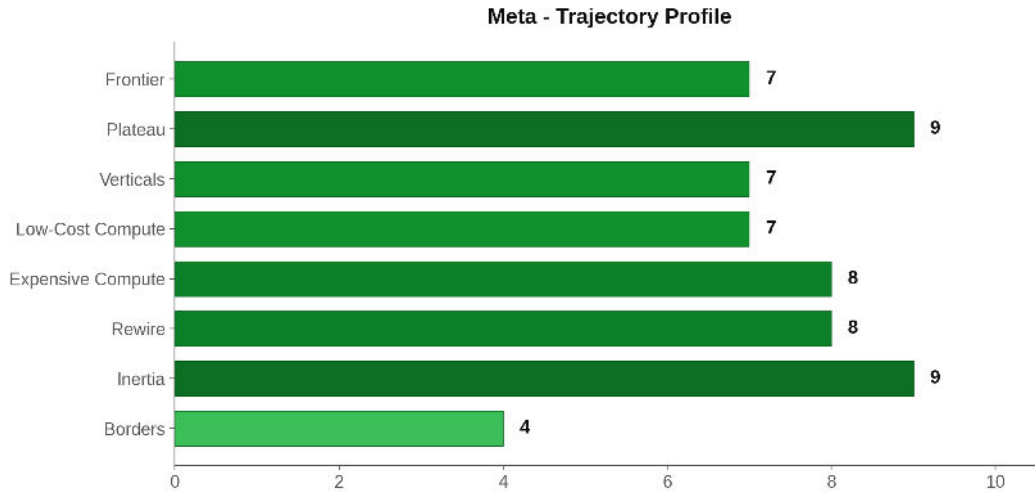
Disruption Potential.

Momentum is steady but not category-defining - Meta does not break out a separate "AI revenue" line, and the AI investment-return question is the central point of the Q1 2026 after-hours sell-off. Novel Capability has shifted from Llama 4 (April 2025; drew open-source community criticism for a perceived rushed and opaque release) to Muse Spark (April 2026; MSL's first model; closed-source). Team Velocity has been disrupted by three reorganisations of the AI org inside 12 months (June 2025 Alexandr Wang hire as Chief AI Officer; August 2025 split into four MSL sub-units; October 2025 cuts of approximately 600 roles) plus the departures of Joelle Pineau (joined Cohere) and Yann LeCun. Category-Shaping Signal is meaningful - Llama set the open-weight reference for two years - but is now contested by the Chinese-origin labs that have moved ahead on rigorous benchmarks.

Trajectory commentary.

Best-positioned for Inertia (the consumer-platform and advertiser-base scale survives a slow-adoption future cleanly; Meta does not need to convert AI capex into AI subscription revenue), for Plateau (the open-weight Llama strategy is the textbook beneficiary of foundation-model commoditisation), and for Expensive Compute (the cash-flow base absorbs inference-cost inflation cleanly). Worst-positioned for Borders - and this is the load-bearing weakness in Meta's Trajectory profile. The Llama 4 Community License explicitly does not apply to individuals domiciled in, or companies with principal place of business in, the EU; Meta is the only major frontier provider in the cohort that refused to sign the EU GPAI Code of Practice. A future of regulatory fragmentation hits Meta harder than any other cohort top-tier vendor.

6.2 Meta (Llama / Meta AI / Superintelligence Labs) (continued)



Eight scenario Trajectory per IM Framework Macro Briefing #IM105 Eight Futures for Agentic AI. Darker green = higher score. Source: IM scoring grid.

Figure 3.2. Meta - Trajectory Profile across eight scenarios (green gradient).

6.3 OpenAI

Defensibility 8.5 / Disruption Potential 9.5 / Robustness 6.56.

Defensibility.

The strongest sub-dimensions are Customer Entrenchment and Distribution & Ecosystem, both at 9. ChatGPT's 800 million weekly active users (publicly stated at DevDay October 2025, with approximately 900 million referenced in early-2026 fundraising materials) is the deepest consumer-AI seat base in the category. The "1 million businesses" deployed figure and the Q1 2026 disclosure that Enterprise revenue now exceeds 40% of total are the enterprise-tier scale evidence. Distribution runs through Microsoft Azure (IP rights through 2032), Apple Intelligence integration in iOS, the ChatGPT App Store and AgentKit announced at DevDay 2025, the standalone consumer apps, and the io / Jony Ive hardware collaboration (\$6.5bn in stock). Strategic Resilience is at 8 - the \$852bn post-money valuation at the \$122bn primary round on 31 March 2026, combined with the \$250bn Azure commitment plus the Stargate forward commitments (within which Oracle's \$300bn deal is the principal counterparty), gives OpenAI a capital position no third-party AI vendor can match, but the Foundation-controlled board structure is unusual.


Disruption Potential.

Momentum is the highest in the cohort at 10 - the \$24bn annualised run-rate at the March 2026 round close (up from approximately \$100m at end-2023) is the fastest revenue ramp in the category's history. Novel Capability is category-defining through GPT-5.x, the Operator agentic-browsing product, the Codex line, and the AgentKit framework. Team Velocity is high. Category-Shaping Signal is the highest in the cohort at 10 - OpenAI sets the cadence the rest of the category is read against, and the GPT-5.5 launch on 23 April 2026 was the inflection that broke the multi-year price-decline narrative at the frontier.

Trajectory commentary.

Best-positioned for Frontier (GPT-5.x leads consumer and enterprise capability benchmarks; the model-and-product alignment with Operator and Codex is the textbook frontier case), for Rewire (Operator and Codex are agentic-native products that ride the restructured-workflow scenario directly), and for Expensive Compute (the \$852bn capital base absorbs inference-cost inflation cleanly). Worst-positioned for Low-Cost Compute - the GPT-5.5 pricing at \$5 / \$30 per million tokens is the explicit evidence that OpenAI is not optimising for a cheaper-inference future - and for Plateau (the value proposition depends on capability lead; commoditisation of foundation-model output compresses the moat against open-weight competitors). The Microsoft AGI verification clause is novel contract language with no precedent.

6.3 OpenAI (continued)



Eight scenario Trajectory per IM Framework Macro Briefing #IM105 Eight Futures for Agentic AI. Darker green = higher score. Source: IM scoring grid.

Figure 3.3. OpenAI - Trajectory Profile across eight scenarios (green gradient).

6.4 Mistral AI

Defensibility 6.0 / Disruption Potential 6.75 / Robustness 6.17.

Defensibility.

The strongest sub-dimension is Distribution & Ecosystem at 7 - hyperscaler partnerships across Microsoft Azure (since February 2024), AWS Bedrock and SageMaker, Google Cloud Vertex AI, IBM watsonx, Nvidia NIM, and Snowflake; plus the late-2025 sovereign-AI stack partnership with SAP and the French and German governments. Customer Entrenchment is moderate at 6, anchored by publicly named European enterprise customers including BNP Paribas, AXA, Stellantis (in-car AI), Veolia, HSBC, and Free Mobile. Knowledge & Data Advantage is moderate at 6 - Mistral's Apache-2.0-licensed open-weight models established the company's reputation, but the open-weight performance floor has been pulled up sharply by Llama, DeepSeek and Qwen. Strategic Resilience is the weakest sub-dimension at 5 - the EUR 11.7bn (approximately \$13.8bn) post-money valuation at the September 2025 Series C is robust but the \$830m debt raise for the Bruyeres-le-Chatel data centre in March 2026 plus the 1.4 GW Paris-region AI-campus joint venture imply a Series D or pre-IPO round in 2026-27.

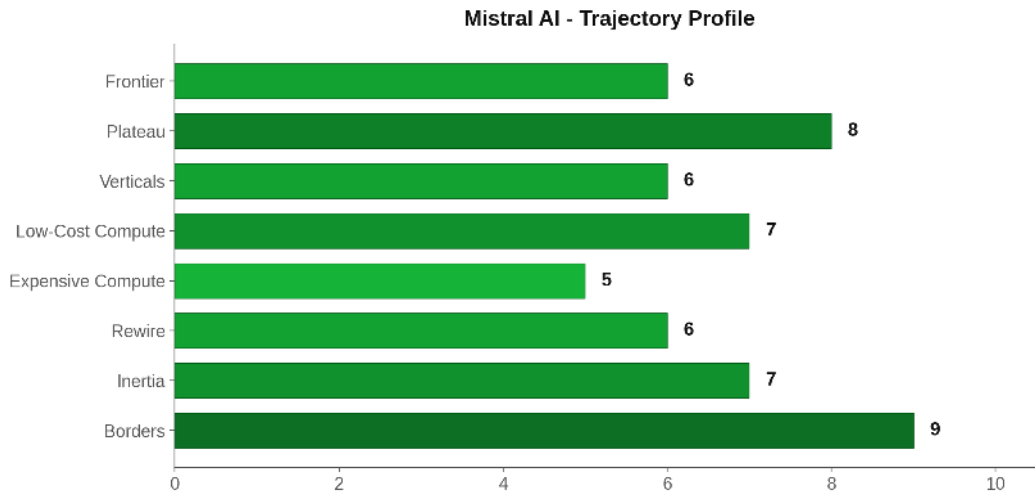
Disruption Potential.

Momentum is real - Arthur Mensch told Bloomberg at Davos that Mistral expects to exceed EUR 1bn in revenue in 2026, up from a reported \$400m ARR in mid-2025 (FT). Novel Capability sits at 6 - Mistral Medium 3 (May 2025) and Mistral Large 3 (May 2026) are credible flagship models, but the company has explicitly stopped competing on capability-frontier benchmarks against the Big Three. Team Velocity is high. Category-Shaping Signal sits at 7 - Mistral is the only credible European foundation-model provider at independent-frontier scale, and Macron's repeated public endorsements of Le Chat over ChatGPT plus the Paris AI Action Summit have framed Mistral as the European national-champion case.

Trajectory commentary.

Best-positioned for Borders - Mistral's Trajectory profile carries the cohort's highest single Borders score at 9, reflecting the French and German government partnerships, the SAP sovereign-AI stack, and the ASML controlling-influence stake (approximately 11% with a seat on Mistral's Strategic Committee). The position is robust to regulatory fragmentation in a way no US cohort vendor's is. Best-positioned also for Plateau (Apache-2.0 lower-tier models position Mistral as an open-weight beneficiary of foundation-model commoditisation). Worst-positioned for Expensive Compute - the \$830m Bruyeres-le-Chatel debt-financed build implies a unit-economics base that inference-cost inflation would tighten on - and for Frontier, the cohort's lowest single Frontier score outside Aleph Alpha.

6.4 Mistral AI (continued)



Eight scenario Trajectory per IM Framework Macro Briefing #IM105 Eight Futures for Agentic AI. Darker green = higher score. Source: IM scoring grid.

Figure 3.4. Mistral AI - Trajectory Profile across eight scenarios (green gradient).

6.5 Anthropic

Defensibility 8.5 / Disruption Potential 9.2 / Robustness 6.16.

Defensibility.

The strongest sub-dimension is Knowledge & Data Advantage at 9 - the Claude family scores at the top of public coding and reasoning benchmarks, and the parent-organisation training pipeline is captive to Claude Code in a way no multi-foundation wrapper can replicate. Strategic Resilience steps up to 9 - the \$965bn post-money valuation now makes Anthropic the cohort's most capitalised pure-play AI lab, ahead of OpenAI at \$852bn. The \$380bn post-money Series G on 12 February 2026 (\$30bn raised), the 28 May 2026 follow-on round of \$65bn at \$965bn post-money (Altimeter-led; FT 28 May 2026), the Google up-to-\$40bn investment plus Broadcom 3.5 GW deal, the Amazon \$13bn cumulative direct investment plus the more-than-\$100bn-over-ten-years AWS commitment, the November 2025 Nvidia-Microsoft-Anthropic alliance with at least \$30bn of Azure spend, and the multibillion-dollar SpaceX data-centre lease, together put Anthropic in a capital position no third-party AI startup can match. Distribution at 8 runs through AWS Bedrock (Claude is the headline model on Bedrock), Google Vertex AI, Microsoft Foundry, GitHub Copilot, plus first-party Claude.ai, Claude Code and Claude for Chrome. Customer Entrenchment at 8 is reinforced by "more than 300,000 business customers" stated at the Series G.

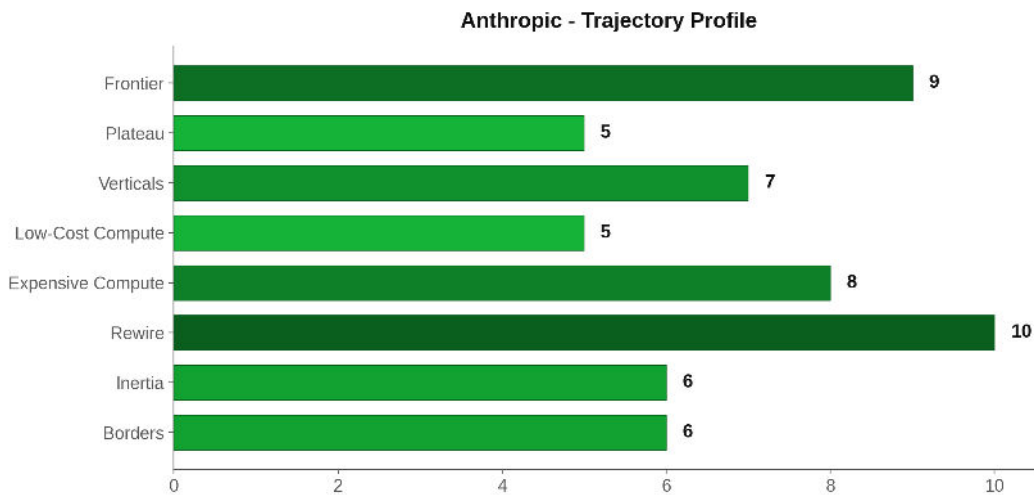
Disruption Potential.

Momentum is the cohort's joint-highest at 10 - Claude Code's run-rate above \$2.5bn (Series G; "more than doubling since the beginning of 2026") and Anthropic's total ARR scaling from approximately \$10bn at the start of 2026 to \$14bn at the Series G in February to a stated \$47bn run-rate by May 2026 (Anthropic, via FT 28 May 2026) sets the agentic-AI revenue cadence the rest of the category is read against. Claude Opus 4.8 launched 28 May 2026 alongside the funding round. Novel Capability is category-defining for the workflow-first, sub-agent, computer-use-adjacent design pattern; MCP (Model Context Protocol) is now an industry-standard tool-use rail. Team Velocity is high. Category-Shaping Signal at 9 is the cohort joint-highest with DeepSeek - Claude Code is the leading evidence point for agentic-native run-rate in the entire category.

Trajectory commentary.

Best-positioned for Rewire - Claude Code's \$2.5bn-plus ARR is the cohort's only Trajectory score of 10, the strongest evidence point in the category that the agentic-native architecture pattern monetises at frontier scale - and for Frontier (model-and-product alignment with the Claude family rides capability frontier directly). Worst-positioned for Plateau (the value proposition leans on Claude's capability lead; commoditisation of foundation-model output compresses the moat directly) and for Low-Cost Compute (Anthropic's pricing model relies on the inference economics the Low-Cost Compute scenario would erode). Counterparty concentration is unusual - Anthropic is simultaneously committed to Amazon, Google, and Microsoft, with SpaceX/Colossus involvement widely reported but not first-party confirmed.

6.5 Anthropic (continued)



Eight scenario Trajectory per IM Framework Macro Briefing #IM105 Eight Futures for Agentic AI. Darker green = higher score. Source: IM scoring grid.

Figure 3.5. Anthropic - Trajectory Profile across eight scenarios (green gradient).

6.6 Cohere

Defensibility 6.2 / Disruption Potential 6.5 / Robustness 6.01.

Defensibility.

The strongest sub-dimension is Customer Entrenchment at 7 - publicly named enterprise customers include Royal Bank of Canada, Bell, Dell, LG, STC Group, Ensemble Health Partners, Palantir, Oracle Fusion Applications (more than 100 GenAI use cases), and the Canadian federal government via the Cohere North sovereign deployment. Strategic Resilience at 6 reflects the cleanest disclosure quality in the cohort outside the audited-filer top tier - the CNBC-cited investor memo (13 February 2026) gives \$240m FY25 ARR, approximately 70% gross margin (up approximately 25 bps YoY), quarter-on-quarter growth above 50% throughout 2025, and a \$7bn valuation post the August / September 2025 rounds. Distribution sits at 6 - Oracle Cloud Infrastructure is the primary partner; also present on Microsoft Azure AI catalog, Google Cloud Vertex AI, and AWS Bedrock. Knowledge & Data Advantage at 6 is moderate; the Aya multilingual research family is a credible open-weight contribution but the production-tier Command A models are API-only.


Disruption Potential.

Momentum is real at 7 - the \$240m FY25 ARR with >50% QoQ growth surpassed the \$200m target, and Aidan Gomez has flagged IPO-readiness "soon" since October 2025. Novel Capability sits at 6 - Command A (March 2025; 111B parameters, 256K context, runs on 2 GPUs at 150% throughput vs Command R+ 08-2024) is a credible flagship and the Cohere North agentic platform (GA August 2025) is the cohort's clearest on-premises and air-gapped sovereign-AI agent platform. Team Velocity at 7 reflects the senior-team upgrade in 2025 (Joelle Pineau as Chief AI Officer from Meta, Francois Chadwick as first CFO from Uber / KPMG). Category-Shaping Signal at 6 - Gomez has framed the position explicitly with "We don't do these blockbuster launch event type things" and "ROI not AGI."

Trajectory commentary.

Best-positioned for Borders - Cohere's Trajectory profile carries the cohort's second-highest Borders score at 8, reflecting the Cohere North sovereign-AI agent platform, the Canadian federal government deployment, and the pending acquisition of Aleph Alpha which would add Bundeswehr clearance and STACKIT runtime access. Best-positioned also for Verticals and Rewire (enterprise-only positioning maps cleanly onto sector-specific and restructured-workflow futures). Worst-positioned for Frontier - Cohere has chosen not to compete on capability-frontier benchmarks against the Big Three - and for Expensive Compute (the enterprise-margin model is exposed to inference-cost inflation in a way the API-first competitors absorb more readily). The pending Cohere-Aleph Alpha merger is the most consequential structural event in this profile.

6.6 Cohere (continued)



Eight scenario Trajectory per IM Framework Macro Briefing #IM105 Eight Futures for Agentic AI. Darker green = higher score. Source: IM scoring grid.

Figure 3.6. Cohere - Trajectory Profile across eight scenarios (green gradient).

6.7 DeepSeek

Defensibility 5.5 / Disruption Potential 9.0 / Robustness 5.63.

Defensibility.

The strongest sub-dimension is Knowledge & Data Advantage at 8 - DeepSeek's technical reports on V3 (arXiv 2412.19437, December 2024), R1 (peer-reviewed Nature paper 17 September 2025) and V3.2 (introducing DeepSeek Sparse Attention; claimed gold-medal performance on 2025 IMO and IOI) are the most detailed primary-source disclosures of training methodology in the cohort, and the company is uniquely positioned at the open-weight cost frontier. Distribution at 6 is real but bounded - the DeepSeek free chatbot app reached #1 free app on US iOS on 27 January 2025 (the "DeepSeek moment"), DeepSeek R1 derivatives on Hugging Face crossed 10 million cumulative downloads, and third-party hosting runs through Together AI, Fireworks, OpenRouter, Chinese cloud platforms, and Nvidia's enterprise stack. Customer Entrenchment at 5 is the weakest in the cohort - DeepSeek does not run a primary enterprise-sales motion. Strategic Resilience at 3 reflects the load-bearing weakness: revenue is effectively undisclosed; funding remains internal from High-Flyer Quantitative Investment Management; SemiAnalysis estimated cumulative buildout at \$1.3-1.6bn through 2025.

Disruption Potential.

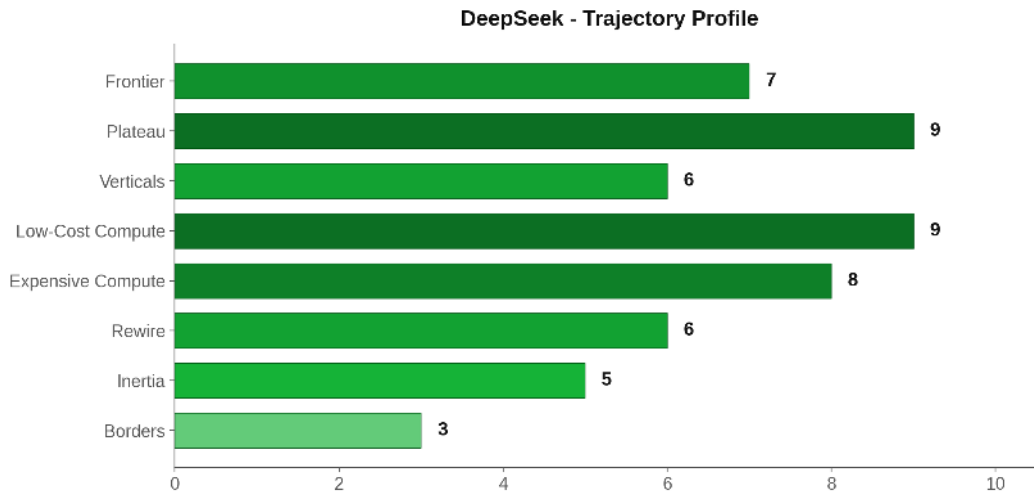
Momentum is the joint-highest in the cohort at 9 (with Anthropic) - the V3, R1 and V3.2 release cadence has reset the open-weight cost frontier. Novel Capability at 9 - DeepSeek Sparse Attention introduced in V3.2-Exp (September 2025) is a category-defining architectural contribution, and the V3 / R1 / V3.2 family's competition-reasoning scores (claimed gold-medal at 2025 IMO and IOI, MATH-500 at 97.3% for R1) are the cohort's strongest open-weight benchmark results. Team Velocity at 9 - Liang Wenfeng's team has shipped at frontier pace despite operating under H2O export-control restrictions through 2025. Category-Shaping Signal is the cohort joint-highest at 9 with Anthropic - the "DeepSeek moment" of January 2025 reset Western capital-market expectations about what a frontier-class model costs to train.

Trajectory commentary.

Best-positioned for Plateau (open-weight, MIT-licensed weights position DeepSeek as the textbook beneficiary of foundation-model commoditisation), for Low-Cost Compute (the V3.2 Sparse Attention architecture and the open-weight inference economics are the cohort's strongest fit to a cheaper-inference future), and for Expensive Compute (the open-weight third-party-hosting model insulates DeepSeek itself from inference-cost inflation).

Worst-positioned for Borders at 3, the cohort's lowest single Borders score - Texas-first US state-level bans, subsequent New York, Virginia, Tennessee, Iowa state bans, federal-device bans, South Korean and Italian national-level app-store restrictions. Worst-positioned also for Inertia - DeepSeek does not own an enterprise-procurement distribution motion.

6.7 DeepSeek (continued)



Eight scenario Trajectory per IM Framework Macro Briefing #IM105 Eight Futures for Agentic AI. Darker green = higher score. Source: IM scoring grid.

Figure 3.7. DeepSeek - Trajectory Profile across eight scenarios (green gradient).

6.8 xAI

Defensibility 5.2 / Disruption Potential 7.5 / Robustness 4.47.

Defensibility.

xAI is, as of 2 February 2026, a wholly-owned subsidiary of SpaceX following an all-stock merger that valued xAI at \$250bn within a \$1.25 trillion combined entity (share exchange at 1 xAI share = 0.1433 SpaceX shares; Musk's stated rationale was orbital-data-centre build-out). For framework purposes we score the xAI unit on its standalone foundation-model trajectory. The strongest sub-dimensions are Knowledge & Data Advantage and Distribution & Ecosystem, both at 6. The X firehose of real-time posts is a training-data and freshness asset uniquely available to xAI through the March 2025 X acquisition, and Grok distribution runs through X (Premium, Premium+, SuperGrok and SuperGrok Heavy tiers) and the grok.com standalone product. Customer Entrenchment at 5 is real in the X-subscription base but moderate at the enterprise tier. Strategic Resilience at 4 reflects the standalone xAI unit's pre-SpaceX-merger profile; the 2 February 2026 SpaceX acquisition has materially upgraded the disclosure regime (xAI 2025 standalone revenue \$3.2bn against \$6.4bn of operating loss now disclosed through the SpaceX S-1). Active litigation exposure: NAACP, SELC and Earthjustice are litigating partners on the unpermitted natural-gas turbines powering Colossus 2; by 11 May 2026 the facility had 46 unpermitted gas turbines on site (Mississippi Today).


Disruption Potential.

Momentum is high at 8 - the SpaceX S-1 has now disclosed xAI standalone 2025 revenue at \$3.2bn against \$6.4bn of operating loss - the first audited-equivalent disclosure of xAI's standalone economics. Novel Capability at 7 - Grok 4 (9 July 2025) was the first model to exceed 50% on Humanity's Last Exam at 50.7% text-only, Grok 4 Heavy reached 58.3% with multi-agent mode, and Grok 4.1 (17 November 2025) claims 65% reduction in hallucinations and topped LMArena EQ-Bench3. Team Velocity at 8. Category-Shaping Signal at 7 - xAI's positioning is "biggest cluster wins" rather than "most efficient training," and the Colossus build-out (122 days to 100,000 H100 GPUs at Colossus 1; target 2 GW total capacity and 555,000 GPUs at Colossus 2, approximately \$18bn GPU purchase) is the most aggressive single-vendor cluster build in the cohort.

Trajectory commentary.

Best-positioned for Frontier at 7 - Grok 4.x rides capability frontier directly on the back of the largest single named training cluster - and for Rewire at 6. Worst-positioned for Verticals, Low-Cost Compute and Borders, all at 4 - xAI does not have a sector-specific data advantage, the cluster-first capex-intensive position is the opposite of a Low-Cost Compute beneficiary, and the regulatory exposure is the cohort's third-highest. The single-CEO-bandwidth-concentration question (Tesla, SpaceX, Neuralink, DOGE / US government, X and xAI) plus the Babuschkin departure (August 2025) and reports of further senior-engineer exits raise key-person risk. xAI signed only the safety-and-security chapter of the EU GPAI Code of Practice.

6.8 xAI (continued)



Eight scenario Trajectory per IM Framework Macro Briefing #IM105 Eight Futures for Agentic AI. Darker green = higher score. Source: IM scoring grid.

Figure 3.8. xAI - Trajectory Profile across eight scenarios (green gradient).

6.9 Aleph Alpha

Defensibility 4.2 / Disruption Potential 4.0 / Robustness 3.69.

Defensibility.

The strongest sub-dimension is Distribution & Ecosystem at 6 - the STACKIT (Schwarz Digits) sovereign cloud platform is the contractually mandated runtime environment for PhariaAI-as-a-Service, and the publicly disclosed customer set includes Bundeswehr (multi-year Cerebras deal announced May 2024), the German Federal Government, Bavarian state administration, SAP, Bosch, BMW, Siemens, and Schwarz Group. Customer Entrenchment at 5 - the German federal and Bundeswehr footprint is genuinely entrenched but the customer count is small and concentrated. Knowledge & Data Advantage at 4 - Aleph Alpha publishes the only training-FLOPs and token-budget figures in the cohort for the Pharia-1-LLM-7B family. Strategic Resilience is the cohort floor at 2 - Aleph Alpha has not publicly disclosed an ARR or revenue figure, German-language coverage characterises the company as having modest revenue and high burn relative to the November 2023 EUR 500m structured raise, Jonas Andrusis stepped down as CEO in October 2025 and confirmed to NZZ in early 2026 "I'm out," and approximately 50 staff were cut in January / February 2026.


Disruption Potential.

Momentum is the cohort floor at 3 - the layoffs, the leadership departures and the pending acquisition are the load-bearing evidence that the standalone trajectory was insufficient. Novel Capability at 5 - the PhariaAI sovereign-AI platform is a credible enterprise OS for governance and on-premises deployment, but Aleph Alpha has explicitly stopped competing on frontier-model performance since mid-2024. Team Velocity at 4. Category-Shaping Signal at 4 - Aleph Alpha was Germany's clearest "European national-champion AI lab" candidate in 2023, and the strategic pivot since 2024 plus the pending acquisition mean the standalone signal has been reset.

Trajectory commentary.

Best-positioned for Borders at 7 - Bundeswehr classified-information clearance, the STACKIT sovereign cloud runtime, GDPR / EU AI Act compliance, on-premises and air-gapped deployment - the position is robust to regulatory fragmentation in a way the standalone Trajectory profile otherwise is not. Worst-positioned for Frontier at 2, the cohort's lowest single Frontier score. Worst-positioned also for Expensive Compute at 3. The pending Cohere acquisition (all-stock at \$20bn combined entity, Aleph Alpha shareholders approximately 10% post-close, \$600m Series E from Schwarz Group as anchor) and the regulatory close path (German Aussenwirtschaftsverordnung; Canadian Investment Canada Act) are the structural resolution this profile is otherwise read as pre-empting. German press has framed the deal as "Germany's AI hope sold to Canada" (heise, FAZ, Handelsblatt).

6.9 Aleph Alpha (continued)



Eight scenario Trajectory per IM Framework Macro Briefing #IM105 Eight Futures for Agentic AI. Darker green = higher score. Source: IM scoring grid.

Figure 3.9. Aleph Alpha - Trajectory Profile across eight scenarios (green gradient).

7. CROSS-COHORT PATTERNS

Four patterns emerge from reading the cohort's scoring grid in aggregate rather than vendor by vendor.

- **The tier gap is the story.** OpenAI, Anthropic, Google DeepMind and Meta have moved structurally into a category of their own. Combined disclosed run-rate revenue is in the multiple-tens-of-billions; combined committed forward compute spend exceeds \$1 trillion through 2032; each carries hyperscaler-grade or parent-hyperscaler balance-sheet support. The remaining five cohort vendors operate at structurally smaller scale on every disclosed dimension. Mistral's reported \$400m ARR and Cohere's \$240m FY25 ARR are within an order of magnitude of each other; both are two orders of magnitude below the top tier.
- **Openness is where the European story now lives.** With the Cohere-Aleph Alpha merger collapsing the European frontier-lab cohort to Mistral plus a transatlantic combined entity, the "Europe versus the US" framing has narrowed sharply. Where European cohort members continue to differentiate is on sovereignty and openness rather than on capability frontier. Open-weight licensing is not the cohort-internal differentiator any more - Llama and DeepSeek dominate the open-weights frontier, and the strongest open-weight models on rigorous benchmarks are Chinese-origin. Sovereign-deployment capability is what the European cohort members have built that the US Big Three cannot match.
- **Compute supply is the binding constraint, not capability.** The cohort's top tier has shifted from a capability-led story to a supply-led story (which lab can secure the compute to ship the next model at scale). Three top-tier vendors have publicly signalled compute-constrained operations in Q1 2026. NVIDIA's Data Center segment at \$75.2bn in Q1 FY27 (+92% YoY), networking at \$14.8bn (+199% YoY) and a \$91bn +/-2% revenue guide for Q2 FY27 are the supply side of the same equation. The 20 May 2026 print confirmed the supply-led reading: roughly \$10bn of additional revenue is being absorbed into the cohort's compute pipeline next quarter.
- **Disclosure-quality variance is asymmetric and has scoring consequences.** The cohort splits sharply on disclosure quality: audited 10-K filers at one end (Alphabet, Meta), company-stated press-release disclosure in the middle (OpenAI, Anthropic), company-stated investor-memo disclosure at the European mid-tier (Mistral, Cohere), audited-via-parent-S-1 since February 2026 (xAI through the SpaceX IPO filing), and undisclosed at the bottom (DeepSeek revenue; Aleph Alpha revenue). The IM Framework's Strategic Resilience scoring weights this variance explicitly. For a buyer making a multi-year integration commitment, the disclosure-quality cell of the scoring grid is the one that deserves the closest read.

8. INDICATORS DASHBOARD

The signals below are the empirical things to watch over the next twelve months. Each one is a piece of public information a reader can go and check. The signals are grouped under four plain-language reader-facing questions rather than under the framework's internal construct names.

How fast is the cohort top tier moving on compute?

- **NVIDIA Data Center revenue quarterly disclosures.** The Q1 FY27 print (20 May 2026) delivered \$75.2bn in Data Center revenue (+92% YoY) and \$14.8bn in networking (+199% YoY) within \$81.6bn of total revenue. Q2 FY27 guide is \$91bn +/-2%. Hyperscalers continue to account for just over 50% of Data Center revenue. Any sequential deceleration off the \$91bn guide would be a material cohort-level signal.
- **AMD Data Center revenue quarterly disclosures.** Next quarterly read in late July 2026; Q1 2026 delivered \$5.8bn (+57% YoY); the cohort-internal MI450 series 6 GW Meta deployment via the Helios rack-scale architecture is the most consequential named customer-win.
- **Hyperscaler capex guidance revisions on each earnings cycle.** Alphabet at \$180-190bn FY2026 with "significantly increase" 2027 guidance; Microsoft at \$190bn calendar 2026; Amazon at approximately \$200bn 2026; Meta at \$125-145bn FY2026; Oracle at \$50bn FY2026. Any single-vendor revision of 10% or more is a material cohort-level signal.
- **Public statements on compute-constrained operations.** Anthropic on Claude Code rate-limit policy; Alphabet on Cloud demand-supply gap; OpenAI on the Stargate-Oracle-Azure forward-commitment ratio; Meta on the +\$107bn quarter-over-quarter non-cancellable commitments line.
- **Specific named-cluster announcements.** Stargate site openings (UAE, Norway, UK, Argentina, the reported India / TCS site); xAI Colossus 2 capacity additions; Anthropic Project Rainier scaling above the current approximately 500,000 Trainium2 chips; Meta Louisiana data centre campus delivery dates; Mistral Bruyeres-le-Chatel commissioning expected Q2 2026.

How fast is the cohort middle tier funding itself?

- **Mistral Series D or pre-IPO round timing.** The \$830m debt raise for Bruyeres-le-Chatel (March 2026), the 1.4 GW Paris-region AI-campus joint venture with Bpifrance, MGX and Nvidia, and the Mistral Compute build-out together imply a Series D or pre-IPO round in 2026-27. A round announcement is the most direct read on whether the European sovereign-AI thesis is being funded at scale.
- **Cohere IPO timing.** Aidan Gomez has publicly flagged "soon" since October 2025 and the February 2026 investor memo reads as IPO-preparation messaging. The Aleph Alpha acquisition close (subject to German FDI and Canadian Investment Canada Act review) is the likely precondition.
- **SpaceX IPO timing and xAI segment disclosure.** With the 2 February 2026 SpaceX-xAI close completing the acquisition, the standalone xAI funding-round cadence is no longer the relevant forward signal. The SpaceX IPO timing (Nasdaq ticker SPCX targeted; reported \$1.75 trillion target valuation), any subsequent SpaceX 10-Q or 10-K segment disclosure naming xAI capital intensity, and any disclosure surfacing the Memphis-Southaven environmental litigation costs are now the live cohort-level signals.
- **DeepSeek any external priced round.** Would break the High-Flyer-only funding pattern and be a structural signal that the founder-and-parent-funded model has reached its scaling ceiling.
- **The Cohere-Aleph Alpha close.** The German Aussenwirtschaftsverordnung clearance and the Canadian Investment Canada Act review are the regulatory gates. Either rejection (genuinely possible in the German case given how the deal has been framed in domestic press) or unconditional clearance is a material cohort signal.

How is the agentic share of foundation-model revenue evolving?

- **Anthropic Claude Code ARR.** Last disclosed at above \$2.5bn at Series G (February 2026); next disclosure expected at Anthropic's next funding round. The single cleanest agentic-SKU revenue datapoint in the category.
- **OpenAI Codex usage metrics.** OpenAI disclosed "more than 2 million weekly users with monthly growth of 70%" at last update; future disclosure cadence is the key signal on whether the agentic-coding segment continues to bifurcate (Claude Code versus Codex versus Cursor).
- **Salesforce Agentforce telemetry.** 771 million Agentic Work Units in Q4 FY2026 (+57% QoQ); 19 trillion tokens processed cumulatively (+5x YoY); Agentforce ARR at \$800m (+169% YoY); 29,000 deals closed (+50% QoQ). The AWU growth rate is the more useful directional indicator than the ARR-as-share-of-Salesforce-total (approximately 2%).
- **The a16z 100-CIO survey quarterly waves.** 2025 wave delivered "enterprise AI spend is growing 75% year-over-year, with innovation budget allocation dropping from 25% to just 7% of total AI spend" and "44% of enterprises now using Anthropic in production, rising to over 63% including testing." The 2026 wave (expected H2 2026) is the next reader-facing read.

- **Anthropic May 2026 agent-billing API rails.** Anthropic's pricing of agentic API usage as a separate line will provide cleaner attribution data within 6-9 months - the cleanest data-architectural change in the cohort for the question of what genuinely-agentic foundation-model revenue actually is.

Where are the regulator pressure points landing?

- **EU AI Act enforcement actions from 2 August 2026 onward.** Twenty-six organisations signed the GPAI Code of Practice including Amazon, Anthropic, Google, IBM, Microsoft, OpenAI, Aleph Alpha, Cohere and Mistral. Meta declined to sign. xAI signed only the safety-and-security chapter. Code signatories gain regulatory predictability under the 2 August 2026 application of enforcement powers; Meta and xAI sit relatively exposed.
- **Next BIS export-control update on advanced chips.** The final rule effective 15 January 2026 revised license review policy for advanced computing semiconductors to China; the Trump December 2025 announcement on H200 exports "in exchange for 25% of the revenues being paid to the US Government" is the most idiosyncratic policy in this space. Further BIS action affecting DeepSeek's chip access (or NVIDIA's China revenue base) is a top-tier cohort signal.
- **CAISI pre-deployment evaluation findings publications.** CAISI agreements with Google DeepMind, Microsoft and xAI announced 5 May 2026 add to existing OpenAI and Anthropic arrangements; CAISI has "completed over 40 evaluations of unreleased models" as of May 2026. Publication of any evaluation finding that materially shifts the public read on a cohort model's capability or safety profile is a high-leverage signal.
- **UK AISI joint-evaluation publications.** The UK and US AISIs conducted joint pre-deployment evaluations of OpenAI o1 and Anthropic Claude 3.5 Sonnet (announced late 2024); the new UK-Anthropic and UK-DeepMind agreements in 2025 set up further joint-evaluation cadence.
- **CAC enforcement actions on Chinese-filed services.** As of December 2025, 748 generative AI services had completed national-level filing including DeepSeek and Baidu's Ernie Bot. CAC 2025 enforcement focus is labeling of AI-generated content; mandatory national standard GB 45438-2025 effective 1 September 2025.
- **xAI Colossus 2 litigation rulings.** Earthjustice emergency injunction filed May 2026 over 46 unpermitted gas turbines; the most concrete near-term regulatory-litigation item in the cohort. A ruling that requires turbine shutdown or permit-compliance retrofitting would have measurable cluster-capacity consequences.

Material moves between now and the Q4 2026 refresh will be tracked in the Information Matters Brief (our weekly Substack digest) and in the quarterly Macro Briefings.

9. WATCHLIST

The Watchlist names vendors active in the foundation-model-providers segment that are not part of the formal scoring in this Category Report. Inclusion is editorial - these are vendors IM tracks and may move into the formal cohort in a future refresh, but where current visibility or scale does not yet warrant the full per-vendor framework treatment.

- **Open-weight labs outside the formal cohort.** Alibaba (Qwen / Tongyi Qianwen, Hangzhou, China) - Qwen3 (April 2025; Apache 2.0; 36 trillion training tokens; 119 languages), Qwen3.5 (February 2026), Qwen3.6-27B (April 2026); strongest open-weight benchmark performance on GPQA Diamond and AIME '24 as of May 2026. Zhipu / GLM (Beijing) - GLM-5 leads SWE-bench Verified at 77.8% for open-weights. Allen AI / AI2 (Seattle) - OLMo 3 (November 2025) plus Tulu 3 405B (January 2025); the strongest fully-open release of 2025. MiniMax (Shanghai) and Z.ai (Beijing / Hangzhou) - named in Epoch AI's frontier-cost open-disclosure data points. Sakana AI (Tokyo) - evolutionary and model-merging research lab.
- **Hyperscaler-internal model labs.** Amazon Nova / Olympus (Seattle) - Amazon's in-house Nova family for AWS Bedrock; not consumer-facing today but a Bedrock-distribution play. Apple Intelligence / Apple Foundation Models (Cupertino) - Apple's on-device and Private Cloud Compute models; not API-distributed; integrated into iOS and macOS.
- **Specialist or vertical foundation-model labs.** Black Forest Labs (Freiburg, Germany) - FLUX image-generation family; specialist generative-media positioning. Stability AI (London) - Stable Diffusion family. AI21 Labs (Tel Aviv) - Jurassic / Jamba family; enterprise-focused. Reka AI (San Francisco / Singapore) - multimodal-first foundation models, enterprise focus.
- **Sovereign / national-champion labs.** Krutrim and Sarvam (Bengaluru) - India sovereign-AI champions. G42 (Abu Dhabi) - sovereign-AI investor and operator; Falcon model family via Technology Innovation Institute. Baidu Ernie, Tencent Hunyuan and ByteDance Doubao (China) - each a credible large-model entity within an ecosystem (search, social, content) that the agentic foundation-model frame does not cleanly cover at company level.

The Watchlist will be refreshed alongside the next Category Report. Vendors that meaningfully change scale, funding posture, or product trajectory in the briefing window will move into the formal cohort.

10. COMPANIES MENTIONED IN THIS REPORT

Legend: * denotes a formally-scored cohort vendor (one of the nine per-vendor profiles in section 6). (W) denotes a Watchlist vendor (named in section 9 but not formally scored in this report). Unmarked entries are supporting players named in the narrative without formal cohort or Watchlist status.

- **AI2 (Allen Institute for AI) (W) - United States.** Independent research institute running the fully-open OLMo model-release programme including OLMo 3 and Tulu 3 405B.
- **AI21 Labs (W) - Israel.** Enterprise-focused foundation-model lab; Jurassic and Jamba families.
- **Aleph Alpha * - Germany.** Sovereign AI platform built around PhariaAI on the STACKIT cloud; subject of all-stock acquisition announced by Cohere on 24 April 2026, not closed as of 16 May 2026.
- **Alibaba (Qwen) (W) - China.** Publicly-listed e-commerce and cloud group; ships the Qwen / Tongyi Qianwen open-weight foundation-model family including Qwen3, Qwen3.5 and Qwen3.6.
- **Alphabet - United States.** Publicly-listed (NASDAQ: GOOGL / GOOG); parent of Google DeepMind, Google Cloud, Search and the Workspace family.
- **AMD - United States.** Publicly-listed data-centre GPU vendor; MI300X, MI325X and MI450 series; named customer-win of the 6 GW Meta deployment on the Helios rack-scale architecture.
- **Amazon (Nova) (W) - United States.** Publicly-listed retailer-and-cloud parent; ships the Nova / Olympus first-party foundation-model family for AWS Bedrock; cumulative \$13bn direct investment in Anthropic plus the Project Rainier compute partnership.
- **Anthropic * - United States.** Foundation-model lab; ships the Claude family (Opus 4.8 launched 28 May 2026, Sonnet 4.6, Haiku 4.5) and Claude Code; \$47bn stated run-rate revenue May 2026; \$65bn May 2026 funding round at \$965bn post-money (Altimeter-led; FT 28 May) surpasses OpenAI's \$852bn.
- **Apple Intelligence (W) - United States.** Publicly-listed consumer-electronics parent; ships on-device foundation models and the Private Cloud Compute backend.
- **ASML - Netherlands.** Publicly-listed semiconductor-lithography equipment maker and lead investor of Mistral's September 2025 Series C with approximately 11% stake.
- **Baidu (Ernie) (W) - China.** Publicly-listed search-and-content giant; ships the Ernie Bot foundation-model family.
- **Black Forest Labs (W) - Germany.** FLUX generative-image-model family; specialist generative-media foundation-model lab.
- **Bpifrance - France.** French sovereign investment fund; participated in Mistral Series C and the 1.4 GW Paris-region AI-campus joint venture.
- **Broadcom - United States.** Publicly-listed semiconductor and infrastructure-software vendor; counterparty in the April 2026 Anthropic 3.5 GW TPU-based compute deal.
- **ByteDance (Doubao) (W) - China.** Privately-held social-and-content giant; ships the Doubao foundation-model family.
- **Cerebras Systems - United States.** Wafer-scale AI supercomputer hardware vendor; CS-3 deployed at Aleph Alpha's alpha ONE data centre.

- **Cohere * - Canada.** Enterprise-only foundation-model lab; ships Command A and Cohere North; \$240m FY25 ARR; acquirer of Aleph Alpha (announced April 2026, pending close).
- **DeepSeek * - China.** Open-weight foundation-model lab; ships V3, R1 and V3.2 under MIT license; founder and CEO Liang Wenfeng; funded internally by High-Flyer Quantitative Investment Management.
- **Earthjustice - United States.** Public-interest environmental-law nonprofit; filed for emergency injunction against xAI Colossus 2 in May 2026.
- **G42 (W) - United Arab Emirates.** Sovereign-AI investor and operator; backs the Falcon foundation-model family via Technology Innovation Institute.
- **GitHub - United States.** Microsoft subsidiary acquired 2018; distribution surface for Anthropic and OpenAI coding models via the Copilot family.
- **GLM / Zhipu (W) - China.** Beijing-based foundation-model lab; GLM-5 leads SWE-bench Verified at 77.8% for open-weights.
- **Google DeepMind * - United Kingdom / United States.** Alphabet's foundation-model lab (merged London-based DeepMind with Google Brain in April 2023); ships the Gemini 3 family, the generative-media stack (Veo 3, Imagen 4, Lyria 2), and scientific systems (AlphaFold 3, AlphaProteo, GNoME).
- **High-Flyer Quantitative Investment Management - China.** Hangzhou-based quantitative hedge fund founded by Liang Wenfeng in 2016; sole funder of DeepSeek.
- **HSBC - United Kingdom.** Publicly-listed bank; named Mistral enterprise customer.
- **IBM (watsonx) - United States.** Publicly-listed enterprise-technology vendor; carries Mistral models on watsonx distribution.
- **Krutrim (W) - India.** National-champion sovereign-AI entity.
- **MGX - United Arab Emirates.** UAE state-sovereign technology investor; participated in OpenAI March 2026 round, Anthropic Series G, xAI Series E, and the Mistral 1.4 GW Paris-region AI-campus joint venture.
- **Meta * - United States.** Publicly-listed (NASDAQ: META) consumer-platform and advertising parent; ships the Llama family (Llama 4 Scout, Llama 4 Maverick) under the Llama 4 Community License and the closed-source Muse Spark via Meta Superintelligence Labs.
- **Microsoft - United States.** Publicly-listed (NASDAQ: MSFT) software-and-cloud parent; OpenAI partner with IP rights through 2032; counterparty to Anthropic Azure spend and Llama distribution on Azure AI Foundry.
- **MiniMax (W) - China.** Shanghai-based foundation-model lab; named in Epoch AI's frontier-cost open-disclosure data points.
- **Mistral AI * - France.** Independent foundation-model lab; ships Mistral Large 3, Mistral Medium 3, Mistral Small 3, Pixtral and Codestral; EUR 11.7bn post-money at September 2025 Series C.
- **NAACP - United States.** National Association for the Advancement of Colored People; litigating party against xAI and subsidiary MZX Tech over unpermitted gas turbines.
- **NVIDIA - United States.** Publicly-listed (NASDAQ: NVDA) data-centre GPU vendor and primary supplier to every cohort lab; Q1 FY27 Data Center revenue \$75.2bn (+92% YoY) within \$81.6bn of total revenue.
- **OpenAI * - United States.** Foundation-model lab structured as OpenAI Group PBC controlled by the OpenAI Foundation; ships GPT-5.x, ChatGPT, the Operator and Codex agentic products, AgentKit, and Sora.

- **Oracle - United States.** Publicly-listed (NYSE: ORCL) database-and-cloud parent; counterparty to OpenAI's \$300bn compute commitment and Cohere's primary distribution partner via OCI.
- **Reka AI (W) - United States / Singapore.** Multimodal-first enterprise foundation-model lab.
- **Royal Bank of Canada - Canada.** Publicly-listed (TSX: RY) bank; named Cohere enterprise customer.
- **Sakana AI (W) - Japan.** Tokyo-based evolutionary and model-merging research lab.
- **Salesforce - United States.** Publicly-listed (NYSE: CRM) enterprise-software parent; Agentforce telemetry (771m AWUs Q4 FY2026; \$800m ARR) cited as the cleanest agentic-SKU disclosure outside Anthropic Claude Code.
- **SAP - Germany.** Publicly-listed (FWB: SAP) enterprise-software parent; sovereign-AI stack partner with Mistral, the French and German governments, and the Aleph Alpha PhariaAI platform.
- **Sarvam (W) - India.** National-champion sovereign-AI entity.
- **Schwarz Group - Germany.** Privately-held retail-and-tech parent (Lidl, Kaufland, STACKIT); largest Aleph Alpha shareholder (>20%) and \$600m anchor of the Cohere Series E concurrent with the April 2026 acquisition.
- **SoftBank - Japan.** Publicly-listed technology investor; approximately \$30bn participation in the OpenAI March 2026 round and Stargate principal.
- **Stability AI (W) - United Kingdom.** Stable Diffusion family; specialist generative-image foundation-model lab.
- **STACKIT (Schwarz Digits) - Germany.** Schwarz Group's sovereign cloud platform; contractually mandated runtime for PhariaAI-as-a-Service.
- **Stellantis - Italy / France.** Publicly-listed automaker; named Mistral enterprise customer for in-car AI.
- **Tencent (Hunyuan) (W) - China.** Publicly-listed social-and-content giant; ships the Hunyuan foundation-model family.
- **xAI * - United States.** Foundation-model lab and parent of Grok; consolidated with X in March 2025; acquired by SpaceX 2 February 2026 in \$250bn all-stock merger and now a wholly-owned SpaceX subsidiary within a \$1.25 trillion combined entity; Colossus 1 and Colossus 2 the largest single named training clusters in the cohort. xAI 2025 standalone revenue disclosed at \$3.2bn against \$6.4bn of operating loss per the SpaceX S-1.

11. GLOSSARY OF TERMS

- **Agentic AI.** Software that takes multi-step actions to complete a task, calling tools and adapting to intermediate results, rather than answering single-turn questions.
- **API (Application Programming Interface).** The programmatic interface through which third-party applications consume a foundation-model provider's models; the primary revenue surface for OpenAI, Anthropic, Mistral, Cohere and DeepSeek's third-party-hosted variants.
- **ARR (Annual Recurring Revenue).** The revenue an annualised subscription book is generating at the time of measurement; the cohort's main public scale measure for subscription-driven vendors.
- **Borders.** One of the Eight Futures; the scenario in which regulatory and geopolitical divergence fragments the global market into geographic spheres.
- **Capex (Capital Expenditure).** Spending on long-lived productive assets - data centres, GPUs, networking, real estate, power infrastructure - that is capitalised on the balance sheet rather than expensed; the cohort top-tier vendors' capex guides for FY2026 are the load-bearing forward indicator of compute supply.
- **Category Report.** An IM publication type that profiles a defined competitive set of vendors with formal 1-10 framework scoring; this report is the second Category Report after #IM107 Coding Agents: The 2026 Landscape.
- **Category-Shaping Signal.** A Disruption Potential sub-dimension in the IM Framework; measures the degree to which a vendor's product or strategy is defining what the category becomes rather than competing inside it.
- **Cohort.** The defined set of vendors selected for scoring in this Category Report; this cohort is nine foundation-model providers.
- **Customer Entrenchment.** A Defensibility sub-dimension in the IM Framework; measures how deeply a vendor's product is embedded in customers' workflows and how costly it would be to switch away.
- **Defensibility.** One of two macro constructs in the IM Framework; the structural depth of a vendor's position against erosion. Has four sub-dimensions: Customer Entrenchment, Knowledge & Data Advantage, Distribution & Ecosystem, and Strategic Resilience.
- **Disruption Potential.** The second macro construct in the IM Framework; the forward-looking ability to shape and capture category direction. Has four sub-dimensions: Momentum, Novel Capability, Team Velocity, and Category-Shaping Signal.
- **Distribution & Ecosystem.** A Defensibility sub-dimension in the IM Framework; measures the depth and breadth of a vendor's channel reach, partner integrations, and platform position.
- **Eight Futures.** The eight scenarios in the IM Trajectory Profile framework: Frontier, Plateau, Verticals, Low-Cost Compute, Expensive Compute, Rewire, Inertia, and Borders. Set in the IM Macro Briefing Eight Futures for Agentic AI.
- **Expensive Compute.** One of the Eight Futures; the scenario in which inference economics tighten and the cost of running agentic workloads rises.
- **Foundation provider.** A company that builds and ships the underlying large-language-model layer on which other agentic AI products depend; the nine cohort vendors in this report.

- **Framework chart.** One of three visualisations the IM Framework reserves for Category Reports: the 2x2 quadrant chart (Defensibility x Disruption Potential), the Trajectory Profile heat-map (eight scenarios across all cohort vendors), and the per-vendor Trajectory bar chart.
- **Frontier.** One of the Eight Futures; the scenario in which foundation-model capability continues to leap forward generation by generation.
- **GPAI (General-Purpose AI).** The category of foundation models covered by the EU AI Act's Articles 53 and 55; GPAI obligations took effect 2 August 2025 and enforcement powers apply from 2 August 2026.
- **IDE (Integrated Development Environment).** The software application a developer writes code in - used in this report when describing distribution surfaces for Anthropic Claude Code, GitHub Copilot and similar agentic-coding products.
- **Inertia.** One of the Eight Futures; the scenario in which the technology arrives but enterprise deployment lags well below forecast.
- **Knowledge & Data Advantage.** A Defensibility sub-dimension in the IM Framework; measures whether a vendor owns a proprietary data asset (corpus, embeddings, fine-tuned model, or unique customer-data network) that other vendors must license or substitute for.
- **Low-Cost Compute.** One of the Eight Futures; the scenario in which open-source models close on closed-source, inference prices collapse, and on-device inference scales.
- **Macro Briefing.** An IM publication type that sets the scenario landscape qualitatively across the agentic AI sector; Eight Futures for Agentic AI is the canonical reference for the eight Trajectory scenarios used in this report.
- **MCP (Model Context Protocol).** Anthropic's emerging standard for connecting AI models to tools and data sources; opened to community 2024-2025 and now a category-wide tool-use rail.
- **MoE (Mixture of Experts).** A model architecture in which only a subset of total parameters is active per inference call; used by Llama 4 Scout / Maverick, DeepSeek V3 / V3.2, Mistral Large 3 and others to expand parameter count without proportionally increasing inference cost.
- **Momentum.** A Disruption Potential sub-dimension in the IM Framework; measures the trajectory of revenue or adoption - how quickly a vendor is growing in observable scale terms.
- **Novel Capability.** A Disruption Potential sub-dimension in the IM Framework; measures whether a vendor's product does something category-defining that competitors do not currently match.
- **Plateau.** One of the Eight Futures; the scenario in which foundation-model output commoditises and capability progress slows.
- **Rewire.** One of the Eight Futures; the scenario in which enterprises restructure their organisations and workflows around agent-mediated operation.
- **Robustness statistic.** The single number per vendor that captures performance across the combined future rather than in any one scenario; calculated as mean of the eight Trajectory scores minus half the standard deviation across them. In this report, robustness is the buyer's optimisation target.

- **RPO (Remaining Performance Obligations).** The contracted-but-not-yet-recognised revenue figure disclosed by cloud and software vendors; Alphabet's Cloud RPO at \$460bn at Q1 2026 and Oracle's \$553bn at Q3 FY2026 are the load-bearing RPO disclosures referenced in this report.
- **SDK (Software Development Kit).** A bundled set of tools, libraries and documentation a software developer uses to build on top of a platform; Anthropic's Claude Agent SDK is the relevant cohort reference.
- **Sector Context Briefing (SCB).** An IM publication type that covers a whole industry vertical qualitatively, without per-vendor 1-10 scoring; distinct from the Category Report format used here.
- **Strategic Resilience.** A Defensibility sub-dimension in the IM Framework; measures a vendor's capital base, parent-company strength, disclosure-quality variance, and ability to survive a downside scenario.
- **Team Velocity.** A Disruption Potential sub-dimension in the IM Framework; measures engineering and execution quality - how fast the vendor's team can ship, iterate, and improve.
- **Trajectory Profile.** A vendor's vector of eight scores, one per Eight Futures scenario; the central forward-looking output of the IM Framework.
- **Verticals.** One of the Eight Futures; the scenario in which sector specialists with proprietary domain data outperform horizontal generalists.
- **Watchlist.** A list of cohort-adjacent vendors named in the report but not formally scored; the Watchlist names the broader population the cohort is drawn from, and vendors may move into the formal cohort at the next Category Report refresh.

12. ABOUT INFORMATION MATTERS

Information Matters is an independent analyst publication covering the agentic AI sector. The Information Matters Brief is our weekly Substack.

Methodology note

This report was produced through a combination of human expertise and oversight supported by an AI research agent specifically designed to carry out detailed market research and forecasts using established methodologies and with access to the latest data and AI models. This research consists of the opinions of Information Matters' research team, human and AI, and the information contained within it should not be considered as statements of fact. None of the information presented here should be taken as investment advice. Reproduction or distribution of this research without written permission from Information Matters Ltd is prohibited. (c) 2026 Information Matters Ltd. All rights reserved.