INFORMATION MATTERS

MACRO BRIEFING // EIGHT FUTURES

Eight Futures for Agentic AI

Eight scenarios for where agentic AI is heading over the next three to five years — the case for each, what to watch, what would falsify them, and the implications for enterprise buyers and acquirers.

May 2026

Martin De Saulles | Principal Analyst

Information Matters | May 2026

Report #IM105

Headline findings

- **Eight plausible scenarios, no single forecast.** We score every vendor against eight named futures rather than betting on one. The vector exposes vendors fragile to specific scenarios that single-future analysis would miss.
- **Plateau and Low-Cost Compute lead our weighting.** Open-source closing on closed-source, inference price collapse, and on-device inference scaling all point to commoditisation as the dominant near-term force.
- **Inertia is the under-discussed scenario.** Every previous enterprise tech wave saw deployment lag the technology; we expect this one to be no different, with major implications for venture-backed challengers.
- **Verticals strengthens steadily.** Sector specialists in healthcare, legal, and financial services are outperforming generalists where it matters; sector-specific data partnerships become the most valuable assets.
- **Borders is real and increasingly material.** EU AI Act enforcement begins August 2026; US-China divergence persists; the global market fragments into geographic spheres.
- **Robustness across all eight is what enterprise buyers should optimise for.** A vendor that scores well in one scenario but poorly in another is a fragile bet over a five-year integration commitment.

EXECUTIVE SUMMARY

Eight plausible futures, scored together

Most analysis of agentic AI implicitly assumes one future — capabilities will keep leaping, or commoditisation will eat the model layer, or regulatory divergence will fragment the market. None of these is provably correct in 2026. All have plausible evidence behind them. Several could partially happen at once.

Information Matters' framework names eight scenarios — Frontier, Plateau, Verticals, Low-Cost Compute, Expensive Compute, Rewire, Inertia, and Borders — and develops each as a coherent future for agentic AI over the next three to five years. Vendors are scored against each scenario, producing a vector that exposes fragility a single-future analysis would miss.

What we believe and what we're watching

Our weighted view as of May 2026: Plateau and Low-Cost Compute lead, with on-device inference reinforcing both. Inertia is the under-discussed scenario most likely to compress aggressive challengers' growth assumptions. Verticals strengthens steadily as sector specialists prove their data advantage. Borders becomes increasingly material as EU AI Act enforcement bites. Rewire happens but unevenly — fast in tech-native sectors, slow elsewhere. Frontier remains a force in narrow capability domains. Expensive Compute is a tail risk we treat seriously.

The interesting strategic implication: vendors well-positioned across multiple scenarios are robust precisely because the combined future favours none in particular. Robustness is what enterprise buyers making multi-year integration commitments should be optimising for.

> *Robustness across all eight scenarios is what an enterprise buyer making a five-year integration commitment is actually optimising for, even when they don't articulate it that way.*

1. FRONTIER

Capabilities keep leaping; foundation-model providers extract the value.

PREMISE

Foundation-model capabilities continue their current trajectory. Each major release every 12 to 18 months unlocks new capabilities — longer context, deeper reasoning, more reliable tool use, more capable computer use. The gap between the best closed-source frontier model and the rest of the field remains meaningful and is not being closed.

THE CASE

Recent capability releases have been substantial rather than incremental: o1 and o3 opened new use cases in reasoning that did not exist commercially in 2024. Inference-time compute opens a separate scaling axis that frontier providers are still climbing. Anthropic, OpenAI, and Google DeepMind hold the bulk of frontier-API revenue and continue to attract the strongest research talent.

DYNAMICS — WINNERS AND LOSERS

Foundation-model providers extract most of the new value. Wrapper companies that depended on a particular generation of model capability find their distinctive value compressed when the frontier moves past them. Companies with deep proprietary moats outside the model layer (Glean, Salesforce, ServiceNow) can still thrive but must continually re-justify why their layer is the right place for value to accumulate.

TWELVE-MONTH INDICATORS TO WATCH

- GPT-5, Claude 5, Gemini 3-class models shipping with capabilities that exceed current best on agentic-task benchmarks.
- Computer-use reliability crossing >80% on multi-step web-task benchmarks.
- Foundation providers shipping vertical agents that compete directly with wrapper companies.
- Closed-source vs open-source gap on commercially-relevant benchmarks remaining stable or widening.

WHAT WOULD FALSIFY THIS

- Best open-source models reaching parity for two consecutive release cycles.
- A major frontier release that materially underwhelms.
- Foundation-provider growth slowing below 50% YoY for two consecutive quarters.

IMPLICATIONS

Buyers should be cautious about wrapper vendors in fast-moving categories. Investors should weight platform plays over wrapper plays. Vendors building on top should invest aggressively in non-model moats — proprietary data, customer entrenchment, regulatory standing, distribution.

2. PLATEAU

Capability levels off; commoditisation wins; non-model moats become the differentiator.

PREMISE

Frontier capability growth slows materially. Open-source models reach "good enough" parity for the bulk of enterprise use cases. Inference prices drop by an order of magnitude as competition intensifies. The model layer becomes infrastructure, like cloud compute became infrastructure a decade ago.

THE CASE

Open-source has closed substantial ground in 2025–26. Llama 3.1, DeepSeek-V3, Qwen 2.5, and Mistral's recent releases are within touching distance of closed-source frontier. Inference prices for capable cloud models in mid-2026 are roughly 10x lower than mid-2024. On-device inference reinforces the trajectory: Apple Intelligence, Gemini Nano, and Microsoft Copilot+ PCs put capable small models on consumer hardware at near-zero marginal cost. Many enterprise use cases work materially well at current capability — "good enough" is the operative phrase.

DYNAMICS — WINNERS AND LOSERS

Wrapper risk inverts. Companies that built thin layers over expensive proprietary models do well because their input cost collapses while their value-added holds steady. Differentiation shifts to what the model itself doesn't provide: data, workflow, distribution, trust, regulatory standing, customer relationships. The companies on Information Matters' defensible list — Salesforce, Glean, Sierra, ServiceNow — are best positioned for Plateau.

TWELVE-MONTH INDICATORS TO WATCH

- Inference prices continuing to drop >5x per year through 2027.
- Open-source models matching closed-source on commercially-relevant benchmarks.
- On-device inference share growing materially across consumer devices.
- Foundation-provider revenue growth slowing below 50% YoY.

WHAT WOULD FALSIFY THIS

- A major reasoning or capability breakthrough that opens new use cases.
- Closed-source vs open-source gap widening rather than narrowing.
- On-device inference failing to find use cases at scale despite hardware capability.

IMPLICATIONS

Buyers should weight non-model moats more heavily. Buyers in privacy-sensitive verticals should evaluate on-device options seriously. Investors should reweight away from foundation-model purebloods toward companies with non-model defensibility and toward hardware vendors capturing on-device value.

3. VERTICALS

Specialisation beats generalisation; sector-specific stacks dominate.

PREMISE

Generic frontier models stop being good enough for serious sector-specific work. Specialised models — trained on or deeply tuned for legal, medical, financial, code, or scientific data — outperform generalist models on the work that matters in each sector. The field fragments into specialised stacks. Sector-specific data partnerships become the most valuable assets in agentic AI.

THE CASE

Real enterprise work is irreducibly domain-specific. We see early evidence in 2026 that specialisation pays: Abridge and Suki in healthcare; Harvey and EvenUp in law; Hippocratic in clinical workflow are outperforming generalist agents on the work that defines their sectors. The data advantage is real — clinical conversations, legal precedent, transactional flow data are sector-specific moats. Verticals strengthens under either Frontier or Plateau, just through different mechanisms.

DYNAMICS — WINNERS AND LOSERS

Vertical AI vendors thrive. Horizontal generalists must either pick a vertical to dominate or accept being the second-best option in everyone's category. Sector-specific data partnerships (LexisNexis, Westlaw, Epic, Cerner, Bloomberg, Refinitiv) become the most valuable assets in the field.

TWELVE-MONTH INDICATORS TO WATCH

- Vertical AI vendors winning enterprise deployments against horizontal generalists.
- Sector-specific data-partnership announcements (exclusive arrangements with Epic, Westlaw, Bloomberg).
- Vertical fine-tuned models outperforming generalist models on sector-specific benchmarks.
- Horizontal vendors announcing vertical-specific product lines as a defensive response.

WHAT WOULD FALSIFY THIS

- A generalist model that outperforms vertical specialists on sector-specific tasks.
- Vertical specialists getting acquired by horizontal generalists at consolidating pace.
- Sector-specific data partnerships proving non-exclusive or commodity-priced.

IMPLICATIONS

Buyers should evaluate vertical specialists alongside horizontal generalists in sector-specific procurement. Investors should weight vertical specialists with strong sector-specific data partnerships highly. Vendors should choose between deep vertical specialisation and clear horizontal positioning; the middle is the weakest place to be.

4. LOW-COST COMPUTE

Inference collapses; on-device scales; high-volume thin-margin businesses thrive.

PREMISE

Compute economics shift dramatically downward. Cloud inference drops 10x as Chinese silicon, ARM-based inference chips, NVIDIA and AMD generational gains, specialty AI accelerators (Groq, Cerebras, SambaNova), distillation breakthroughs, and energy-efficiency improvements compound. On-device inference scales materially as Apple Silicon, Qualcomm Snapdragon X, and AMD AI accelerators put capable models on consumer hardware. Jevon's paradox: total volume scales faster than per-unit price drops.

THE CASE

Inference prices for capable cloud models in mid-2026 are roughly 10x lower than mid-2024 with no flattening. Specialty hardware is producing genuinely lower-cost inference for specific workloads. On-device inference is the most aggressive form of cost reduction — when capable models run on hardware the user already owns, the marginal cost is near zero. Distillation makes small models commercially viable on consumer devices.

DYNAMICS — WINNERS AND LOSERS

High-volume thin-margin businesses become enormous. Consumer agentic AI becomes a real category. Hyperscalers benefit on the cloud side; hardware vendors (Apple, Qualcomm, ARM, AMD) capture more value on the device side. Privacy-preserving on-device applications flourish. Foundation-model providers face price pressure but volume compensation.

TWELVE-MONTH INDICATORS TO WATCH

- Cloud inference price per million tokens continuing to drop >5x per year.
- Specialty hardware capturing >10% of measurable inference market.
- On-device inference share growing across Apple Intelligence, Gemini Nano, Copilot+ PCs.
- Consumer agentic products reaching >100M MAU at sustainable unit economics.

WHAT WOULD FALSIFY THIS

- Inference prices stabilising or rising for two consecutive quarters.
- Energy or chip-supply constraints biting (pushing toward Expensive Compute).
- On-device inference failing to find use cases at scale.

IMPLICATIONS

Buyers benefit from cheaper inference and should expect their bills to fall over time. Buyers in privacy-sensitive sectors should evaluate on-device options. Investors should weight scale plays heavily and consider hardware vendors. Vendors should design for scale — consumption-based pricing, scalable infrastructure, distribution channels that monetise volume.

5. EXPENSIVE COMPUTE

Energy and chip constraints bite; capital-rich incumbents win.

PREMISE

Compute economics shift dramatically upward. Energy constraints intensify as data-centre capacity hits grid limits. Geopolitical chip restrictions tighten. GPU shortages persist. Inference per token rises rather than falls. Compute becomes the operative bottleneck on agentic-AI deployment.

THE CASE

Energy is the binding constraint that under-discussed analyses have been pointing at. Data centres are increasingly limited by grid capacity rather than capital. Microsoft-Constellation, Google's geothermal, and dedicated power-purchase agreements are responses to a real bottleneck. Geopolitical chip restrictions are tightening, not loosening. Hyperscaler capex is at unprecedented levels — a sign of demand outstripping current supply.

DYNAMICS — WINNERS AND LOSERS

Capital-rich incumbents win. Hyperscalers benefit. Anthropic and OpenAI absorb cost rises through hyperscaler partnerships. Thin-margin and pre-revenue agentic vendors face existential pressure on burn rate. Free-tier consumer agentic products either monetise sharply or shut down. Enterprise-contracted vendors with strong gross margins absorb cost rises more comfortably than consumption-priced or free-tier products.

TWELVE-MONTH INDICATORS TO WATCH

- Inference prices stabilising or rising across major providers.
- New nuclear, geothermal, or other firm-power agreements signed by AI providers.
- Geopolitical chip restrictions tightening (new export controls, EU chip act).
- Free-tier or thin-margin agentic products announcing pivots, layoffs, or closures.

WHAT WOULD FALSIFY THIS

- Inference prices resuming their >5x-per-year decline through 2027.
- Energy constraints visibly easing as new firm-power comes online.
- Specialty hardware breakthroughs that bring inference cost down dramatically.

IMPLICATIONS

Buyers should pay attention to vendor compute strategy. Investors should weight balance-sheet strength heavily and be cautious of free-tier plays without monetisation paths. Vendors should ensure their pricing model can absorb compute inflation and consider strategic compute-supply arrangements.

6. REWIRE

Agentic AI restructures organisations; machines become customers.

PREMISE

Agentic AI begins to restructure how organisations operate, on a meaningfully shorter timeline than enterprise software waves typically take. SaaS categories absorbed into orchestration layers. New categories emerge in which the buyer is itself an agent — the machine-as-customer premise. Procurement decisions mediated through automated systems rather than human evaluation cycles.

THE CASE

Salesforce has restructured significantly around agentic AI deployment. ServiceNow is making similar moves. Several large enterprises are publicly piloting agent-mediated procurement. Internal corporate restructuring around agentic capability is beginning. The 2010s cloud-restructuring wave took most of the decade to fully manifest; agentic AI is at a similar inflection point and the question is cadence rather than direction.

DYNAMICS — WINNERS AND LOSERS

Vendors whose products were designed for human procurement struggle. Agent-native or agent-friendly vendors thrive. Orchestration layers eat individual SaaS tools from above. Brand and demand-generation models tuned for human attention lose effectiveness; programmatic discoverability and machine-readable APIs become the new distribution.

TWELVE-MONTH INDICATORS TO WATCH

- Major enterprise announcements of organisational restructuring with headcount implications.
- Documented cases of agent-mediated procurement at material scale.
- B2B SaaS revenue patterns shifting toward consumption-based and API-first revenue mix.
- Major SaaS vendors announcing agent-native architecture rewrites.

WHAT WOULD FALSIFY THIS

- Two consecutive years of unchanged enterprise organisational structures (Inertia wins).
- Agent-mediated procurement remaining experimental without enterprise-wide adoption.
- B2B SaaS revenue growing at historical rates with unchanged buyer-side composition.

IMPLICATIONS

Buyers should evaluate vendors for agent-native architecture and API-first design. Investors should weight agent-native architectures and API-first companies. Vendors should design for both human and agent users with API-first as the primary surface. The longer a vendor stays UI-first, the harder the architectural retrofit becomes.

7. INERTIA

Agentic AI is real, but enterprise deployment lags; incumbents win.

PREMISE

Agentic AI is technically capable but enterprise deployment lags substantially. Change management, governance, training, compliance, and integration with existing systems all take longer than tech-enthusiast predictions assume. Three to five years out, agentic AI is in production at fewer enterprises and at more limited depth than 2026 valuations and timelines imply. The slow-adoption counterpart to Rewire.

THE CASE

Every previous enterprise tech wave had this dynamic. Cloud was forecast to displace on-premise within five years and took fifteen. ERP rollouts took five years on average against eighteen-month forecasts. Even SaaS adoption was slower than initial projections. Recent agentic AI surveys show pilot-to-production conversion rates remaining low. Governance, compliance, and integration are the actual bottlenecks, not capability. Change management for autonomous agents is structurally harder than for tools that augment humans.

DYNAMICS — WINNERS AND LOSERS

Incumbents win because customers don't switch as fast as the technology would allow. Salesforce, ServiceNow, Microsoft, Oracle, SAP all benefit. Aggressive challengers face cash-flow pressure — startups whose burn rates assumed faster adoption hit cash crunches before the market catches up. Long-sales-cycle vendors with traditional enterprise relationships outperform. Patient capital wins. Customer entrenchment matters more under Inertia than under any other scenario.

TWELVE-MONTH INDICATORS TO WATCH

- Pilot-to-production conversion rates remaining low (consistently below 30% in major categories).
- Tier-1 SaaS incumbent churn rates staying flat or declining despite agentic challengers.
- Buy-side surveys reporting governance and compliance as primary bottlenecks (not capability).
- Aggressive 2025-era agentic challengers raising down rounds, announcing layoffs, or running out of cash.

WHAT WOULD FALSIFY THIS

- Rapid replacement of major SaaS categories within 24–36 months.
- Tier-1 SaaS incumbents losing material market share to agentic-native challengers.
- Pilot-to-production conversion rates accelerating sharply.

IMPLICATIONS

Buyers should expect their own agentic AI deployments to take longer than vendor claims and plan budgets accordingly. Trust existing vendor relationships during the agentic transition rather than betting on greenfield challengers. Investors should weight customer entrenchment, capital efficiency, and patient capital heavily. Vendors should ensure their cash position can sustain a longer-than-expected sales cycle.

8. BORDERS

Regulatory regimes diverge sharply; the global market fragments.

PREMISE

Regulatory approaches to AI fragment markedly across geographies. The EU AI Act enforcement tightens. The US oscillates between administrations on AI governance. China continues with its own regulatory framework. Cross-border data flows become more restricted; model deployment becomes regionalised; audit and compliance requirements diverge sharply by jurisdiction.

THE CASE

The EU AI Act is in force; enforcement of the General-Purpose AI provisions begins August 2026 with penalties up to 7% of global turnover. US export controls have tightened since 2022 and continue evolving. China's AI governance framework increasingly favours domestic vendors. The structural pattern follows GDPR's Balkanisation of cloud — faster, with higher stakes, and more geopolitical entanglement (chips, models, training data).

DYNAMICS — WINNERS AND LOSERS

Vendors with the resources to maintain three or more regional product lines thrive — Microsoft, Google, AWS at the top. Local champions in each major geography prosper within their borders: Mistral, Cohere, Aleph Alpha in Europe; DeepSeek, Baidu, Alibaba's Qwen in China. Mid-sized vendors with global ambition but without compliance budget for multiple regimes get squeezed. Buyers with operations across geographies face vendor-rationalisation pressure favouring multi-nationals.

TWELVE-MONTH INDICATORS TO WATCH

- EU AI Act enforcement actions producing material fines in 2026–27.
- US, EU, and Chinese regulatory frameworks diverging in observable ways.
- Regional data-residency commitments becoming standard for major AI vendors.
- Local-champion AI vendors gaining material market share within their geographies.

WHAT WOULD FALSIFY THIS

- EU AI Act enforcement remaining symbolic without material fines.
- Regulatory convergence: EU, US, UK frameworks aligning on common standards.
- Free flow of AI services across major geographies continuing without friction.

IMPLICATIONS

Buyers operating across geographies should weight multi-national AI vendors heavily, even at the cost of best-in-class capability. Investors should be cautious of mid-sized AI startups with global ambition but without capital and compliance infrastructure for multiple regimes; consider local-champion plays. Vendors should make explicit decisions about which regimes to operate in and invest seriously in compliance infrastructure.

CROSS-SCENARIO INTERACTIONS

Which scenarios reinforce; which are in tension

The eight scenarios are analytically separable but not mutually exclusive. Some combinations are mutually reinforcing; some are tensions that resolve only one way. Reading the framework requires understanding how the scenarios interact.

SCENARIOS THAT REINFORCE EACH OTHER

- **Plateau and Low-Cost Compute.** Capability commoditisation and falling inference prices are different aspects of the same broader dynamic. The combined scenario is what most enterprise buyers are implicitly preparing for.
- **Plateau and Verticals.** When the model layer commoditises, sector-specific data and domain expertise become the only remaining differentiators.
- **Inertia and Borders.** Both add deployment friction; regulatory complexity is one of the specific reasons enterprise adoption lags. Together they produce a world where incumbents win disproportionately.
- **Inertia and Plateau.** If technology commoditises but adoption lags, incumbents add cheap commoditised capability to existing products and challengers run out of cash before they catch market traction. The harshest combined scenario for venture-backed agentic challengers.

SCENARIOS IN TENSION

- **Frontier vs Plateau.** Direct tension. Partially compatible: pre-training scaling can plateau while inference-time-compute keeps climbing.
- **Low-Cost vs Expensive Compute.** Direct tension by construction. Cannot both dominate at the same time.
- **Rewire vs Inertia.** The central question of adoption pace. Mutually exclusive at the macro level; partial coexistence is possible in different sectors.

PLAUSIBLE COMBINED SCENARIO FOR 2027–28

Plateau partially obtains, Low-Cost Compute partially obtains, Inertia is widespread in regulated and process-heavy sectors while Rewire happens faster in tech-native sectors, Verticals strengthens steadily, Borders becomes increasingly material. Frontier remains a force in narrow domains. Expensive Compute is a tail risk that cannot be dismissed.

INFORMATION MATTERS' VIEW

Where our weight currently sits

Information Matters does not predict which scenario will play out. The framework's purpose is to help buyers and investors navigate uncertainty, not to claim certainty about the future. We do communicate where our weight currently sits, with explicit reasoning, so readers can disagree or update.

OUR WEIGHTED VIEW (MAY 2026)

- **Plateau — high weight.** Open-source closing on closed-source, inference price collapse, on-device inference scaling, and enterprise "good enough" sentiment all point to commoditisation as a real force.
- **Low-Cost Compute — high weight.** Hardware roadmap is favourable; distillation is real; on-device inference reinforcing the trajectory of inference prices downward.
- **Inertia — high weight, possibly under-discussed.** Historical pattern is consistent across every prior wave; pilot-to-production conversion rates remain low; governance and compliance bottlenecks are real.
- **Verticals — increasing weight.** Vertical specialists outperforming generalists; data-partnership advantages in legal, healthcare, finance are real.
- **Borders — increasing weight.** EU AI Act enforcement begins August 2026; US-China divergence persists; structural pattern follows GDPR's Balkanisation of cloud.
- **Rewire — moderate weight, slow timeline.** Real but happening unevenly; we expect a mixed pattern with Rewire-pace adoption in some sectors and Inertia-pace in others.
- **Frontier — declining weight.** Capability is still improving but the pace of differentiated leap appears to be slowing; reasoning keeps Frontier alive in narrow domains.
- **Expensive Compute — moderate weight as tail risk.** Energy bottleneck and geopolitical chip restrictions are real but their materialisation could go either way over a 3-year horizon.

WHERE THIS VIEW COULD BE WRONG

If a major capability release in late 2026 or 2027 dramatically expands what agentic AI can do, Frontier weight increases. If energy constraints bite faster than expected, Expensive Compute weight increases. If pilot-to-production conversion rates accelerate sharply, Inertia weight decreases and Rewire weight increases. Held with confidence proportional to evidence; updated quarterly.

SCENARIO OUTLOOK

IM's qualitative weighting at a glance

The chart below visualises the qualitative weighting set out in the previous section. Bars are sorted by weight; longer and darker bars carry more weight in our interpretation of Trajectory Profiles. Weights are illustrative and refreshed quarterly.

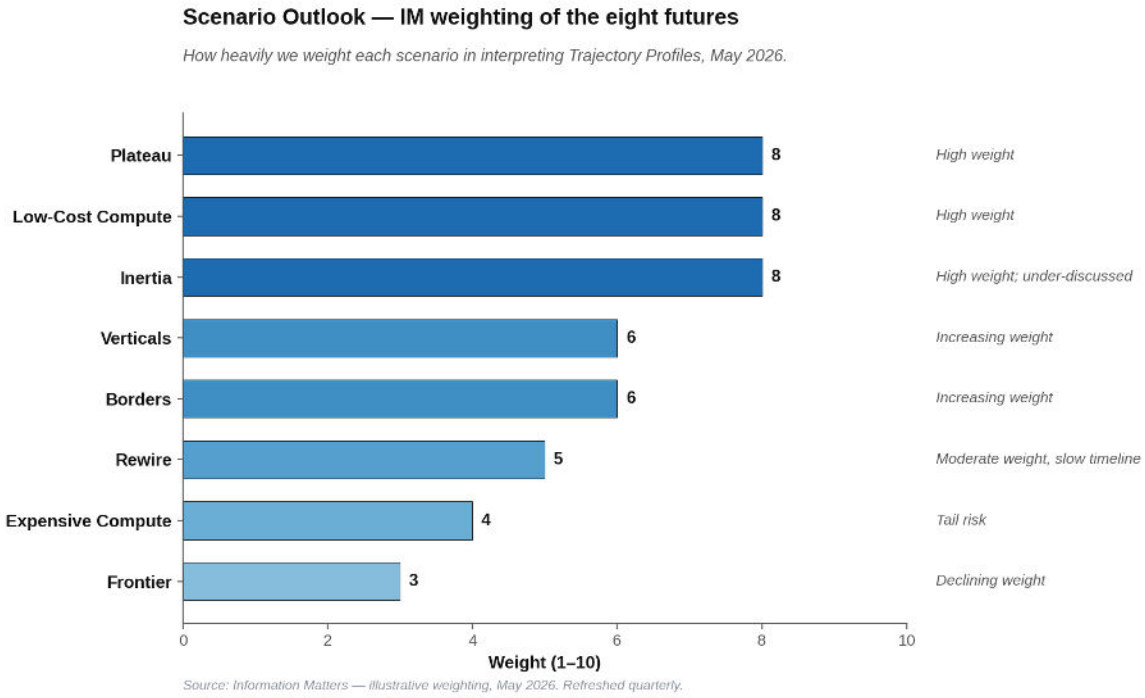

TIME HORIZONS

When each scenario is most likely to materially affect vendors

The eight scenarios do not all play out on the same timeline. Some are already manifest and persistent (Inertia, Verticals); some have a defined start date (Borders, with EU AI Act enforcement beginning August 2026); some are longer-arc scenarios that will unfold over the next five years (Rewire). The chart below maps the indicative period during which each scenario is most likely to drive vendor outcomes.

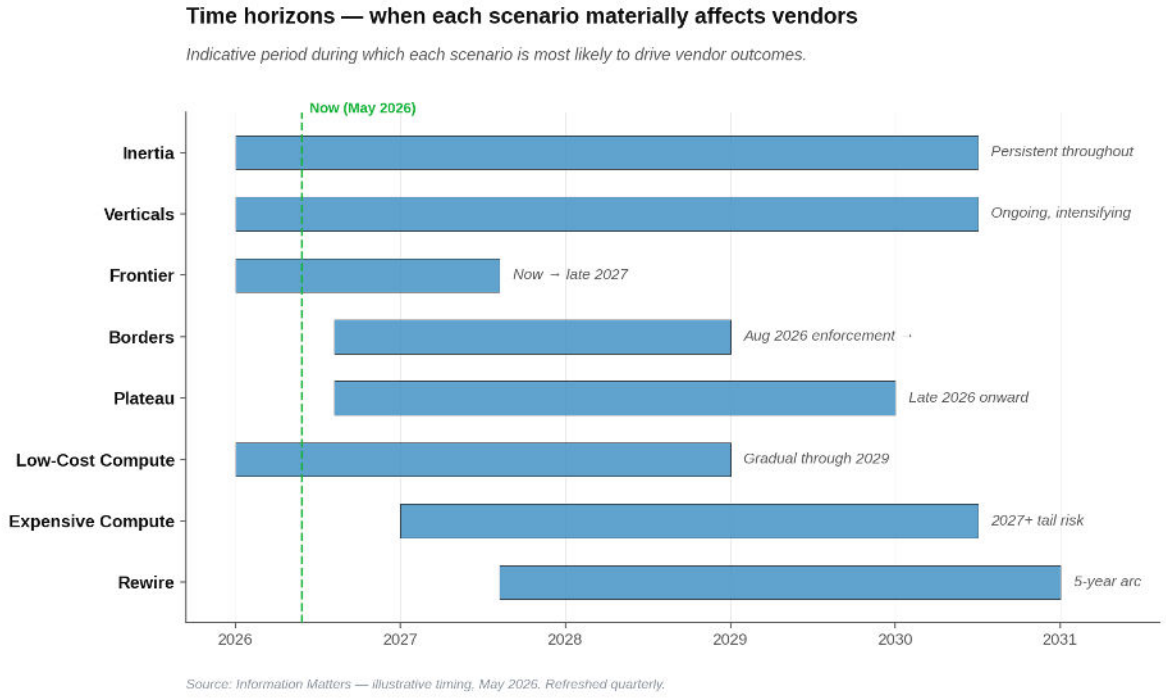

INDICATORS DASHBOARD

What to watch over the next twelve months

Scenario	Headline indicators
Frontier	GPT-5/Claude-5/Gemini-3 release cadence; computer-use reliability >80%; foundation providers shipping vertical agents.
Plateau	Cloud inference price trajectory; open-source parity on commercial benchmarks; on-device inference adoption; enterprise "good enough" sentiment.
Verticals	Vertical specialists winning enterprise deployments; sector-specific data partnerships; vertical fine-tuned model performance.
Low-Cost Compute	Cloud inference price decline rate; specialty hardware market share; on-device adoption; consumer agentic product MAU.
Expensive Compute	Inference price stabilisation or rise; hyperscaler capex; firm-power agreements; chip-export-restriction tightening.
Rewire	Enterprise organisational restructuring; agent-mediated procurement at scale; B2B SaaS revenue mix shifting.
Inertia	Pilot-to-production conversion rates; tier-1 SaaS incumbent churn; aggressive challengers raising down rounds; long-cycle enterprise software vendor agentic-adjacent revenue.
Borders	EU AI Act enforcement actions; US/EU/China regulatory divergence; regional data-residency standardisation; local-champion AI vendor market share.

ABOUT INFORMATION MATTERS

Information Matters is an independent analyst publication covering the agentic AI sector. We publish a quarterly market outlook, thematic deep-dives, and standing trackers. The Information Matters Brief is the weekly Substack. Methodology and data are disclosed in full across all our outputs.

Methodology note

This briefing 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 as well as our proprietary database of agentic AI companies and the public technology, regulatory, and economic literature relevant to each scenario. 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.