**INFORMATION MATTERS**

**ANALYST RESEARCH // CODING AGENTS**

# Coding Agents: The 2026 Landscape

Twelve vendors scored across the IM Framework — Defensibility, Disruption Potential, and Trajectory Profile across the eight scenarios from our Eight Futures Macro Briefing.

**May 2026**

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

Information Matters | May 2026

**Report #IM107**

## Headline findings

- **Foundation-model-shipped coding agents reshape the cohort top.** GitHub Copilot and Anthropic Claude Code occupy the two highest-robustness positions in our scoring — Copilot on incumbent distribution and balance-sheet resilience, Claude Code on capability-leading momentum and novel-capability scoring. The category's top tier is no longer pure-play startups; it is foundation-providers and platform incumbents shipping vertical coding products.
- **Cursor is the asymmetric case — highest disruption potential, lowest defensibility-adjusted robustness in the top five.** Anysphere's 10/10 on Category-Shaping Signal and 10/10 on Momentum is matched by below-cohort scores on multiple Trajectory scenarios where wrapper-thin business models price poorly (Frontier, Inertia, Expensive Compute). The cohort's most-discussed vendor is also its most-fragile to specific futures.
- **Autonomous-SWE vendors carry the cohort's Novel Capability ceiling but its Defensibility floor.** Cognition (Devin), Magic, and Lovable score the highest in the cohort on Novel Capability (9, 8, 8) and the lowest on Customer Entrenchment and Distribution & Ecosystem. The structural risk for this segment is that capability leadership without distribution depth is a fragile platform on which to build a multi-year integration commitment.
- **JetBrains AI is the cohort's quiet incumbent.** Robustness 5.7 ranks third behind only Copilot and Claude Code, with the highest Customer Entrenchment score outside the platform incumbents. The vendor most overlooked by the venture-narrative crowd is structurally well-positioned across the eight scenarios.
- **Inertia and Borders favour the incumbents disproportionately.** The two Trajectory scenarios where established distribution and balance-sheet strength matter most are also the scenarios where Copilot, JetBrains AI, and Claude Code outscore the rest of the cohort by the widest margin. Buyers building five-year integration commitments should treat these scenarios as load-bearing in vendor selection.
- **Robustness is the buyer's optimisation target, not headline momentum.** The cohort's robustness range (3.4 to 6.9) is wider than its momentum range (4 to 10). A vendor with strong momentum but weak robustness is a fragile bet for an enterprise buyer making a multi-year commitment. Our framework's central editorial claim, applied to this cohort: optimise for the vector, not the headline number.

# 1. EXECUTIVE SUMMARY

This report covers twelve coding-agent vendors active in the segment as of May 2026, profiled and scored through the full IM Framework: Defensibility (four sub-dimensions), Disruption Potential (four sub-dimensions), and Trajectory Profile across the eight scenarios from #IM105. Scores are qualitative analyst judgements on a 1–10 scale; the underlying numerical grid is the canonical input to the three Framework charts (the 2x2 positioning quadrant on the cover, the cross-company Trajectory Profile heat-map, and the per-vendor Trajectory bar charts on each profile page).

Two analytical threads carry through the report. The first is the *moat geography* of the coding-agent cohort. Defensibility concentrates with the platform-scale incumbents — GitHub Copilot at Microsoft, JetBrains AI at JetBrains — and with the foundation-model-shipped capability layer that is Anthropic Claude Code. Disruption Potential concentrates with the venture-funded mid-tier and with the autonomous-SWE segment: Cursor scores the highest single Disruption Potential in the cohort, with Cognition / Devin, Lovable, and Magic carrying the Novel Capability ceiling. The middle of the cohort — Sourcegraph, Replit, Codeium, Tabnine — sits between the two pressures with neither end's structural advantage.

The second thread is *robustness as the buyer's optimisation target* . The cohort's robustness statistic (mean of eight Trajectory scores minus half the standard deviation across them) ranges from 6.94 to 3.40. The two highest-robustness vendors are GitHub Copilot and Anthropic Claude Code, both of which score consistently across the eight scenarios; Copilot is structurally favoured in Inertia and Borders specifically, the two scenarios where established distribution and balance-sheet strength matter most. The lowest-robustness vendor, Cognition (Devin), is also the cohort's highest scorer on Novel Capability — exactly the fragility pattern the framework is designed to expose. Cursor, the cohort's headline momentum case, ranks seventh on robustness despite a 9.2-average Disruption Potential score; its low scores on Frontier, Inertia, and Expensive Compute are the three specific scenarios where wrapper-thin business models price poorly.

The category is consolidating. Microsoft's GitHub Copilot, Anthropic Claude Code, and the JetBrains AI roadmap each represent foundation-provider or platform-scale entries into a segment that until 2024 was dominated by startups. The cohort's mid-tier — the Series B / Series C IDE-and-assistant vendors and the more capital-intensive autonomous-SWE vendors.

## 2. THE CODING-AGENT CATEGORY IN 2026

The coding-agent segment is the most-funded and most-public-facing vertical inside agentic AI. Our Q1 2026 market-sizing work put coding agents at the $5–7 billion range for 2026, with two named vendors — Cursor at approximately $2 billion ARR and Anthropic Claude Code at a $2.5 billion run-rate disclosed alongside Anthropic's February 2026 Series G — together accounting for roughly half that segment. GitHub Copilot, embedded in Microsoft's Office 365 channel, ended Microsoft's Q2 FY2026 with 4.7 million paid subscriptions, a scale that no pure-play startup can replicate.

The cohort's mid-tier faces a specific structural question: gross-margin economics for pure-play coding agents have not been publicly disclosed by any major vendor, and the headline ARR scaling has been valued on a 75–85 per cent software-margin comparator. The trigger event most likely to settle the question is an Anysphere (Cursor) S-1 filing during 2026; we expect the disclosed gross margin to land in the 50–60 per cent range, which would materially re-price the venture-funded segment.

Three structural facts shape the cohort.

The **foundation-model providers have entered the category directly** . Anthropic Claude Code is the most visible case but not the only one — OpenAI's Codex products, Google's Gemini Code Assist and the agentic Jules tool, and the broader Big Three pattern of shipping vertical coding agents on top of the same foundation model that the third-party startups depend on. The implication is that the pure-play coding-agent vendors are competing for distribution against the platforms whose foundation-model output they themselves resell.

The **platform incumbents are not standing still** . GitHub Copilot has shipped agentic features through 2025–2026 that match what the venture-funded mid-tier launched eighteen months earlier, with the procurement and Microsoft-channel distribution advantages that the mid-tier has no equivalent for. JetBrains AI has rolled out the Junie agentic capability inside the JetBrains IDE family, addressing the professional-developer base that has historically been JetBrains' core. Sourcegraph's Cody product addresses the same enterprise developer segment from a code-search foundation.

The **autonomous-SWE segment has reset the upper bound of what "coding agent" means** . Cognition's Devin, Magic's long-context engineering agent, Factory's droid-pattern teams, Lovable's text-to-app workflow — these are not autocomplete tools with an agentic veneer. They are products that, when they work, complete software-engineering tasks end-to-end with limited human supervision. The framework treats this segment with high Novel Capability scoring but currently low Defensibility scoring; the question over the briefing horizon is whether either the capability holds against foundation-provider competition or the defensibility builds before it has to be defended.

## 3. 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 in the customer base), Knowledge & Data Advantage (proprietary corpora or unique data assets), Distribution & Ecosystem (channel reach and partner integration), and Strategic Resilience (capital base and parent-company strength). **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 Eight Futures Macro Briefing (#IM105). 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). A vendor's Trajectory vector is the central output of the framework; the **robustness statistic** — the mean of the eight scores minus half the standard deviation across them — is the single number that captures how a vendor performs across the combined future rather than in any one scenario.

In this report, robustness is the buyer's optimisation target. A high-momentum vendor with low robustness is in a fragile position whatever its current trajectory looks like.

---

# 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. The four named quadrants — Dominant Innovators (top-right), Disruptive Challengers (top-left), Established Incumbents (bottom-right), Emerging Players (bottom-left) — frame the read.

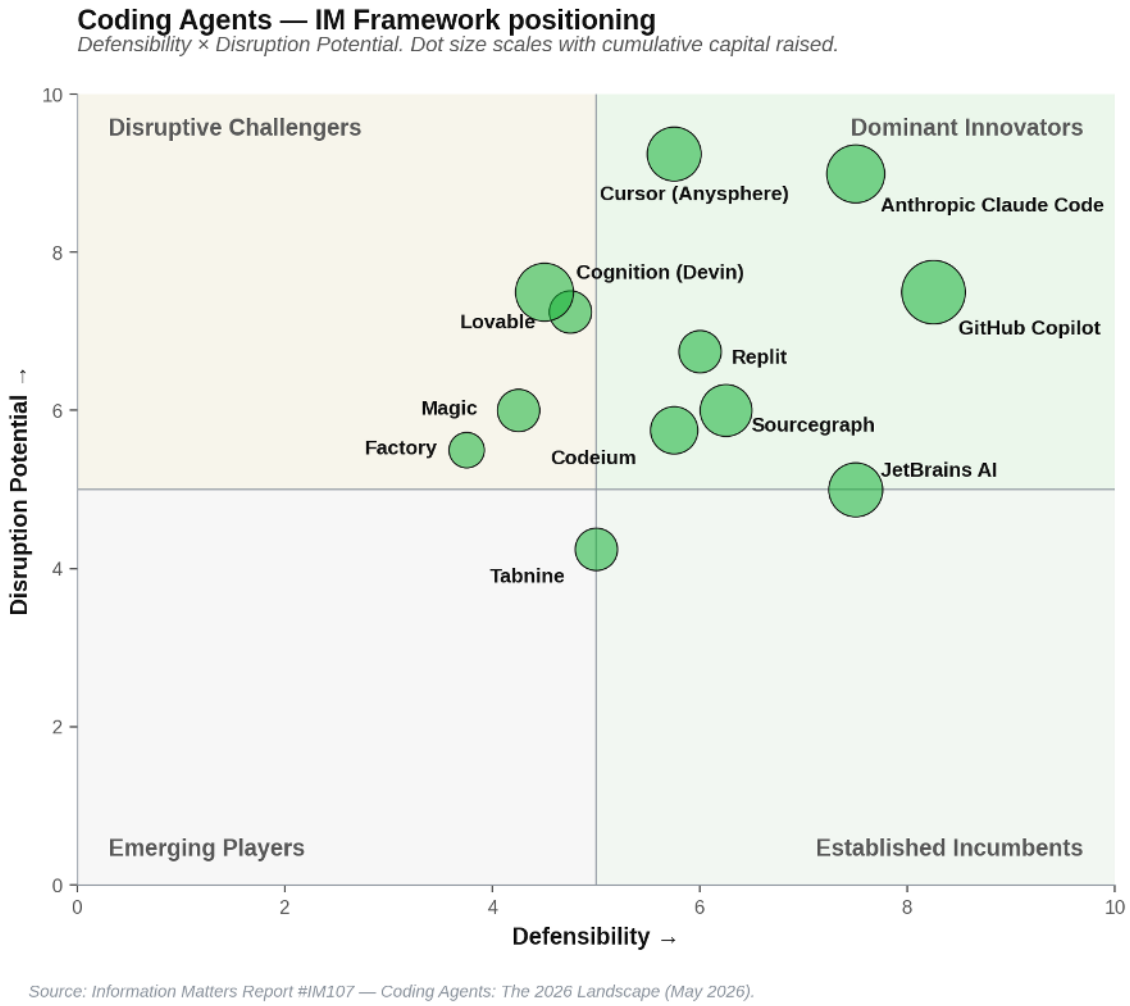


Source: Information Matters Report #IM107 — Coding Agents: The 2026 Landscape (May 2026).

*Fig. 1 — Coding Agents IM Framework positioning (2x2 quadrant). Dot size scales with cumulative capital raised. Source: Information Matters #IM107.*

**Dominant Innovators (top-right) — Anthropic Claude Code, GitHub Copilot.** Both score above 7 on Defensibility and above 7 on Disruption Potential. Their position in the same quadrant is the cohort's central market-structure fact: the two vendors that combine durable distribution and capital base with first-tier momentum and capability are also the two vendors that the venture-funded middle tier is most exposed to.

**Disruptive Challengers (top-left) — Cursor, Cognition (Devin), Lovable.** All three score high on Disruption Potential (9.2, 7.5, 7.2 respectively on the macro-construct) but below the

cohort median on Defensibility. This is the framework's classic "fragile-on-disruption" position: vendors driving the conversation but without the structural depth that would protect them through a category compression.

**Established Incumbents (bottom-right) — JetBrains AI, Sourcegraph.** Strong Defensibility, moderate Disruption Potential. JetBrains in particular sits as the cohort's clearest Established Incumbent: 7.5 on Defensibility, 5.0 on Disruption Potential, third-highest robustness in the entire cohort. The pattern is the textbook incumbent advantage — the venture-narrative crowd does not write about JetBrains and the buyer-procurement crowd reliably does.

**Emerging Players (bottom-left) — Tabnine, Codeium, Replit, Magic, Factory.** Five vendors sit below the cohort median on both axes. The framework treats this position as the most exposed to compression in the briefing horizon: insufficient distribution to be incumbent-safe, insufficient capability or category-shaping work to compete with the Disruptive Challengers.

The watchlist of cohort-adjacent vendors (Section 9) sits outside the formal quadrant scoring but is named in the body of this report because several of those vendors will move into the formal cohort in the next refresh.

---

## 5. TRAJECTORY PROFILES — THE CROSS-COHORT VIEW

The Trajectory Profile heat-map (page following this section) 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 which encodes the same data with the IM brand colour.

The cohort splits cleanly into three groups when read across the eight scenarios. **A robust-everywhere top tier** — GitHub Copilot and Anthropic Claude Code — scores consistently across every scenario, because their parents own the distribution channels (Microsoft 365, the Anthropic API) and the capital to absorb most futures. **A quiet incumbent tier** — JetBrains AI, Sourcegraph and Replit — has a lower ceiling but no scenario where the floor falls dangerously, because each holds a defensible workflow position (the JetBrains IDE family, enterprise code-search, the browser-native developer environment) that survives most paths. **A high-variance challenger tier** — Cursor, Lovable, Magic, Cognition's Devin and Factory — pairs strong scores in agentic-native scenarios (Rewire, Low-Cost Compute) with weak scores in the scenarios that reward incumbent scale (Inertia, Frontier, Borders). The framework treats the variance as the substance: the same vendor that the venture-narrative crowd calls category-defining is the vendor most exposed when deployment lags or regulation fragments the market. The robustness statistic captures both at once, and is the buyer's optimisation target for a multi-year commitment.

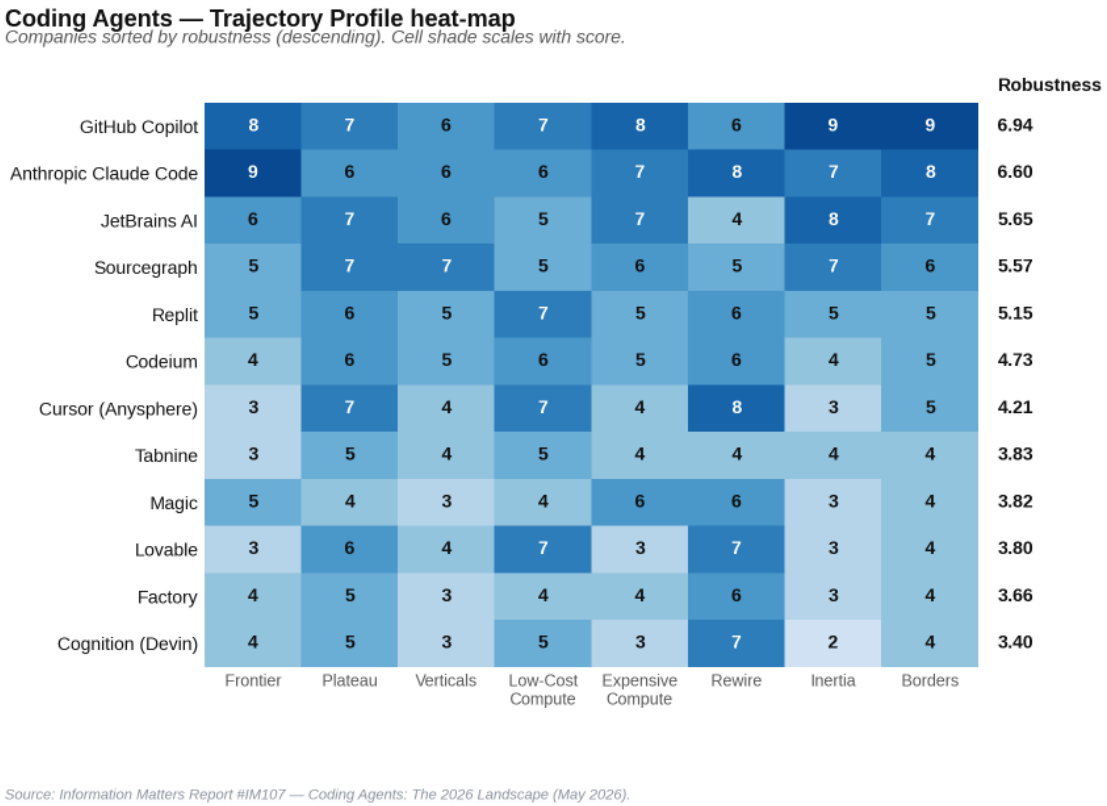


*Fig. 2 — Trajectory Profile heat-map across the eight scenarios. Companies sorted by robustness. Source: Information Matters #IM107.*

## 6. VENDOR COHORT — TWELVE PROFILES

Twelve per-vendor profiles follow, ordered by Trajectory Profile robustness. Each carries the macro-construct scores in narrative form and the per-vendor Trajectory bar chart on the page.

## 6.1 GitHub Copilot (Microsoft)

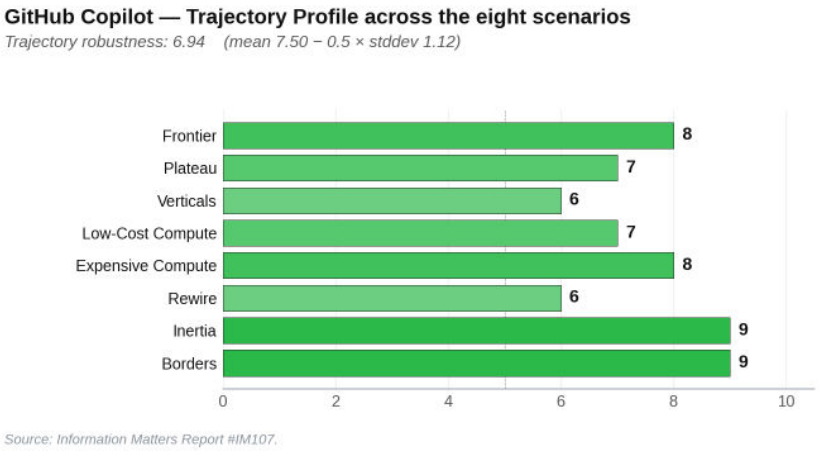


*Trajectory Profile — GitHub Copilot (Microsoft). Source: Information Matters #IM107.*

*Defensibility 8.2 / Disruption Potential 7.5 / Robustness 6.94.*

*Defensibility.* The strongest sub-dimension is Customer Entrenchment: 4.7 million paid subscriptions at the end of Microsoft's Q2 FY2026 is the deepest seat base in the cohort and embeds GitHub Copilot inside the procurement, identity, and policy frameworks of Microsoft 365 enterprise accounts. Distribution & Ecosystem is similarly first-tier through the Microsoft Azure and Visual Studio channels and the GitHub developer-base default. Strategic Resilience is parent-balance-sheet sized. The weakest sub-dimension is Knowledge & Data Advantage — Copilot's underlying capability depends on the OpenAI / Microsoft model partnership rather than on a proprietary corpus that Copilot itself controls.

*Disruption Potential.* Momentum is high but not category-defining; the seat-base is enormous but the per-seat-revenue growth runs on enterprise procurement timelines rather than the venture-scale rates the pure-play startups post. Novel Capability has improved markedly through the 2025–2026 rollout of agentic features (Copilot Workspace, the Coding Agent, Copilot Extensions) but the framework reads these as catch-up moves rather than category-shaping ones. Team Velocity is parent-organisation-sized rather than startup-sized. Category-Shaping Signal is strong by definition — Copilot defined the category in 2021.

*Trajectory commentary.* Best-positioned for **Inertia** (the textbook incumbent advantage; enterprise procurement defaults to Microsoft when adoption is slow) and **Borders** (Microsoft's regional compliance and data-residency footprint is incumbent-grade across every major jurisdiction). Worst-positioned for **Rewire** (large-org change-management friction inside Microsoft's existing distribution apparatus is real) and **Verticals** (Copilot's value proposition is horizontal, and a verticalised future favours sector-specific competitors).

## 6.2 Anthropic Claude Code

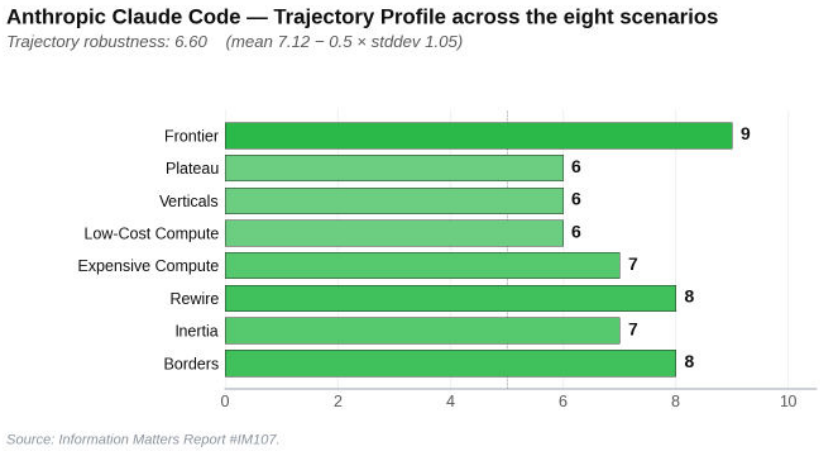


*Trajectory Profile — Anthropic Claude Code. Source: Information Matters #IM107.*

*Defensibility 7.5 / Disruption Potential 9.0 / Robustness 6.60.*

*Defensibility.* The strongest sub-dimension is Strategic Resilience: Anthropic's February 2026 Series G, the $2.5 billion Claude Code run-rate disclosed alongside it, and the parent-company run-rate of approximately $30 billion reached by April 2026 together put Claude Code in a capital and balance-sheet position that no third-party coding-agent startup can match. The Claude Code product itself is roughly one-twelfth of Anthropic's overall revenue scale, which means the product's runway is not constrained by its own unit economics in the way a pure-play vendor's would be. Knowledge & Data Advantage is reinforced by Claude's underlying capability — the Anthropic model family scores at the top of public coding benchmarks and the parent-organisation training pipeline is captive to Claude Code in a way that the multi-foundation wrappers cannot replicate. Customer Entrenchment is the weakest sub-dimension at this stage — Claude Code's user base is large and growing but newer, and switching cost has not yet compounded the way Copilot's seat-based lock-in has.

*Disruption Potential.* Momentum is the highest in the foundation-provider segment — Claude Code's $2.5 billion run-rate is the most consequential coding-agent disclosure in the briefing window. Novel Capability is category-defining for the workflow-first, sub-agent, computer-use-adjacent design pattern. Team Velocity is Anthropic engineering. Category-Shaping Signal is high — Claude Code is reshaping what foundation-provider coding agents look like, and the third-party wrapper segment now competes against a target that ships its own capability releases.

*Trajectory commentary.* Best-positioned for **Frontier** (Claude Code rides the capability frontier directly, with Anthropic as the deepest possible alignment of model and product) and **Rewire** (agentic-native architecture from the off, sub-agent and tool-use patterns line up with the Rewire scenario). Worst-positioned for **Plateau** (the value proposition leans on Claude's capability lead; commoditisation of foundation-model output compresses the moat) and **Low-Cost Compute** (Anthropic's pricing model relies on the inference economics the Low-Cost Compute scenario would erode).

## 6.3 JetBrains AI

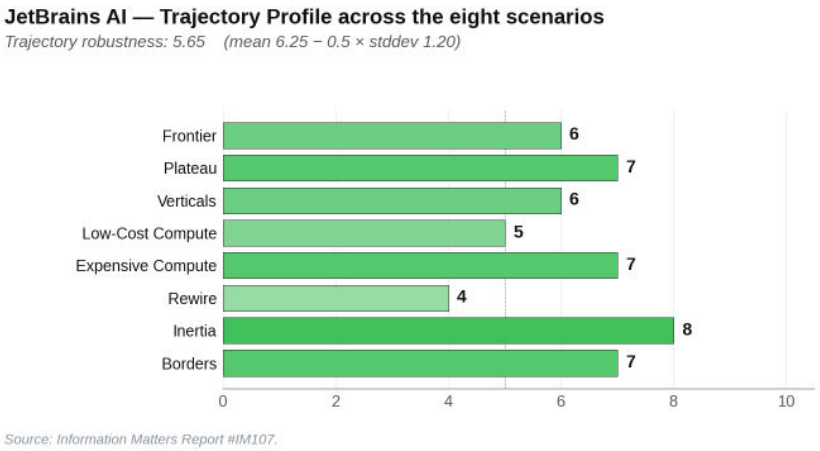


*Trajectory Profile — JetBrains AI. Source: Information Matters #IM107.*

*Defensibility 7.5 / Disruption Potential 5.0 / Robustness 5.65.*

*Defensibility.* The strongest sub-dimension is Customer Entrenchment: the professional-developer base across the JetBrains IDE family (IntelliJ IDEA, PyCharm, WebStorm, GoLand, Rider) is the largest paid-IDE installed base outside Visual Studio, and a developer who chooses a JetBrains IDE is one who has already paid for tooling rather than defaulting to the free options. Distribution & Ecosystem is similarly first-tier through that IDE family's two-decade compounding of enterprise-developer relationships. Strategic Resilience is reinforced by JetBrains being privately held, profitable, and Czech-headquartered — a balance-sheet posture more like a German Mittelstand company than a venture-funded startup. The weakest sub-dimension is Knowledge & Data Advantage; JetBrains' AI line uses third-party foundation models and the IDE-context advantage is real but bounded.

*Disruption Potential.* Momentum is steady rather than category-defining; the JetBrains AI Assistant adoption curve has tracked the IDE adoption curve more than driving its own. Novel Capability has improved through the Junie agentic line but the framework reads it as following rather than leading. Team Velocity is high engineering quality on a slower cadence. Category-Shaping Signal is muted — JetBrains is structurally a follower in the AI conversation, by design.

*Trajectory commentary.* Best-positioned for **Inertia** (deep installed-base seats do not churn quickly even in a slow-adoption future) and **Plateau** (commoditisation of foundation-model output is broadly neutral for vendors whose value is the IDE rather than the model). Worst-positioned for **Rewire** (incumbent IDE architecture has the hardest path to agentic-native restructuring) and **Frontier** (capability-frontier scenarios reward the model-and-product alignment that JetBrains' third-party-model approach cannot match).

## 6.4 Sourcegraph

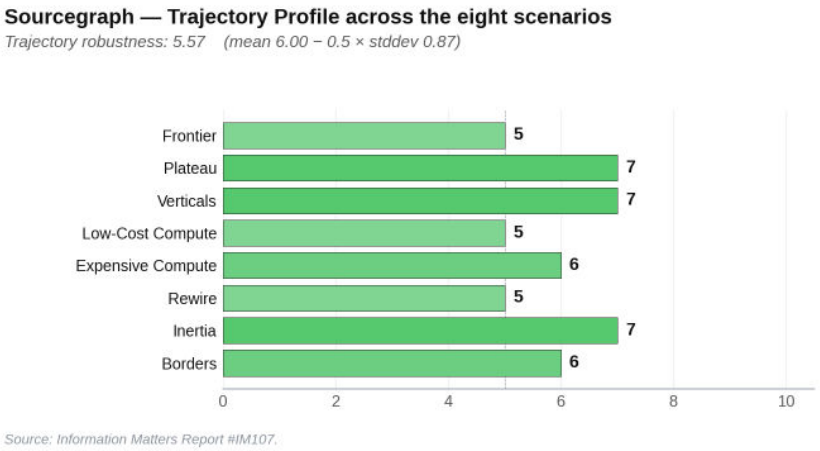


*Trajectory Profile — Sourcegraph. Source: Information Matters #IM107.*

*Defensibility 6.2 / Disruption Potential 6.0 / Robustness 5.57.*

*Defensibility.* The strongest sub-dimension is Knowledge & Data Advantage: Sourcegraph's enterprise code-search infrastructure indexes large customer code corpora at scale, and the Cody product activates that index for the customer's own codebase rather than for a generic foundation-model context. Customer Entrenchment is meaningful within the enterprise base — code-search deployments are sticky once integrated into developer workflows. Distribution & Ecosystem is enterprise-sales-led rather than developer-grassroots. Strategic Resilience is moderate; Sourcegraph is well-funded but not at the foundation-provider scale.

*Disruption Potential.* Momentum is steady. Novel Capability is solid through Cody's multi-foundation-model design and the code-context engineering. Team Velocity is high. Category-Shaping Signal is moderate — Sourcegraph defined the enterprise-code-search-plus-AI niche but the niche is not the centre of the agentic-coding conversation.

*Trajectory commentary.* Best-positioned for **Plateau** (proprietary code-context becomes more valuable when foundation models commoditise) and **Verticals** (enterprise code corpora are themselves a vertical-specific data asset). Worst-positioned for **Rewire** (the code-search-plus-AI architecture is bounded by its own design heritage) and **Frontier** (the value proposition is not capability-frontier dependent).

## 6.5 Replit

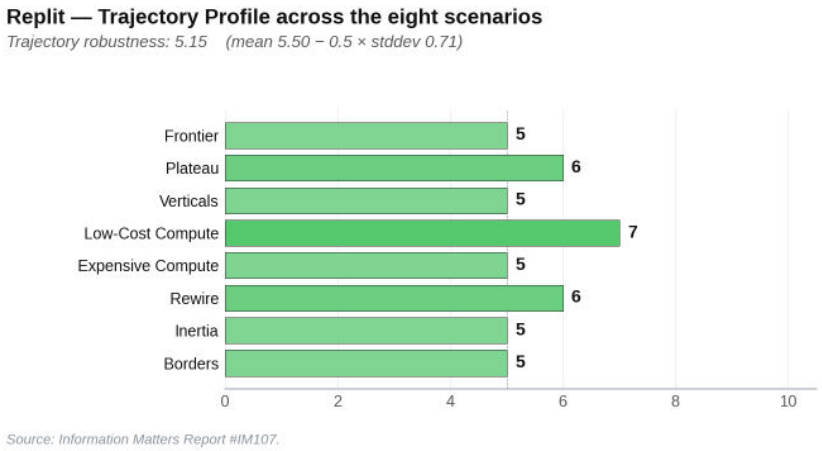


*Trajectory Profile — Replit. Source: Information Matters #IM107.*

*Defensibility 6.0 / Disruption Potential 6.8 / Robustness 5.15.*

*Defensibility.* Strongest sub-dimension is Distribution & Ecosystem: Replit's browser-native dev environment and community of millions of developers is a distribution moat the IDE-based competitors do not have. Customer Entrenchment is meaningful in the long-tail developer base but thinner at the enterprise tier. Knowledge & Data Advantage is moderate — Replit's agentic workflow data is accumulating but not yet a categorical moat. Strategic Resilience is reinforced by recent funding but Replit remains exposed to unit-economics questions that the platform-incumbent cohort does not face.

*Disruption Potential.* Momentum is real — the Replit Agent product has materially repositioned the company through 2025–2026 from an online IDE into a full-stack agentic dev environment. Novel Capability is high through Replit Agent's end-to-end app-build workflow. Team Velocity is high. Category-Shaping Signal is moderate; Replit is co-defining the browser-native agentic-dev pattern with Lovable but neither vendor dominates the segment yet.

*Trajectory commentary.* Best-positioned for **Low-Cost Compute** (browser-native and consumer-adjacent positioning benefits from on-device and cheaper inference) and **Rewire** (Replit Agent's end-to-end design is the textbook agentic-native architecture).
Worst-positioned for **Borders** (consumer-facing global product faces patchwork regulatory exposure) and **Inertia** (the enterprise procurement vector Replit is now leaning into is exactly the segment Inertia compresses).

## 6.6 Codeium

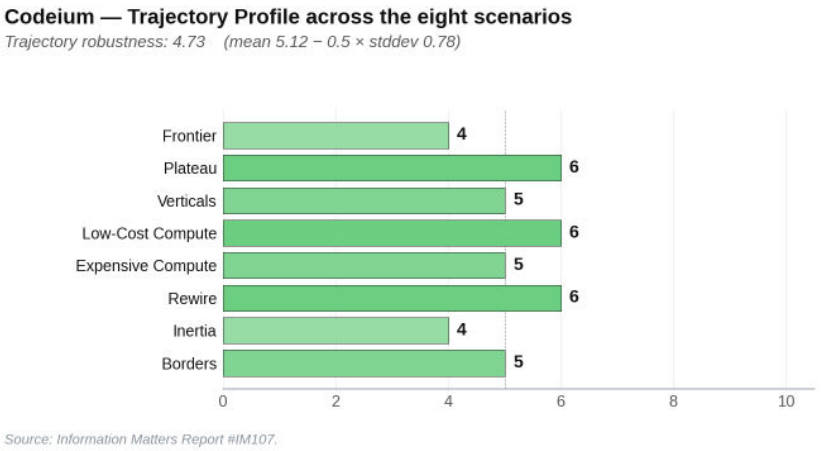


*Trajectory Profile — Codeium. Source: Information Matters #IM107.*

*Defensibility 5.8 / Disruption Potential 5.8 / Robustness 4.73.*

*Defensibility.* Strongest sub-dimension is Customer Entrenchment in the enterprise-secured, on-premise-capable segment that Codeium has carved out. Distribution & Ecosystem is solid through the enterprise-direct-sales motion that Codeium has built around its on-premise-capable positioning, even after the Windsurf IDE — which Codeium originated and shipped — was divested to Cognition Labs in May 2025. Knowledge & Data Advantage is moderate — Codeium does not own a proprietary corpus and depends on customer-context retrieval. Strategic Resilience is reasonable post the recent funding rounds.

*Disruption Potential.* Momentum is moderate. Novel Capability is meaningful in the on-premise-deployable agentic surface but framework-shaping rather than category-defining. Team Velocity is steady. Category-Shaping Signal is moderate; Codeium is the enterprise-procurement-friendly choice in a category where most competitors lean consumer or developer-grassroots.

*Trajectory commentary.* Best-positioned for **Plateau** (enterprise-friendly, on-prem-capable architectures are more valuable when foundation models commoditise) and **Rewire** (the agentic-native IDE design is well-positioned for restructured workflows). Worst-positioned for **Inertia** (Codeium's growth depends on enterprise-procurement velocity that Inertia would compress) and **Frontier** (capability-frontier futures favour vendors with closer model partnerships).

## 6.7 Cursor (Anysphere)

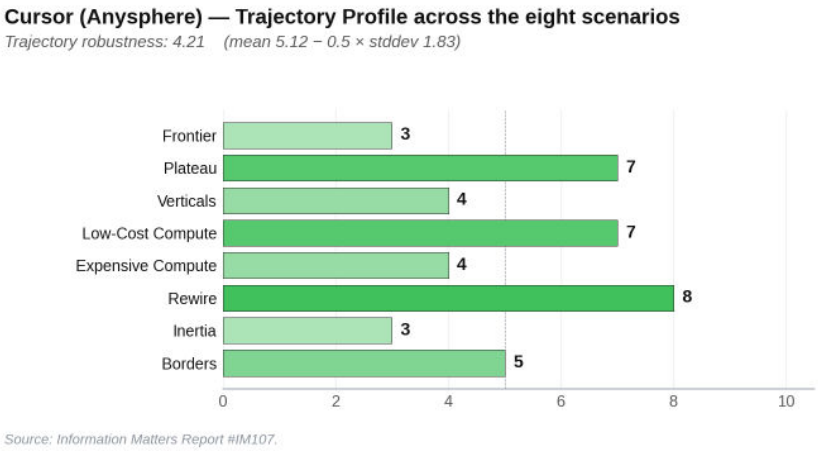


*Trajectory Profile — Cursor (Anysphere). Source: Information Matters #IM107.*

*Defensibility 5.8 / Disruption Potential 9.2 / Robustness 4.21.*

*Defensibility.* Strongest sub-dimension is Customer Entrenchment: Cursor's user base is the most-engaged in the cohort by developer-time-spent measures, and the workflow customisations (chat history, project context, custom agents) accumulate switching cost. Distribution & Ecosystem is direct-to-developer with growing enterprise reach but no platform-channel equivalent to Copilot or JetBrains. Knowledge & Data Advantage is the weakest sub-dimension — Cursor wraps multiple foundation models without owning a proprietary corpus. Strategic Resilience is well-capitalised at Series D+ scale; the open analytical question is the path to durable gross-margin economics, which no major coding-agent vendor has yet publicly disclosed.

*Disruption Potential.* Momentum is the highest in the cohort — Cursor reached approximately $2 billion ARR by February 2026 with the fastest growth curve in the segment. Novel Capability is real through the multi-cursor, tab-completion, composer-mode, and agent-mode product line. Team Velocity is high. Category-Shaping Signal is the highest in the cohort — Cursor reset enterprise expectations for AI coding tools and every subsequent product in the category has been built against the Cursor experience as the reference.

*Trajectory commentary.* Best-positioned for **Rewire** (agentic-native architecture from the off; Cursor's product is structurally aligned with the Rewire scenario's restructured-workflow assumption) and **Low-Cost Compute** (efficient inference on the customer-deployed surface is the Cursor product's central technical premise). Worst-positioned for **Frontier** (wrapper-thin business model competes with the foundation providers that ship their own coding agents), **Inertia** (the $2 billion ARR growth thesis requires aggressive enterprise-procurement conversion that Inertia would deny), and **Expensive Compute** (inference-cost inversion would compress the margin profile the equity is currently priced on).

## 6.8 Tabnine

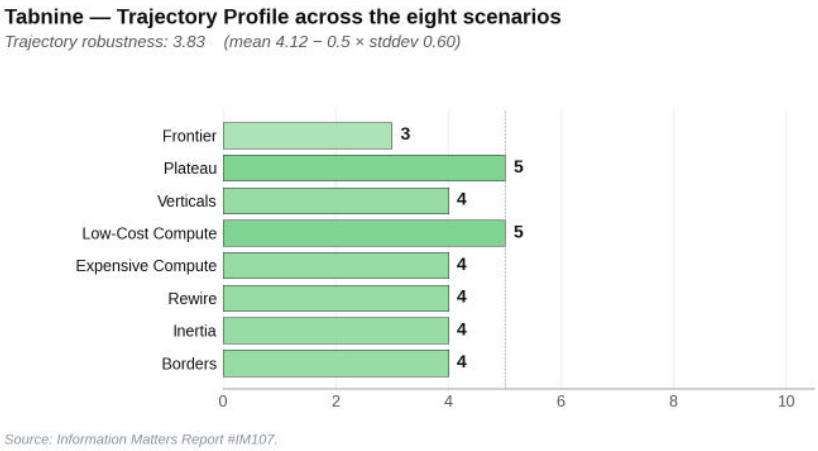


*Trajectory Profile — Tabnine. Source: Information Matters #IM107.*

*Defensibility 5.0 / Disruption Potential 4.2 / Robustness 3.83.*

*Defensibility.* Strongest sub-dimension is Customer Entrenchment in the privacy-first, self-hostable segment that has been Tabnine's identity since the autocomplete generation. Knowledge & Data Advantage is moderate — Tabnine has trained domain-specific models but the moat is bounded by the foundation-model-based competition. Distribution & Ecosystem is steady through the existing customer base. Strategic Resilience is the cohort-floor segment; Tabnine is profitable but at a different scale than the venture-funded peers.

*Disruption Potential.* Momentum is low — Tabnine's positioning is incumbent rather than aggressive growth. Novel Capability is moderate in the privacy-first segment but the segment itself is not where the agentic-AI conversation is happening. Team Velocity is steady. Category-Shaping Signal is low; Tabnine is defining the privacy-and-self-hosted niche rather than the broader category.

*Trajectory commentary.* Best-positioned for **Borders** (the self-hostable, privacy-first design pattern is the natural beneficiary of regulatory fragmentation) and **Plateau** (commoditisation of foundation-model output reinforces the privacy-and-deployment-flexibility moat).
Worst-positioned for **Frontier** (capability-frontier scenarios favour close model partnerships that Tabnine does not have) and **Rewire** (the autocomplete-anchored product architecture has the hardest path to agentic-native restructuring).

## 6.9 Magic

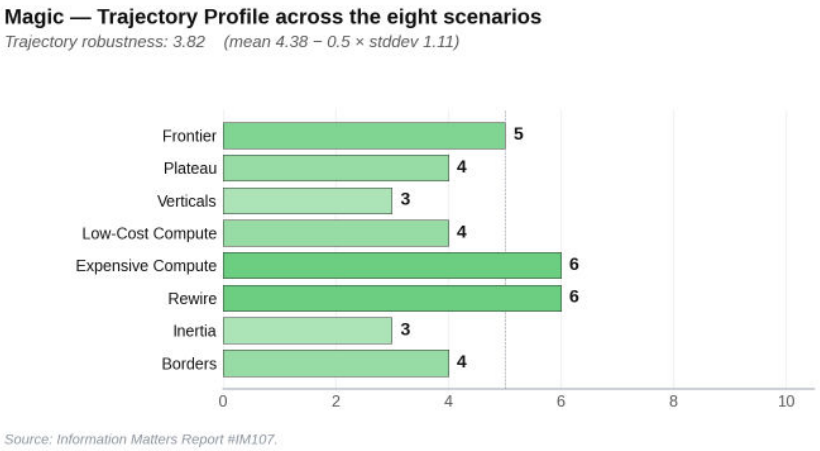


*Trajectory Profile — Magic. Source: Information Matters #IM107.*

*Defensibility 4.2 / Disruption Potential 6.0 / Robustness 3.82.*

*Defensibility.* The strongest sub-dimension is Strategic Resilience: Magic raised $470 million cumulative across 2024–2025 with backing from Eric Schmidt and Nat Friedman among others, a capital base that gives the team multi-year runway in a high-burn segment. Knowledge & Data Advantage is genuinely interesting — Magic's LTM-2-mini architecture targets very-long-context inference (100M token claimed context window) as a structural advantage that could matter at the right enterprise use cases. The other Defensibility sub-dimensions sit at the cohort floor; Magic does not have the customer base or the distribution position the more-established vendors carry.

*Disruption Potential.* Momentum is moderate — the product is not yet at general availability and the company has been deliberately quiet through 2025. Novel Capability is high (long-context architecture is a distinctive technical bet). Team Velocity is solid. Category-Shaping Signal is moderate; the long-context bet may yet redefine what enterprise-coding-agent context windows look like, but the influence has not materialised at scale.

*Trajectory commentary.* Best-positioned for **Expensive Compute** (the long-context architecture is structurally aligned with the more-expensive-inference future Magic's design optimises for) and **Rewire** (long-context agentic workflows fit the restructured-engineering scenario). Worst-positioned for **Inertia** (a capital-intensive, capability-betting startup needs enterprise procurement to clear at a pace Inertia would deny) and **Verticals** (Magic's positioning is horizontal long-context, not sector-specific data).

## 6.10 Lovable

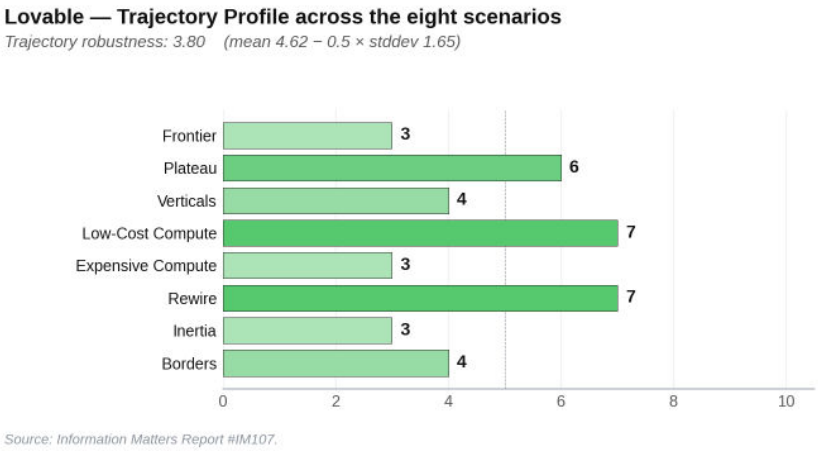


*Trajectory Profile — Lovable. Source: Information Matters #IM107.*

*Defensibility 4.8 / Disruption Potential 7.2 / Robustness 3.80.*

*Defensibility.* The strongest sub-dimension is the combination of Distribution & Ecosystem and Knowledge & Data Advantage — Lovable's text-to-app workflow has accumulated a long-tail of build-history and prompt-pattern data, and the consumer-and-startup distribution motion is materially different from the IDE-anchored cohort. Customer Entrenchment is moderate; the user base is large but the depth of any single relationship is shallow by enterprise standards. Strategic Resilience is reinforced by recent funding but at venture-startup scale.

*Disruption Potential.* Momentum is high — Lovable's product-market fit is real in the no-code-adjacent text-to-app segment. Novel Capability is high; the product does something genuinely category-defining (text-to-working-app in minutes) that the IDE-based competitors do not attempt. Team Velocity is high. Category-Shaping Signal is meaningful — Lovable is co-defining the text-to-app category with Replit and a handful of others.

*Trajectory commentary.* Best-positioned for **Rewire** (the entire product is built around the agentic-native, prompt-to-deployment workflow) and **Low-Cost Compute** (Lovable's consumer-startup economics benefit from the cheaper-inference future). Worst-positioned for **Inertia** (consumer-startup adoption rates are exactly the metric Inertia compresses) and **Expensive Compute** (the venture-startup unit economics require the current or cheaper inference cost regime).

## 6.11 Factory

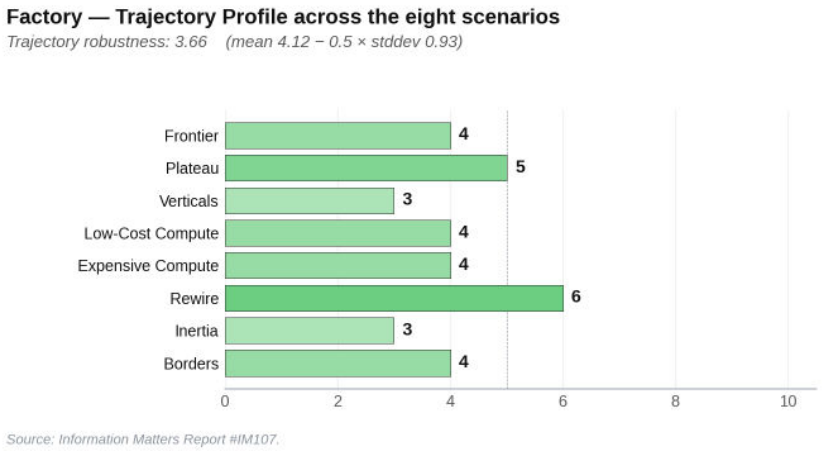


*Trajectory Profile — Factory. Source: Information Matters #IM107.*

*Defensibility 3.8 / Disruption Potential 5.5 / Robustness 3.66.*

*Defensibility.* The strongest sub-dimension is Strategic Resilience: Factory's Series A funding gives the team meaningful runway, though at the cohort floor on capital. Knowledge & Data Advantage is moderate — the droid-team product architecture accumulates execution data with each customer deployment. Customer Entrenchment and Distribution & Ecosystem sit at the cohort floor; Factory is an emerging-tier vendor in an enterprise-procurement-led segment.

*Disruption Potential.* Momentum is moderate. Novel Capability is meaningful through the droid-team architecture (pre-configured engineering-task agents organised into teams). Team Velocity is high. Category-Shaping Signal is moderate; Factory is part of the autonomous-SWE conversation but has not yet defined a sub-segment of it.

*Trajectory commentary.* Best-positioned for **Rewire** (the droid-team architecture is the textbook agentic-native restructuring) and **Plateau** (commoditisation of foundation models is broadly neutral for vendors whose value is the workflow orchestration, not the model). Worst-positioned for **Inertia** (an emerging-tier autonomous-SWE vendor needs aggressive enterprise procurement that Inertia denies) and **Verticals** (Factory's positioning is horizontal-SWE, not sector-specific).

## 6.12 Cognition (Devin)

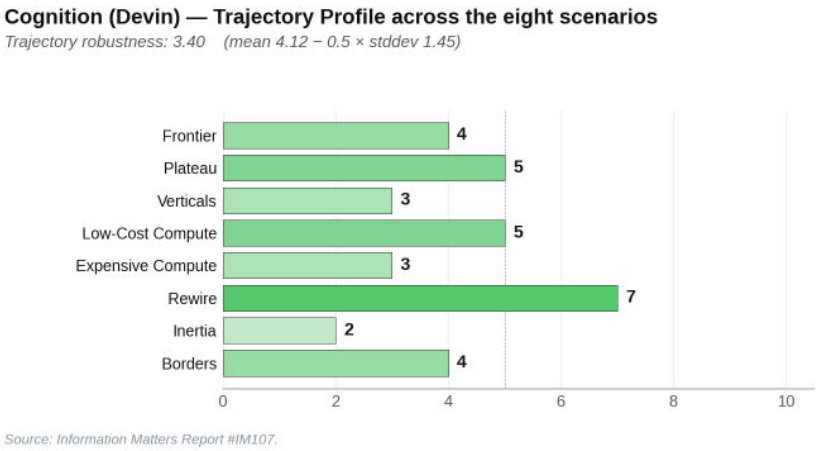


*Trajectory Profile — Cognition (Devin). Source: Information Matters #IM107.*

*Defensibility 4.5 / Disruption Potential 7.5 / Robustness 3.40.*

*Defensibility.* The strongest sub-dimension is Strategic Resilience: Cognition raised $200 million+ cumulative through 2024–2025 and the May 2025 acquisition of Windsurf added distribution depth that the prior pure-play position lacked. Customer Entrenchment and Distribution & Ecosystem sit at the cohort floor; Devin's enterprise deployments are pilot-grade rather than at-scale, and the value proposition of an autonomous engineer is one that enterprise procurement is still working out how to commit to.

*Disruption Potential.* Momentum is moderate — Devin's headline launch generated category-defining attention but the conversion of that attention into seats-and-revenue at scale is the open question. Novel Capability is the highest single score in the cohort at 9 — Devin defined what "autonomous SWE" means in commercial form. Team Velocity is high. Category-Shaping Signal is the second-highest in the cohort behind Cursor — Devin reset what coding-agent autonomy looks like.

*Trajectory commentary.* Best-positioned for **Rewire** (Devin is the textbook agentic-native product; the Rewire scenario is the one Devin's design assumes) and to a lesser extent for **Frontier** (Devin rides capability releases directly). Worst-positioned for **Inertia** (a high-burn autonomous-engineer vendor cannot sustain through a multi-year deployment lag), **Expensive Compute** (the unit economics of running Devin at scale are exposed to inference cost), and **Verticals** (Devin is horizontal-SWE and a verticalising market favours specialists).

# 7. CROSS-COHORT PATTERNS

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

**The foundation-provider entry has changed the cohort's top tier permanently.** Anthropic Claude Code and GitHub Copilot are the cohort's two highest-robustness vendors, both with double-digit-billion-dollar parent organisations and direct exposure to the model layer that the third-party vendors resell. This was not the case in the 2024 cohort, when the venture-funded pure-plays occupied the top spots and the foundation providers had not yet shipped vertical agents. The 2026 cohort's top tier is structurally different. Buyers committing to a coding-agent vendor over a five-year horizon are now choosing whether to commit to a foundation-provider product (with the model-and-product alignment advantage) or to a third-party product (with the multi-foundation flexibility advantage) — a choice the 2024 cohort did not present.

**The autonomous-SWE segment has a capability-versus-defensibility asymmetry that the framework specifically exposes.** Cognition (Devin), Magic, Lovable, and Factory all score in the top half of the cohort on Novel Capability (6–9) and the bottom half on Defensibility (3.8–4.8). The framework's editorial position on this asymmetry, made explicit in the headline-findings bullets, is that capability leadership without distribution depth is a fragile platform on which to build a multi-year integration commitment. The corollary for buyers: the autonomous-SWE vendors are most useful as proof-of-concept partners or as tactical deployments where the workflow is bounded — they are not yet load-bearing strategic vendors for enterprise software stacks.

**The mid-tier IDE-and-assistants segment carries the cohort's compression risk.** Sourcegraph, Replit, Codeium, and Tabnine sit in the cohort middle by Defensibility and Disruption Potential alike. None of them has the platform-incumbent depth that Copilot and JetBrains carry; none has the category-shaping signal that Cursor or Claude Code carry; none has the novel-capability ceiling that Cognition or Lovable carry. Our previous work on the coding category identified this segment as the most exposed to the gross-margin re-pricing. The Category Report's per-vendor scoring narrows the call: the four mid-tier vendors are not equally exposed (Sourcegraph and Replit score materially higher on robustness than Codeium and Tabnine), but all four face the same structural pressure of competing against platform-incumbent distribution above and capability-leading challengers below.

The watchlist (Section 9) names a further set of cohort-adjacent vendors that may move into the formal scoring in the next refresh — including Augment Code, Sweep, Z.ai's coding products, OpenHands, and the older autocomplete tier that has been steadily out-competed since 2024.

# 8. INDICATORS DASHBOARD

The single most-asked question of any IM Category Report is: what would change the call? 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; together they are the evidence we expect to refresh the report against at the Q4 2026 update.

The signals are grouped under three plain-language questions, rather than under the framework's internal construct names, so the dashboard reads as a watch-list rather than a methodology cross-reference.

## How fast is the cohort top tier moving?

- **Foundation-model coding-benchmark releases.** GPT-5, Claude-5, Gemini-3, and equivalent next-generation models. Where the foundation-provider products (Claude Code, OpenAI Codex, Gemini Code Assist) sit on the standard agentic coding benchmarks — particularly SWE-bench and the computer-use evaluations.
- **Quarterly revenue or run-rate disclosures from the cohort top.** Anthropic Claude Code, Microsoft GitHub Copilot, and Anysphere (Cursor) are the three most likely to disclose. Any of the three reporting an ARR or run-rate number materially above or below current expectations would move the top-tier scoring.
- **Shipping cadence.** Major-feature releases from the top five vendors; the rate at which agentic capabilities are reaching general availability.

## Where is the pressure on the middle tier?

- **An Anysphere (Cursor) S-1 filing.** Anysphere remains the cohort's clearest candidate for a 2026 public-equity event. A filing would publish a coding-agent gross margin for the first time.
- **Coding-vendor margin disclosures.** Any vendor-level gross-margin disclosure, in advance of the Anysphere event or alongside it, is the load-bearing signal for the margin-reset thesis.
- **Foundation-provider vertical-agent launches into the coding surface.** Microsoft, Anthropic, Google, and OpenAI each shipping deeper coding-specific agents directly into the surfaces (IDEs, Word, browser environments) the third-party cohort depends on.
- **Down rounds, acquisitions, or wind-downs at the venture-funded mid-tier.** Cognition, Magic, Lovable, Factory, and the smaller autonomous-SWE vendors are the most exposed. Any change of capital posture in this segment — particularly down rounds or acquihires — is a leading indicator of the consolidation thesis playing out.

## How is the buyer side responding?

- **Pilot-to-production conversion rates** for coding agents in Fortune-500 deployments. The Deloitte and Gartner numbers on enterprise agentic-AI deployment provide the cohort-level baseline; coding-specific evidence would refine the Inertia scoring.
- **Enterprise organisational restructuring around agentic engineering.** Named cases of engineering organisations reshaping team structures or workflow design around agent-mediated software development; this is the Rewire scenario in observable form.
- **Tier-1 SaaS incumbent seat disclosures.** Microsoft Copilot, JetBrains AI, and Sourcegraph seat or licence-base disclosures; renewal-rate trends at the Series-C and below tier.
- **Regulatory action affecting the cohort.** EU AI Act enforcement actions touching coding-agent vendors specifically; US export-control evolution affecting the foundation-model supply that the cohort depends on; any regional-champion coding-agent vendor reaching meaningful market share in its home jurisdiction.

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 coding-agent segment that are not part of the formal scoring in this Category Report. Inclusion in the Watchlist 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.

**IDE-and-assistants tier.** Augment Code (privacy-and-enterprise-focused, US, well-funded but still building distribution); Z.ai (Chinese coding-agent vendor, publicly listed in mainland China, addressing the China-domestic market in a Borders-divergent scenario); BLACKBOX.AI (Canadian, developer-focused); Cline (open-source agentic IDE); Roo Code (open-source fork lineage); Supermaven (small-team autocomplete with novel architecture).

**Autonomous-SWE tier.** Sweep (Series B, GitHub-issue-to-PR workflow); OpenHands (open-source autonomous SWE, formerly OpenDevin); Codegen (acquired); Charlie Labs, Cosine, Bloom, BootLoop, BASE44 (acquired by Wix, June 2025), Anything, AutonomyAI, Twill, Zencoder, Capy, Refresh AI, Fine, Hiverge, Sutro, iGent AI, Maya Labs, Tempo (alive-and-active emerging-tier vendors in the autonomous-SWE space).

**Foundation-provider coding products not formally scored.** OpenAI Codex / Codex CLI line; Google Gemini Code Assist and the agentic Jules; Alibaba's Qwen-coder line — all foundation-provider products that compete in the cohort but where the framework treats the parent foundation provider as the analytical unit rather than the coding product specifically.

**Code-quality and DevOps adjacencies.** AutoCodeRover (acquired), Backline AI, Ellipsis, Jazzberry, Propel, QualGent, Thunder Code, cubic (all in the Code Quality & Security sub-bucket); Ciroos, StarSling, a37 (DevOps & Infrastructure); Ardent, Nozomio Labs, Sourcebot, nao Labs (Data & Specialised); Anycode, Clidey, CodeBuddy, CodePal, Delty, FairMind, Humane, IBM Watson Group, LetMeCheck.ai, Potpie, Warp, XLooop, devlo (the Vertical (Coding & Technical) catch-all).

**Substrate-tier infrastructure not in cohort.** Anthropic (the parent of Claude Code, scored at product level not company level), Microsoft and GitHub (the parent of Copilot, scored at product level), JetBrains (the parent of JetBrains AI, scored at product level), Google (substrate provider not scored as a coding-agent vendor), OpenAI (substrate provider).

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

- **a37** — *United States* . Emerging-tier DevOps-and-infrastructure agent vendor; named in the Watchlist.
- **Alibaba (Qwen-coder)** — *China* . Publicly-listed foundation provider; ships Qwen-coder, a foundation-provider coding product not separately scored.
- **Anthropic (Claude Code)** — *United States* . Foundation-model provider; ships Claude Code as a foundation-provider-shipped coding agent on the Claude Agent SDK.
- **Anycode** — *United States* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **Anysphere** — *United States* . Parent company of Cursor; Series D+ venture-funded coding-agent vendor.
- **Anything** — *United States* . Emerging-tier autonomous-SWE vendor.
- **Ardent** — *United States* . Emerging-tier vendor in the data-and-specialised coding-agent bucket.
- **Augment Code** — *United States* . Privacy- and enterprise-focused IDE-and-assistants vendor; well-funded but still building distribution.
- **AutoCodeRover** — *Singapore* . Acquired code-quality-and-security agent.
- **AutonomyAI** — *United States* . Emerging-tier autonomous-SWE vendor.
- **Backline AI** — *United States* . Code-quality-and-security agent vendor.
- **BASE44** — *Israel* (Maor Shlomo, founder). "Vibe coding" platform that lets users build applications without code; acquired by Wix for approximately $80 million in June 2025 and continuing to operate as a distinct product under the Wix umbrella.
- **BLACKBOX.AI** — *Canada* . Developer-focused IDE-and-assistants vendor.
- **Bloom** — *United States* . Emerging-tier autonomous-SWE vendor.
- **BootLoop** — *United States* . Emerging-tier autonomous-SWE vendor.
- **Capy** — *United States* . Emerging-tier autonomous-SWE vendor.
- **Charlie Labs** — *United States* . Emerging-tier autonomous-SWE vendor.
- **Ciroos** — *United States* . Emerging-tier DevOps-and-infrastructure agent vendor.
- **Cline** — *Open-source* . Open-source agentic IDE in the IDE-and-assistants tier.
- **Clidey** — *United Kingdom* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **CodeBuddy** — *United States* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **CodePal** — *Israel* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **Codegen** — *United States* . Acquired autonomous-SWE vendor.

- **Codeium** — *United States* (Exafunction Inc.). Series C IDE-and-assistants vendor; originated the Windsurf IDE, which Cognition Labs acquired in May 2025; positioned as enterprise-secured and on-premise capable.
- **Cognition (Devin)** — *United States* (Cognition Labs). Late-stage / pre-IPO vendor; ships Devin, the category-defining autonomous software-engineering agent.
- **Cosine** — *United States* . Emerging-tier autonomous-SWE vendor.
- **cubic** — *United Kingdom* . Code-quality-and-security agent vendor.
- **Cursor (Anysphere)** — *United States* . Series D+ IDE-based coding-agent product from Anysphere; reached approximately $2 billion ARR by February 2026.
- **Delty** — *United States* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **Devin** — *United States* . The product name of Cognition's autonomous software-engineering agent; see Cognition (Devin).
- **devlo** — *United States* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **Ellipsis** — *United States* . Code-quality-and-security agent vendor.
- **Factory** — *United States* (Factory AI). Series A autonomous-SWE vendor; ships droid-team agents for software-engineering tasks.
- **FairMind** — *Italy* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **Fine** — *Israel* . Emerging-tier autonomous-SWE vendor.
- **GitHub** — *United States* (Microsoft subsidiary, acquired 2018). Owns the GitHub Copilot product; in the IM library tagged as the parent of the Copilot product entry.
- **GitHub Copilot (Microsoft)** — *United States* . The largest-installed-base AI coding assistant; 4.7 million paid subscriptions disclosed at the end of Microsoft's Q2 FY2026.
- **Google (Gemini Code Assist, Jules)** — *United States* . Publicly-listed foundation provider; ships Gemini Code Assist (IDE-based) and the agentic Jules tool — foundation-provider coding products not separately scored.
- **Hiverge** — *United Kingdom* . Emerging-tier autonomous-SWE vendor.
- **Humane** — *United States* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **IBM Watson Group** — *United States* . Coding-agent presence in the Vertical (Coding & Technical) catch-all bucket.
- **iGent AI** — *United Kingdom* . Emerging-tier autonomous-SWE vendor.
- **Jazzberry** — *United States* . Code-quality-and-security agent vendor.
- **JetBrains** — *Czech Republic* . Bootstrapped IDE maker; parent company of JetBrains AI.
- **JetBrains AI** — *Czech Republic* . The AI line at JetBrains, spanning the inline AI Assistant and the agentic Junie product, distributed across the JetBrains IDE family.

- **LetMeCheck.ai** — *India* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **Lovable** — *Sweden* (Lovable AB). Series B text-to-app workflow agent; the consumer-and-startup adjacency in the cohort.
- **Magic** — *United States* (Magic Inc.). Series B autonomous-SWE vendor; long-context architecture (LTM-2-mini, claimed 100M-token context window); backing from Eric Schmidt and Nat Friedman among others.
- **Maya Labs** — *India* . Emerging-tier autonomous-SWE vendor.
- **Microsoft** — *United States* . Publicly-listed; owns GitHub and through it the GitHub Copilot product.
- **nao Labs** — *United States* . Emerging-tier vendor in the data-and-specialised coding-agent bucket.
- **Nozomio Labs** — *United States* . Emerging-tier vendor in the data-and-specialised coding-agent bucket.
- **OpenAI (Codex)** — *United States* . Late-stage / pre-IPO foundation provider; ships the Codex and Codex CLI line as foundation-provider coding products not separately scored.
- **OpenHands** — *Open-source* . Open-source autonomous-SWE platform; formerly OpenDevin.
- **Potpie** — *United States* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **Propel** — *United States* . Code-quality-and-security agent vendor.
- **QualGent** — *United States* . Code-quality-and-security agent vendor.
- **RBCx** — *Canada* . RBC Bank's tech-banking arm; named in this report as the lender behind the March 2026 Spellbook debt facility.
- **Refresh AI** — *United States* . Emerging-tier autonomous-SWE vendor.
- **Replit** — *United States* (Replit Inc.). Series B browser-native developer environment; ships Replit Agent as a full-stack agentic build product.
- **Roo Code** — *United States* . Open-source agentic IDE in the IDE-and-assistants tier (fork lineage from earlier open-source tools).
- **Sourcebot** — *United States* . Emerging-tier vendor in the data-and-specialised coding-agent bucket.
- **Sourcegraph** — *United States* (Sourcegraph Inc.). Series D enterprise code-search vendor; ships Cody as the AI assistant on top of the code-search infrastructure.
- **Spellbook** — *Canada* . Mid-tier Word-native legal-AI vendor; named here for the parallel mid-tier compression pattern in the legal sector.
- **StarSling** — *United States* . Emerging-tier DevOps-and-infrastructure agent vendor.
- **Supermaven** — *United States* . Small-team IDE-and-assistants vendor with a distinctive autocomplete architecture.

- **Sutro** — *United Kingdom* . Emerging-tier autonomous-SWE vendor.
- **Sweep** — *United States* . Series B autonomous-SWE vendor; GitHub-issue-to-pull-request workflow.
- **Tabnine** — *Israel* (Tabnine Ltd). Series B privacy-first, self-hostable autocomplete-and-agent vendor.
- **Tempo** — *Canada* . Emerging-tier autonomous-SWE vendor.
- **Thunder Code** — *France* . Code-quality-and-security agent vendor.
- **Twill** — *United States* . Emerging-tier autonomous-SWE vendor.
- **Warp** — *United States* . Terminal-centric coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **Windsurf** — *United States* . Coding-agent IDE originated by Codeium (Exafunction Inc.); acquired by Cognition Labs in May 2025.
- **XLooop** — *United States* . Emerging-tier coding-agent vendor in the Vertical (Coding & Technical) catch-all bucket.
- **Z.ai** — *China* . Publicly-listed coding-agent vendor addressing the China-domestic market.
- **Zencoder** — *United States* . Emerging-tier autonomous-SWE vendor.

## 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.

**Anysphere.** The parent company of the Cursor coding-agent product; we name both the parent and the product in this report.

**ARR (Annual Recurring Revenue).** The revenue an annualised subscription book is generating at the time of measurement; the cohort's main public scale measure.

**Autonomous SWE.** A coding-agent product designed to complete software-engineering tasks end-to-end with limited human supervision, rather than acting as an assistant inside a developer's IDE.

**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.

**Category Report.** An IM publication type that profiles a defined competitive set of vendors with formal 1–10 framework scoring; this report is the format's reference example.

**Cohort.** The defined set of vendors selected for scoring in this Category Report; this cohort is twelve vendors.

**Customer Entrenchment.** A Defensibility sub-dimension in the IM Framework; measures how deeply a vendor's product is embedded in the customer's workflow 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; 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.

**Borders.** One of the Eight Futures; the scenario in which regulatory and geopolitical divergence fragments the global market into geographic spheres.

**Expensive Compute.** One of the Eight Futures; the scenario in which inference economics tighten and the cost of running agentic workloads rises.

**Frontier.** One of the Eight Futures; the scenario in which foundation-model capability continues to leap forward generation by generation.

**Inertia.** One of the Eight Futures; the scenario in which the technology arrives but enterprise deployment lags well below forecast — the "under-discussed scenario" in IM's weighting.

**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.

**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.

**Verticals.** One of the Eight Futures; the scenario in which sector specialists with proprietary domain data outperform horizontal generalists.

**Foundation provider.** A company that builds and ships the underlying large-language-model layer on which other agentic AI products depend; in this report Anthropic, OpenAI, Google, Microsoft (via its OpenAI partnership), and Alibaba are the named foundation providers.

**Framework chart.** One of the three visualisations the IM Framework reserves for Category Reports: the 2x2 quadrant (Defensibility × Disruption Potential), the per-vendor Trajectory bar chart (single vendor across the eight scenarios), and the Trajectory Profile heat-map (all cohort vendors across the eight scenarios).

**IDE (Integrated Development Environment).** The software application a developer writes code in — for example Visual Studio Code, JetBrains' IntelliJ family, Cursor, or browser-based environments like Replit.

**Knowledge & Data Advantage.** A Defensibility sub-dimension; 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.

**MCP (Model Context Protocol).** An emerging standard for connecting AI models to tools and data sources. Anthropic's MCP is the named version of the protocol referenced where relevant in this report.

**Momentum.** A Disruption Potential sub-dimension; measures the trajectory of revenue or adoption — how quickly a vendor is growing in observable scale terms.

**Novel Capability.** A Disruption Potential sub-dimension; measures whether a vendor's product does something category-defining that competitors do not currently match.

**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.

**Sector Context Briefing (SCB).** An IM publication type that covers a whole industry vertical qualitatively, without per-vendor 1–10 scoring.

**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 reference in this report.

**Strategic Resilience.** A Defensibility sub-dimension; measures a vendor's capital base, parent-company strength, and ability to survive a downside scenario.

**SWE-bench.** A public benchmark for coding agents that measures whether the agent can resolve real GitHub issues end-to-end; the most widely cited evaluation in the cohort.

**Team Velocity.** A Disruption Potential sub-dimension; 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.

**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 as well as our proprietary database of agentic AI companies, with the coding-agent cohort drawn from 83 vendors across six Vertical (Coding) sub-buckets in the database and scored against the IM Framework in May 2026. 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. 