INFORMATION MATTERS

ANALYST RESEARCH // CODING AGENTS

Coding agents: the margin reset is coming

Why the first pure-play S-1 will reset the implicit valuation of the coding-agent category — and when we expect it.

May 2026

Martin De Saulles | Principal Analyst
Fiona | AI Research Analyst
Information Matters | May 2026

Report #IM104

Headline findings

- **The cleanest growth story, no margin disclosure.** $2B Cursor ARR; $2.5B Claude Code; 4.7M paid Copilot subscriptions. None is a margin disclosure.
- **Anysphere's S-1 is the trigger.** Q3 2026 central case, with a one-quarter slip available. The only major coding-agent vendor whose business is monolithically a coding agent.
- **50-60% predicted disclosed margin.** Materially below the 75-85% SaaS comparator the equity is currently being priced against.
- **The category bifurcates.** Specialists with vertical or workflow defensibility (Cognition, Sourcegraph, Lovable, Tabnine, the QA tier) survive. The autocomplete middle is repriced or absorbed.
- **Anthropic's filing is consequential, not the reset.** It will disclose the inputs to every coding-agent vendor's COGS. The coding-agent margin reset itself requires a pure-play disclosure.
- **Discipline of the house.** Our 50-60% range is reasoned from inference economics, not sourced from a primary disclosure. We will publish the reconciliation when the disclosure arrives.

THE GROWTH STORY

The numbers everyone agrees on, and the question they don't answer

The coding-agent category has the cleanest growth story in agentic AI. The numbers are not in dispute, and they are not the question this report is interested in.

Cursor reached $2 billion in run-rate revenue in February 2026, the trajectory disclosed via CNBC's coverage of its November 2025 funding round (then at approximately $1 billion ARR) and confirmed at $2 billion in CEO Michael Truell's March 2026 Fortune interview. GitHub Copilot ended Microsoft's Q2 FY2026 with 4.7 million paid subscriptions disclosed on Microsoft's January 2026 earnings call, against the approximately 90 per cent Fortune 100 adoption figure Satya Nadella had disclosed on the preceding Q4 FY25 call in July 2025. Claude Code was disclosed at $2.5 billion run-rate revenue alongside Anthropic's February 2026 Series G announcement — Anthropic's first formal break-out of the product line in a corporate-finance instrument. The JetBrains 2025 Developer Ecosystem Survey, with 24,534 respondents, found that 85 per cent of working developers use AI tools regularly and 62 per cent use at least one AI coding agent. JetBrains' April 2026 workplace AI tools survey added the relevant counterweight: 74 per cent of developers had adopted specialised AI developer tools by January 2026 against the much smaller share of employers that report a formal AI tools policy in place.

The scale is real, the adoption signal is the strongest in any sub-segment of agentic AI, and the revenue trajectories at the top of the category are without precedent in enterprise software history. The Cursor curve from zero to $2 billion ARR is the fastest such climb the analyst community has documented, and we said so in our Q1-2026 report (§2.3). None of that is the contrarian claim of this piece.

The contrarian claim is that the category is being valued on revenue scale and growth velocity while the question that determines whether the equity is worth what it is currently priced at — gross margin — has not been answered by any of the major vendors. Until it is, the sector is being underwritten on an assumption that the unit economics of an application-layer agent that makes heavy frontier-model inference calls will arrive at SaaS-grade margins. Our reading is that they will not, that the first IPO filing or formal earnings break-out which discloses gross margin will be the trigger that prices that distinction in, and that the timing of that disclosure is closer than the consensus narrative implies.

THE MISSING QUESTION

Why no major coding-agent vendor has disclosed gross margin

The published case rests on three numbers: Cursor's $2 billion ARR with disclosed headcount under 100, GitHub Copilot's 4.7 million seats, and Claude Code's reported run-rate. None is a margin disclosure. The Cursor revenue-per-employee figure — better than $10 million per employee on the most aggressive published headcount of fewer than 100 staff, or closer to $13 million per employee on the IM library figure of 150 — is being read as a productivity miracle. It can equally be read as a structural dependence on a single category of variable cost — frontier-model inference — that does not appear on the disclosed denominator. Both readings are consistent with the data. The market is choosing the first.

The Q1-2026 report (§2.5) consolidated the primary-source data on inference economics. The conclusion held there and holds here. Anthropic and OpenAI have each disclosed flagship-model price cuts in the 60-80 per cent range over the eighteen months ending Q1 2026. That price-curve compression helps application-layer vendors who can stabilise per-user inference cost faster than per-user revenue compresses; it hurts those whose users learn to consume more tokens per session as the agents get more capable. The behavioural data points to the second pattern. JetBrains' April 2026 survey found developers who had adopted agentic coding tools were running them on a meaningfully larger share of their working time than developers who had adopted only autocomplete-style tools twelve months earlier. The unit of consumption has changed. A token is not the right unit for a margin model when the user's session length is doubling.

What we do not yet have is a disclosed gross margin from a major coding-agent vendor. Not from Cursor — pre-IPO. Not from Anthropic via Claude Code — Anthropic does not yet break out unit economics by product surface. Not from GitHub Copilot — Microsoft consolidates Copilot revenue inside its Productivity & Business Processes segment and reports against bundled offerings rather than a coding-agent-specific COGS line. The closest available reference points are indirect: Microsoft's Q2 FY2026 earnings called out commercial cloud gross margin in the high 60s percentage and explicitly noted "AI infrastructure scaling" as a margin headwind without quantifying the per-product impact; Salesforce has flagged consumption-based agent economics as a structural input to gross margin in its FY26 Q1 commentary, without disclosing per-product COGS for Agentforce. These are the most direct primary signals available, and they support the direction of the argument without supplying the number.

TWO COHORTS INSIDE THE CATEGORY

The structure of the 83 is bimodal in a way the headline narrative obscures

The Information Matters internal library, at the date of this report, contains 83 companies whose primary or principal positioning is in the coding-agent stack. The split across our six coding-stack subcategories is visible in Figure 1. The geographic distribution is dominated by the US (60 of 83), with smaller cohorts in the UK, Israel, France, Singapore, India, China, Italy, Canada, and the Netherlands. Funding stage is heavily weighted to the early end: only a handful of companies in the cohort sit at Series C or later, with the majority either pre-Series-A or unlabelled in our funding-stage taxonomy.

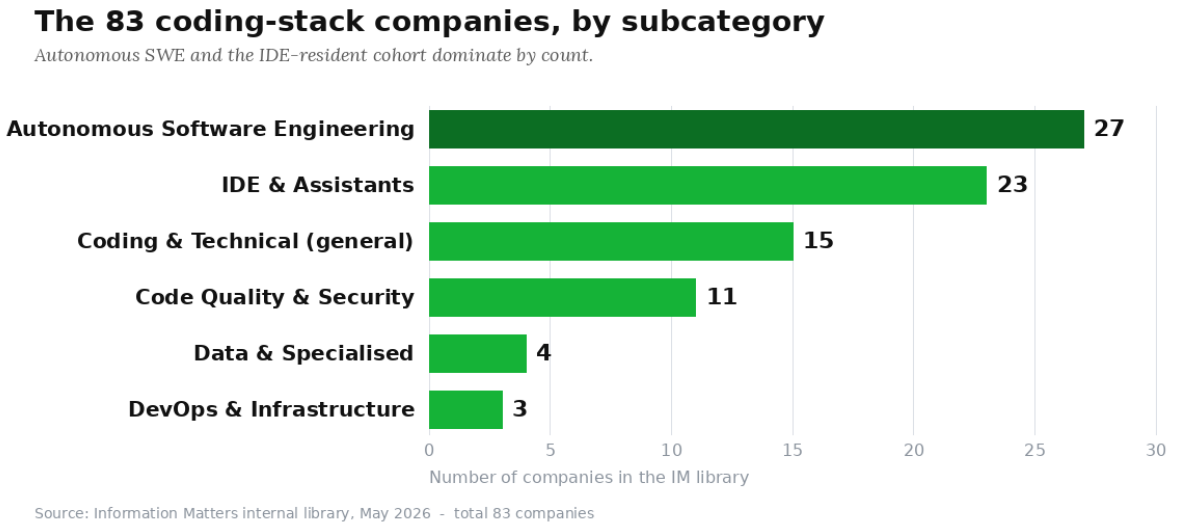

Figure 1. Subcategory composition of the 83 coding-stack companies in the IM library.

That distribution is the structural fact of the category. Two cohorts are visible inside the 83: a defensible-specialisation cohort with a vertical, workflow, or deployment posture outside the autocomplete primitive; and a cohort whose product is generic AI autocomplete or thin agentic wrappers around the frontier-model APIs, with no defensible specialisation against the buyers who will absorb the surface.

Cohort 1 — the specialists

Five names stand out from the IM library as cohort-1 specialists. They are not the only candidates, but they are the cleanest cases.

- **Cognition.** (US, late-stage). Devin reached approximately $73 million ARR on a standalone basis by June 2025; following Cognition's July 2025 acquisition of Windsurf, the combined post-acquisition run-rate has been reported at $150-155 million. Cumulative funding stands at approximately $1 billion through the September 2025 Series D at a $10.2 billion valuation. The strongest independent IPO candidate in the autonomous-SWE tier on a multi-year horizon.
- **Sourcegraph.** (US). Code intelligence as the moat — indexing, retrieval, the search-graph primitives that make context windows tractable for large codebases. The agentic surface is layered on that infrastructure rather than competing for the autocomplete primitive.
- **Lovable.** (Sweden, Series B). Reached $200 million ARR in November 2025 from a 2023-vintage company; approximately $650 million in cumulative funding through the December 2025 Series B at a $6.6 billion valuation. The defensibility argument is workflow lock-in. Different buyer, different production loop, different competitive set than Cursor.
- **Thunder Code** (France) and AutoCodeRover (Singapore, recently acquired). The QA-and-debugging specialists. Smaller and more vertical than the autocomplete cohort, but selling into a procurement frame (test, regression, security) where coding agents are not the budget centre and the revenue-per-customer profile is more sustainable. AutoCodeRover's acquisition by Sonar in February 2025 is the directional signal.
- **Tabnine** (Israel) and Refact (UK). The privacy- and self-hosted-deployment specialists. Tabnine's revenue trajectory inside the IM library is materially below the headline names; the defensibility argument is structural — regulated and air-gapped enterprise buyers will not deploy a Cursor-style agent that calls a frontier model over the public internet.

Cohort 2 — the structurally vulnerable

Five names from the IM library, drawn deliberately from a longer set, illustrative rather than exhaustive. The Windsurf trajectory is the clearest existing data point that the consolidation pattern is already in train at the autocomplete surface.

OpenAI announced a $3 billion Windsurf acquisition agreement in May 2025 — at the time, the largest acquisition OpenAI had committed to. The deal collapsed in July 2025 over the question of whether OpenAI's commercial arrangement with Microsoft would force disclosure of Windsurf intellectual property to Microsoft. Google then absorbed Windsurf's CEO Varun Mohan, co-founder Douglas Chen, and a senior research cohort into Google DeepMind under a $2.4 billion non-exclusive licensing arrangement, leaving the company without its founding leadership but with its product, brand, and the bulk of its engineering organisation intact. Cognition acquired the residual entity — IP, product, trademark, and the approximately 210 remaining employees — in July 2025 at a reported $250 million, against a Windsurf book of approximately $82 million ARR at close.

Three observations follow. First, the headline acquisition price compressed from $3 billion to $250 million for the corporate entity itself across two failed quarters — pricing in the IP-disclosure risk and the loss of senior leadership. Second, the eventual buyer was a cohort-1 specialist (Cognition), not a hyperscaler or foundation-model provider, which complicates the simpler reading of autocomplete consolidation as platform-absorption-only; a specialist with a defensible procurement frame can also be the absorber, and the bifurcation outcome we expect into 2027 will include some cross-cohort acquisitions of this shape. Third, the $82 million ARR base did not, in this product surface, prove adequate to support an independent runway in front of GitHub Copilot's installed distribution; that is the cohort-2 structural finding the saga makes concrete.

Other cohort-2 names in the IM library include Supermaven, Cline, Roo Code (all US autocomplete and IDE-assistant primitives without a visible vertical or workflow specialisation), Tempo (Canada), and several US-based autocomplete entrants from the 2024-25 cohort whose product surface overlaps directly with what Cursor, Copilot, and Claude Code already do. The honest qualifier is that public information on most of these companies is limited. The point is not to forecast individual outcomes; it is that the structure of the 83 is bimodal in a way the headline narrative obscures. The middle of the distribution — companies large enough to be visible but not large enough to be defensible — is where the margin-reset event will hit hardest.

THE TRIGGER EVENT

An Anysphere (Cursor) S-1 in Q3 2026

The cleanest test is Cursor. Anysphere is the only major coding-agent vendor whose business is plausibly close to a public-equity event during 2026, and Cursor's product is monolithically a coding agent — there is no enterprise platform layer obscuring the unit economics, no bundled cloud business, no foundation-model training revenue subsidising the application surface. An S-1 filing from Anysphere would publish a coding-agent gross margin without ambiguity for the first time. The CEO's March 2026 Fortune interview is the data point against this thesis: Truell told Fortune the company is focused on long-term mission rather than IPO. We have read that interview as a deflection of the timing question, not a denial of the eventual filing. The pattern of advisory engagement consistent with S-1 preparation — the breadth of legal counsel, the cadence of bank introductions — is what an outside observer would expect at Anysphere's scale and revenue trajectory, although none of it is public-disclosed and this part of the prediction is qualitative.

Anthropic via Claude Code is the candidate the consensus narrative reaches for first. We think it is the wrong candidate. Anthropic is plausibly preparing for a public-equity event during late 2026 or 2027 — Wilson Sonsini representation was reported by the Financial Times in early December 2025 — but the disclosure surface in any Anthropic S-1 will not be a coding-agent margin in isolation. The Anthropic filing is consequential for the category — it will disclose the margin profile of the underlying inference infrastructure, which is the input to every coding-agent vendor's COGS — but it will not be the coding-agent margin reset itself. The reset requires a pure-play disclosure.

We hold the view that the trigger event lands in the third quarter of 2026: an Anysphere (Cursor) S-1 filed during Q3, with the road-show during Q4, and the listing during Q4 2026 or the early weeks of Q1 2027. The contingent reading is a one-quarter delay if the macro window deteriorates and the filing slips. The earlier reading — a Q2 2026 filing — is also possible but not what we currently weight as the central case.

The disclosed margin and the response sequence

The disclosed margin in that filing will be the data point that matters. We expect it to land in the 50-60 per cent range on a properly-loaded gross-margin calculation that includes the inference-cost component at non-discounted commercial rates. That is materially below the 75-85 per cent range that a SaaS comparator would price into the equity — a 15-25 point gap that the disclosure event prices in (Figure 2). The first read on the filing will be the margin number; the second, more consequential read will be the rate of margin compression or expansion implied by the cohort-level usage trajectory disclosed in the same document.

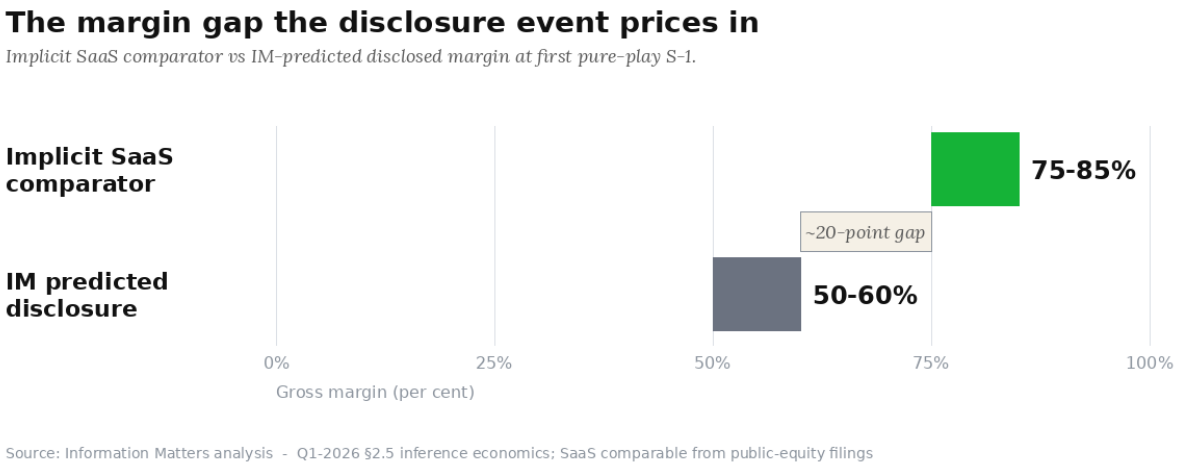

Figure 2. The margin gap. SaaS comparator vs IM-predicted disclosed margin.

The response across the rest of the sector will be ordered. The first wave, in the days after the filing, is the cross-comparable analyst response: read-throughs to GitHub Copilot's implicit margin, to the Claude Code revenue contribution inside Anthropic's then-current disclosure, to the smaller pure-plays. The second wave, over the following quarter, is the procurement response: enterprise buyers re-evaluating the long-term price trajectory of their coding-agent contracts. The third wave, into 2027, is the strategic response: which of the foundation-model providers buys which of the coding-agent specialists, and at what revaluation against the cohort's pre-trigger comparables.

WHAT HAPPENS NEXT

Three structural outcomes through 2027

The window for an independent coding-agent vendor to establish a defensible position is narrowing. We put it at nine to twelve months from the date of this report. After the trigger, the window for the second cohort closes — not because the products become bad, but because the equity-funding and acquisition multiples that allowed those companies to operate at sub-scale revenue will compress against the disclosed margin profile of their better-funded peers.

- **The specialists survive.** Companies that have chosen a vertical, a workflow layer, or a deployment posture outside the autocomplete primitive operate independently into 2028. The revenue ceiling for survivors is lower than the headline cohort but the margin profile is structurally better.
- **The middle is repriced or absorbed.** The autocomplete and generic-agent layer consolidates into Microsoft via GitHub, Google, the foundation-model providers' first-party coding offerings, and JetBrains for the IDE-resident audience. Several of the names in our second cohort are acquired during 2027, on terms set against the post-disclosure comparable rather than the pre-disclosure one.
- **The headline names diverge.** Cursor and Cognition do not converge. We expect Cursor to remain independent through 2027 with margin pressure visible in subsequent filings, and Cognition to either find a strategic acquirer or continue private at higher capital-cost burn; the former is the more likely on our reading.

OUR VIEW

The contrarian read of this report is not that coding agents are going to fail. It is that the category is going to bifurcate, that the mechanism of bifurcation is the public disclosure of gross margin, and that the timing is closer than the equity-market positioning currently implies. The single highest-signal observation point in 2026 is the Q3 disclosure window. If a major independent coding-agent vendor files an S-1 between July and October 2026, the prediction in this report is testable on the day the filing is made public. If the filing slips to 2027, the prediction is testable a quarter later. The position holds in either case; what changes is when the market reprices, not whether.

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 over 500 companies engaged in the agentic AI space, of which 83 sit in the coding-agent subset that grounds this report. The disclosed-margin range cited in this report (50-60 per cent) is reasoned from the inference-economics evidence in our Q1-2026 Agentic AI Market Report (§2.5) and the public price-curve disclosures from Anthropic and OpenAI; it is not sourced from a primary disclosure by a coding-agent vendor. 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.