

Anysphere (Cursor): The fastest reported revenue ramp in software, built on a margin and a moat it does not own — and now a SpaceX equity story, not a SaaS one.

1. Executive Summary & Verdict

Verdict: PROCEED-WITH-CONDITIONS — and only as a pre-close, structure-aware position. As of June 26, 2026, Anysphere is no longer a clean standalone-SaaS investment: SpaceX signed a definitive agreement on June 16, 2026 to acquire the company for a reported \$60B in all-stock consideration, expected to close Q3 2026 [company-overview-8] [company-overview-9]. Any position taken now is, in substance, a bet on (a) the deal closing on the stated terms, (b) where your entry point sits in the preference stack, and (c) the liquidity and value of SpaceX stock — **not** on Cursor's independent unit economics.

The three load-bearing reasons for PROCEED-WITH-CONDITIONS rather than an unqualified PROCEED:

- 1. The revenue is real and historically exceptional; the gross margin underneath it is not.** ARR moved from ~\$100M (Jan 2025) to ~\$2B (Feb 2026) [traction-3][quant-3]. But independent estimates put 2025 gross margin near **-30%**, with one firm reporting ~\$650M in annual Anthropic spend against ~\$500M revenue [market-7][adversarial-5]. The company reached only "slight" gross-margin profitability by April 2026, and only on enterprise accounts; individual-developer plans remain loss-making [traction-5][quant-4].
- 2. Cursor's two largest model suppliers are also its two most credible existential competitors** — and the disintermediation trigger is *coupled to the margin fix* (see §6, the central insight of this report). Both suppliers have demonstrated willingness to weaponize API access (Anthropic cut off OpenAI Aug 2025, Windsurf June 2025, xAI-via-Cursor Jan 2026) [competitive-3][competitive-4][competitive-5]. This is structural intermediation risk, not a hypothetical.
- 3. The verdict is gated on the deal, not the business.** A pre-close secondary buyer must underwrite the equity waterfall, the \$10B break-up fee, the stock-vs-cash structure, lock-ups, and the regulatory close — not SaaS multiples (see §8.E).

The bull case, stated as crisply as the bear case (so the PROCEED half is as defensible as the WITH-CONDITIONS half):

- **Category optionality is enormous and Cursor sits at its center.** If ASP per developer rises 5–10× as agentic coding displaces autocomplete, the TAM doubles or triples, and Cursor owns the highest-engagement surface in the category [market-4][market-6].
- **Composer is a credible margin fix with early evidence.** The April 2026 swing to "slight" gross-margin profitability is attributed to Composer absorbing inference volume and cheaper open-weight routing — i.e., the thesis is already partially proven, not purely aspirational [traction-5][quant-4].
- **The ramp and the mix are both genuinely best-in-class.** Reported \$100M→\$2B in ~13 months with enterprise share rising 25%→60% is, if even directionally accurate, the strongest demand signal in application-layer software [traction-1][traction-5].

What would flip the verdict:

- **To PROCEED (unqualified):** audited financials showing a credible path to 50%+ blended gross margin, driven by a verifiable Composer substitution rate (proprietary model handling the majority of inference tokens), plus disclosed enterprise NRR >120% and top-10-logo concentration under ~25% of ARR.
- **To HOLD/PASS:** the SpaceX deal collapses (re-pricing risk on a business losing money at the individual tier), OR Anthropic/OpenAI materially reprice or restrict API access before Composer can absorb the volume, OR audited gross margin proves structurally worse than the "slight profitability" framing implies, OR the equity waterfall shows your entry point is materially subordinated under stacked preferences.

This report is analytical opinion on public information. There is **no credible allegation of fraud or wrongdoing** here, and none is made or implied.

2. Company Overview

Legal entity: Anysphere, Inc., a Delaware C-corporation, HQ at 2261 Market Street, San Francisco; trades as Cursor [company-overview-1].

Founding: Incorporated January 2022 by four MIT computer-science students — **Michael Truell (CEO)**, **Sualeh Asif (CPO)**, **Arvid Lunnemark**, and **Aman Sanger** [company-overview-2][company-overview-3]. The Cursor editor launched publicly in March 2023; the team graduated the OpenAI Startup Fund accelerator that year [company-overview-2].

Funding trajectory (all figures company-disclosed or reporter-attributed; no audited financials exist — treat as indicative). The **Amount** column is the capital *raised* in the round; the **Post-Money** column is the resulting valuation — they are not the same number and should not be read together:

Round	Date	Amount Raised	Lead(s)	Post-Money	Source
Pre-Seed	Apr 2022	\$0.4M	—	—	[company-overview-2]
Seed	Oct 2023	\$8M	OpenAI Startup Fund	n/d	[company-overview-2] [company-overview-4]
Series A	Aug 2024	\$60M+	a16z + Thrive	~\$400M	[company-overview-5]
Series B	Dec 2024	\$105M	Thrive	~\$2.6B	[company-overview-4] [company-overview-6]
Series C	Jun 2025	\$900M	Thrive	\$9.9B	[company-overview-4] [company-overview-7]
Series D	Nov 2025	\$2.3B	Accel + Coatue	\$29.3B	[company-overview-3]
Pending acquisition	Jun 16, 2026	\$60B all-stock (consideration, not a raise)	SpaceX	\$60B implied	[company-overview-8] [company-overview-9]

Note: round labeling diverges across sources (some list a Benchmark-led Series B at ~\$400M; others fold it into Series A). The valuation waypoints — ~\$400M → \$2.6B → \$9.9B → \$29.3B — are consistent across sources even where round names are not. Total venture raised pre-acquisition is approximately **\$3.3B** [traction-5]. Named cap-table participants include OpenAI Startup Fund, a16z, Thrive, Accel, Coatue, DST Global, NVIDIA, Google, Battery Ventures, and Patrick Collison (angel) [company-overview-3][company-overview-4][company-overview-7].

The valuation step-up is itself a flag: the Series A post-money of ~\$400M (Aug 2024) → \$29.3B post-money (Nov 2025) is a ~73× re-rating in ~15 months on company-controlled information [company-overview-3]. (To be explicit: $\$29.3B \div \$400M = 73.25\times$; the \$400M is a *valuation*, not the size of any round.) That is a growth-option price, not an earnings price (see §8).

Headcount: ~60 (Mar 2025) → 300+ at Series D (Nov 2025) → ~300–500 (estimates vary) by Apr 2026 [company-overview-2][company-overview-3][company-overview-12].

3. Market & TAM

The AI code-tools category is real, compounding, and growing faster than general enterprise software. The most defensible third-party sizing (Mordor) puts the narrow "AI code tools" market at **\$7.37B (2025) → \$9.35B (2026) → \$29.96B by 2031 at 26.2% CAGR**, with Research and Markets directionally consistent at ~23.7% CAGR [market-1][market-2]. Broader agentic/enterprise-inclusive aggregations cite ~\$12.8B in 2026 revenue, roughly double 2024 — treat the wider figure as directional, not auditable [market-3].

TAM vs. SAM reality: Global developer headcount is roughly 28–32M. At \$20/month/seat, a fully-penetrated individual-developer SAM is ~\$7B/year — consistent with the research consensus. Enterprise licensing at 2–5× per seat pushes plausible SAM to ~\$15–25B [market-3A-inference]. Growth beyond that requires either usage-based agentic price expansion or category expansion to non-developer workers. Both are plausible; neither is certain.

Durable tailwinds: a structural developer-supply gap; foundation models crossing capability thresholds (92% HumanEval in 2024) that moved AI coding from sidebar to default infrastructure; and the agentic shift — Anthropic's Claude Code alone reportedly contributed ~4% of all public GitHub commits (135,000/day) by early 2026, with SemiAnalysis projecting >20% by end-2026 [market-4][market-6]. *Inference: the agentic shift is the real TAM optionality — if ASP per developer rises 5–10× as autocomplete commoditizes, the category doubles or triples. The bull case lives here, but it depends on inference costs falling, which is not yet empirically certain.*

4. Product & Technology

Cursor is a **VS Code fork**, not a plugin — a distinction that matters. Forking gives Anysphere editor-substrate control that the VS Code extension API denies to GitHub Copilot: the **Shadow Workspace** (a hidden environment that pre-tests AI output against linters before the user sees it), inline multi-file diff views, and deep completion-state integration [product-1][product-2]. Three proprietary layers exist, with **materially different moat depth:**

1. Tab completion model (strongest moat). Since March 2024, autocomplete runs a **custom sparse model** trained to predict *edits* (git-diff-like) rather than text — not a frontier-LLM call [product-5]. Context window expanded to 13,000 tokens ("Fusion"); median latency 260ms (from 475ms); it predicts "jumps" to the next edit point [product-5]. The methodology was reinforced by the **Supermaven acquisition (Nov 2024)** [product-6][product-7]. Company claims of "> 1B edited characters/day" are self-reported and unverifiable [product-5]. *Inference: the real-usage data flywheel here is the deepest, hardest-to-replicate asset in the stack.*

2. Codebase retrieval/indexing (moderate moat). AST-aware chunking via tree-sitter, a **custom code-trained embedding model**, Merkle-tree incremental re-indexing, and Turbopuffer vector storage with client-side path obfuscation [product-8]. The trained embedding weights and engineering iterations are proprietary; the architecture (tree-sitter,

cosine search, commodity vector store) is replicable in months by a well-resourced competitor [product-8][product-2].

3. Composer / agent model (emerging, high-trajectory, unproven). Composer launched in Cursor 2.0 (Oct 2025) as the first in-house frontier-class model; Composer 2 (Mar 2026) is built on a **continued-pretraining + RL recipe over the Qwen 2.5 Coder open-weight base**, claiming 250 tokens/sec and SWE-bench Multilingual 73.7 (self-reported, "CursorBench") [product-9][product-10]. SpaceX-linked training-compute references are **unconfirmed — treat as rumor** [product-10]. This is the strategic crux: in-house models attack the cost-of-revenue problem that is the company's defining weakness.

IP / provenance flag (buyer will raise this). Composer's open-weight base is **Qwen 2.5 Coder — an Alibaba-origin model**. For a company about to be owned by a US defense-adjacent entity (SpaceX/xAI), a continued-pretrain on a Chinese open-weight foundation carries non-trivial **procurement, export-control, and provenance** questions, plus Qwen license-term obligations on derivatives. This is not a disqualifier — open-weight base models are standard practice — but it is exactly the kind of supply-chain provenance item a defense-adjacent acquirer's counsel will diligence, and it is unaddressed in public materials [product-9][product-10].

Moat verdict: *real but narrow* in the Tab model; *wide but decaying* in UX-iteration lead (VS Code Agent Mode went GA March 2026 [product-2]); *emerging but unproven* in Composer. The VS Code fork itself is **not** a moat — it is a commodity starting point Microsoft can out-integrate natively. Cursor's structural vulnerability: Anthropic, OpenAI, and Microsoft are all building editor-native agent surfaces that remove the need for a third-party fork — and Cursor currently *pays* its largest competitive threats to exist [product-2][competitive-1].

5. Team & Founders

A technically elite, very young team running at the edge of its operating experience.

- **Michael Truell (CEO, ~25 at Series D).** Coding from age 11; Google intern at 18 working on language models for feed ranking; backed early by Ali Partovi/Neo Scholars [team-3]. Technical credibility is real; he has never run a company through a downturn, a platform shift, or an at-scale enterprise sales motion [team-3].
- **Sualeh Asif (CPO).** Three-time IMO representative for Pakistan (bronze); competitive-programming pedigree directly relevant to code generation [team-4].
- **Aman Sanger.** Broadest commercial exposure (Google, Bridgewater, You.com; ran an AI consultancy pre-Cursor); role variously reported as COO [team-4].
- **Arvid Lunnemark — departed October 2025**, weeks before the \$2.3B Series D, to found Integrous Research (AI safety) [team-5]. **Yellow flag:** the timing (co-founder exit immediately preceding the largest-ever financing) warrants understanding departure terms and any IP/equity issues. The stated mission-driven reason is coherent; no public indication of conflict [team-5].

Operator maturity / org signals: Non-founder COO Jordan Topoleski (2024, thin public profile); Head of Security Travis McPeak hired only July 2025 — **late by enterprise-SaaS standards** for a company already at ~\$500M ARR and pushing into Fortune 500 [team-7] [team-8][team-9]. Engineering-first, no-PM, in-person-only (SF/North Beach) culture drives extraordinary velocity but carries prioritization/documentation risk as enterprise complexity grows [team-11]. Glassdoor ~4.7/5 (comp scored lowest at 4.2) [team-11].

Key-person risk (most material team item): Truell concentration is high (public face, primary spokesperson, first company). The remaining three founders give reasonable functional coverage, but there is **no visible seasoned-executive "adult supervision" hire** from a scaled software company [team-3]. Under pending SpaceX/xAI ownership, retention of an SF-based, in-person team is an open question given SpaceX's Texas/California structure [team-11][team-13].

Governance: Below average for a \$29B company. Only Sameer Gandhi (Accel) is a publicly named board member; Thrive/a16z seats are inference [team-12]. **Investor-competitor conflict:** Google and NVIDIA are both Series D investors *and* suppliers/potential competitors (Gemini Code Assist; GPUs) — they may hold visibility into Cursor's unit economics and model-consumption data [team-6][adversarial-9]. A sharp buyer would probe information-rights structure.

Conditional related-party flag (promoted from Q-list). If the unconfirmed SpaceX/Colossus training-compute arrangement was *already live pre-deal*, the acquisition is not arm's-length price discovery — it is a vendor buying a captive customer, which materially changes the fairness-opinion lens and the credibility of the \$60B basis [product-10][team-13]. **Currently unconfirmed and therefore tagged conditional**, but it is the single related-party knot a buyer's counsel must untie before relying on the headline valuation. See Risk Register #9 and Q11.#6.

6. Competitive Landscape

The foundation-lab disintermediation risk is the primary structural risk — not a secondary concern. Cursor pays retail API rates to Anthropic and OpenAI, which then sell directly to developers at margins Cursor cannot match. One Fortune-cited VC: "*burning \$1 to make 90 cents isn't a business*"; another described Anthropic's posture as trying "*to drown out Cursor*" [competitive-1].

The central insight of this report — the kill-condition is coupled to the margin fix. The obvious question a sharp reader asks is: *if Anthropic can disintermediate Cursor, why hasn't it already?* The answer is the equilibrium that makes this thesis precise. Cursor is currently **Anthropic's single largest API customer (~\$650M/year)** [market-7][adversarial-5]. Anthropic's rational move today is to *monetize* the threat, not starve it — milk the largest check it has while building Claude Code alongside. **That equilibrium inverts the day Composer's substitution rate makes Cursor a *shrinking* Anthropic customer.** At that moment Anthropic's incentive flips from "milk" to "starve," because Cursor is no longer worth

more as a customer than as a corpse. **The implication is the sharp, non-obvious part: disintermediation accelerates precisely when the margin fix works.** The two are not separate risks — they are the same mechanism viewed from two ends. Every dollar of inference Composer wins back from Anthropic improves gross margin *and* moves Cursor toward the threshold where its largest supplier's incentive to cut it off turns on. A buyer cannot underwrite the margin recovery and the supplier relationship independently; they are one coupled variable.

Ranked by threat:

1. **Anthropic / Claude Code (Critical).** Terminal-native agent launched Feb 2025; reported ~\$2.5B run rate and 300,000+ business customers by early 2026 — **single-sourced [competitive-1], and carries the same ±30% skepticism band applied to Cursor's own figures** (the competitor's numbers get no free pass). A JetBrains April 2026 survey shows Claude Code "most-loved" at **46%** vs. Cursor **19%** vs. Copilot 9% [competitive-2] [adversarial-3]. Workplace adoption grew 6× (3%→18%) in ~nine months [competitive-2]. Critically, **Anthropic has repeatedly severed access** — OpenAI (Aug 2025), Windsurf (June 2025), xAI-via-Cursor (Jan 2026) [competitive-3][competitive-4][competitive-5]. Cursor President Oskar Schulz conceded: "*The IDE isn't the right form factor anymore for a world where you can produce 10 times more code*" [competitive-1].
2. **OpenAI / Codex CLI (High).** Open-source terminal agent on GPT-5.5 reported #1 on Terminal-Bench 2.1 at **83.4%** vs. Claude Code 78.9% [competitive-6]. **Treat as a single-sourced, self-/vendor-reported benchmark and band it accordingly** — benchmark leaderboards are marketing surfaces and we apply the same skepticism we apply to Composer's self-reported CursorBench number; we do not present one as fact and the other as claim. OpenAI also holds an IDE asset via Windsurf.
3. **GitHub Copilot / Microsoft (High in enterprise).** 4.7M+ paid subscribers, ~42% market share by headcount, ~90% Fortune 100 penetration, pre-approved in most Fortune 500 procurement [competitive-7][competitive-8][competitive-2]. **Weakness:** developer satisfaction is poor (9% most-loved) — wins on procurement, not preference [competitive-2]. *Inference: a distribution moat Cursor cannot outspend, but the quality gap keeps Copilot from displacing Cursor in developer-led orgs.*
4. **Windsurf (Medium, evolving).** ~\$82M ARR then stalled; acquisition terms **explicitly contested** across sources (\$3B OpenAI vs. \$250M Cognition) — do not rely on either figure [competitive-10][competitive-11][competitive-12]. Now ships SWE-1.5 and undercuts at \$15/month [competitive-12].

Structural read: Two of the four leading players have been acquired by foundation labs within 12 months. The "who wins the independent slot?" question may already be answered — and Cursor's own answer was to sell.

7. Traction & Financial Signals

ARR trajectory (recency-tagged; every figure company-supplied or analyst-estimated — none audited). ARR-basis note for the valuation math: this report uses the **May 2026 ~\$4B run-rate (Sacra)** as the single denominator for all current EV/ARR calculations in §8.D, footnoted there; the Dec 2026 >\$6B figure is company guidance and is discounted, not used for valuation.

Date	ARR	Source	Reliability
Jan 2025	~\$100M	Sacra estimate [traction-1]	Secondary (analyst)
Jun 2025	>\$500M	Company at Series C [traction-2]	Primary, unaudited
Nov 2025	>\$1B	Company at Series D [traction-3]	Primary, unaudited
Feb 2026	~\$2B	Fundraising press [traction-4]	Primary, unaudited
May 2026	~\$4B run-rate	Sacra [traction-5]	Secondary — used as §8.D denominator
Dec 2026 (guidance)	>\$6B	Company guidance [adversarial-9]	Promotional — discounted, not used for valuation

\$100M → \$2B in ~13 months implies ~20× YoY; the company described revenue "doubling ~every two months" through mid-2025 [traction-1][traction-2]. If accurate, this would be the **reported fastest** \$1M→\$500M B2B ramp, beating Wiz/Deel/Ramp by 6–8 months — a superlative resting on **unaudited self-reported ARR**, so we state it once, here, and not as fact [traction-6]. *Inference: growth is decelerating off a larger base (5× in H1 2025, then ~4× over the following ~9 months) — still extraordinary, but watch for the plateau.*

Customer mix: ~360,000 individual paying users at ~\$276 ACV (Mar 2025) [traction-1]; 50,000+ engineering teams (Dec 2025) [traction-3]; ~64–70% Fortune 1000 penetration (company-disclosed/Bloomberg-via-Sacra) [traction-5][traction-8]; named logos include NVIDIA, Uber, Adobe, Salesforce, PwC, Stripe, Ramp, Brex [traction-5]. **Enterprise revenue share rose from ~25% (late 2024) to ~60% (early 2026)** — a structurally healthier, higher-ACV mix [traction-5].

Concentration / revenue-quality caveat (the asterisk on "enterprise = durable"). A rising enterprise mix is healthier *only if it is not concentrated*. **Top-N logo concentration is undisclosed and is a first-order unknown:** one lost seven-figure logo ≠ one churned \$20/mo developer, so the enterprise shift *changes the shape of churn risk rather than removing it*. Compounding this, several named logos (**NVIDIA, Google**) are also **investors**, raising a relatedness question — are they arm's-length customers or strategically-motivated

ones whose spend would not survive a relationship change? Until top-10-ARR concentration is disclosed, "enterprise = durable" is asserted, not established [traction-5][adversarial-9]. (See Q10 and missing-data register #9.)

Retention — the critical unknown, with a defensible proxy. No NRR has been disclosed; Contrary Research noted (late 2025) that "retention and long-term engagement have not been validated" [traction-9]. But a buyer pays for triangulation: freemium→paid conversion of ~36% is exceptional [traction-1], yet the **June 2025 pricing backlash and documented individual-tier churn to Windsurf** landed *precisely as enterprise mix was rising* [competitive-13][quant-2]. *Inference: individual gross retention likely took a hit in H2 2025 while enterprise expansion masked it — i.e., a blended NRR that looks acceptable may be enterprise expansion papering over individual-tier contraction. The blended number, if disclosed, would understate the divergence between the two cohorts.* This is a defensible read from sourced facts, not a claim of fact.

Cash: ~\$1B reserves with "single-digit monthly cash burn" as of Dec 2025 — comfortable runway independent of the deal [traction-9].

8. Quantitative Read & Unit Economics

This is where a sharp buyer earns or loses money. [Show the math.](#)

A. Per-seat economics vs. model COGS — with a sensitivity table, because one assumption drives the whole conclusion. Claude Sonnet pricing is ~\$3/\$15 per 1M tokens (input/output) [quant-1]; a single long-context agentic session can cost **\$0.50–\$2.00** in pure API fees [quant-1]. The linchpin is requests/month per active Pro developer. **The ~400 figure below is an illustrative analyst assumption [quant-2], not a sourced datapoint — so here is the sensitivity, which is the actual arithmetic:**

Pro tier (\$20/mo)	@ \$0.08/request	@ \$0.15/request
150 req/mo	COGS \$12 → +40% margin	COGS \$22 → -12% margin
250 req/mo (≈breakeven)	COGS \$20 → ~0%	COGS \$38 → -88%
400 req/mo (base assumption)	COGS \$32 → -60%	COGS \$60 → -200%

The *direction* is robust — the Pro tier only clears breakeven in the low-usage / low-cost-per-request corner ($\leq \sim 150$ req at the cheap end), and goes deeply negative for any genuinely heavy agentic user. This is corroborated by the independent $\sim -30\%$ blended estimate (§8.B). Light read: **the median engaged Pro user is plausibly unprofitable, and the heavy user is structurally so** [quant-2].

- *Business tier* ($\sim \$40/mo$), *same usage*: still negative at 400 req, but materially better and breakeven-positive across more of the grid.
- *Ultra tier* ($\sim \$200/mo$): positive at typical usage, more so if the user runs Composer rather than Claude [quant-2].

This arithmetic is the June 2025 pricing controversy. Cursor swapped fixed "500 fast requests" for a credit pool at API rates; effective delivered compute fell ~55% ($\approx 500 \rightarrow \approx 225$ Sonnet-equivalent requests at the same \$20), effective price rose **20x+** for heavy agentic users, one HN user reported \$350 overage in a week, and CEO Truell apologized and issued refunds (July 4, 2025) [quant-2][competitive-13][traction-pricing]. **This was a margin-defense move, not a comms failure** — selling agentic AI at autocomplete prices was never sustainable [market-8][competitive-4].

B. Blended gross margin. Independent estimate (Fundamental, mid-2025): ~\$650M Anthropic spend vs. ~\$500M revenue $\rightarrow \approx$ **-30% gross margin** [market-7][adversarial-5]. A Newcomer source claimed Cursor was "spending 100% of its revenue on Anthropic" [adversarial-4]. **(Both are critical, single-source, interested-adjacent estimates and carry the $\pm 30\%$ band.)** By April 2026, Sacra reports only **"slight" gross-margin profitability** — enterprise positive, individual-tier still loss-making — attributed to Composer absorbing inference volume and cheaper open-weight routing (e.g., Kimi K2) [traction-5][quant-4]. *This remains ~30+ points below the 70–80% SaaS baseline* [adversarial-8].

C. Revenue per employee vs. named peers (Cursor at ~\$1B ARR / ~300 employees, **Nov 2025** basis; comps tagged with vintage):

Company	ARR (as-of)	Employees (as-of)	Rev/Employee
Cursor	~\$1B (Nov 2025)	~300 (Nov 2025)	~\$3.3M
Stripe	~\$20B (2024 est.)	~8,000 (2024)	~\$2.5M
Figma	~\$600M (2023)	~2,100 (2023)	~\$286K
Median public SaaS	— (2024–25)	—	~\$395K

[quant-rev-emp][quant-8]. Exceptional — **but misleading for an inference-heavy business, and the vintage caveat matters** (Figma's figure is a 2023 snapshot; do not read it as contemporaneous). The deeper point: the "employees" doing much of the work are Anthropic's and OpenAI's GPUs, whose cost sits in **COGS, not headcount**. *Revenue/employee flatters any company that rents its core input — it is a vanity metric here, not a strength.*

D. Valuation/ARR multiple vs. comps (single ARR basis per round, footnoted; the SpaceX line uses the **May 2026 ~\$4B run-rate** from §7, not the discounted Dec guidance):

Round	Valuation	ARR (basis)	EV/ARR
Series C (Jun 2025)	\$9.9B	~\$500M (Jun 2025 co.)	~20x
Series D (Nov 2025)	\$29.3B	~\$1B (Nov 2025 co.)	~29x
Apr 2026 talks	~\$50B	~\$2B (Feb 2026)	~25x
SpaceX (Jun 2026)	\$60B	~\$4B (May 2026 run-rate, Sacra)	~15x

High-growth public SaaS comps trade at ~~8–15× forward ARR~~ [quant-5]. The SpaceX print lands at the ~~**top of that band (15×)**~~ — richer than it looks only if you accept the run-rate denominator; on a conservative ~\$3B basis it is ~20×, i.e., 1.3–2.5× above the comp midpoint [quant-5]. To justify the **\$29.3B** Series D on fundamentals (≈7% discount, 25× terminal), an analyst needs ~\$5–6B ARR within five years at 60%+ gross margins — requiring three things *simultaneously* true: 80%+ annual ARR growth through 2027, gross margin expanding from ~0–30% to 60%+ (i.e., Composer displacing the majority of third-party inference), and the IDE remaining the primary developer surface against CI/CD-native agents [quant-5][adversarial-2]. *No software business has sustained the first at this scale for five years; the second is credible but unproven; the third is actively contested.*

E. Deal structure & entry economics — the actual deal a secondary buyer underwrites.

§1 names the deal as the whole thesis, so the structure math cannot be left in the missing-data register. The framing (mechanics; exact terms are undisclosed and **must be confirmed against the definitive agreement** — see Q11.#7):

- **All-stock means conversion, not cash.** Common and preferred convert to SpaceX shares at a fixed exchange ratio. The **\$60B / \$29.3B = ~2.04× step-up** off the last round looks clean *at the headline* — but the headline is the pre-preference number. What each holder actually receives depends entirely on the preference stack, which is undisclosed.
- **Where the preferences sit determines common-holder economics.** Total venture raised is \$3.3B [traction-5]. If the stack is ~~1× non-participating~~ (founder-friendly, common case in hot rounds), prefs take ~~~\$3.3B off the top~~ — ~~**5.5%~~ of \$60B** — and ~94.5% flows pro-rata to common and the option pool; the 2.04× step-up is roughly real for common. If instead late rounds carry **participating preferences, multiple liquidation prefs, or senior stacked seniority** (plausible given the \$29.3B Series D was a defensive, late-cycle round), late investors double-dip and **common-holder step-up compresses materially below 2.04×** — potentially to near-zero for early common in a downside re-pricing. **This is the single most important number a pre-close secondary buyer needs and it is not public.**
- **Bridge to the break-up fee (\$9.4).** The **\$10B break-up fee (~17% of deal value)** is unusually large; its size is the clearest market signal that *both sides* see a non-trivial probability the deal does not close — which is precisely why a secondary position is a bet on close, not on Cursor's P&L.

Inference, explicitly flagged: the entire "2.04× clean step-up" reading collapses or holds on preference terms no public source discloses. A buyer who underwrites the headline without the waterfall is underwriting a number that may not be theirs.

9. Risk Register & Red Flags (ranked)

#	Risk	Severity	Confidence	Evidence
1	Structural gross-margin compression — money-losing at the individual tier; ~-30% blended in 2025; only "slight" profitability by Apr 2026	Critical	High	[market-7] [adversarial-5] [traction-5] [quant-4]
2	Supplier-as-competitor / API weaponization, coupled to the margin fix — Anthropic & OpenAI compete directly and have severed access before; the disintermediation trigger <i>fires when Composer succeeds</i> (§6)	High	High	[competitive-1] [competitive-3] [competitive-4] [competitive-5]
3	No audited financials — every ARR, margin, and valuation figure is company-controlled or analyst- estimated	High	High	all lanes
4	Deal-execution / structure risk — \$60B all-stock, \$10B break-up fee (~17%), Q3 close, illiquid SpaceX paper, undisclosed preference waterfall (§8.E)	High	High	[company- overview-8] [company- overview-9] [traction-5]
5	Undisclosed retention/NRR + customer concentration — enterprise cohort expansion is the dominant unknown and explicitly "not validated"; top-10-logo concentration undisclosed; NVIDIA/Google are investor-customers	High	Medium	[traction-9] [traction-5] [adversarial-9]
6	Moat shallowness — VS Code fork is commodity; native VS Code Agent Mode GA Mar 2026	Medium- High	Medium	[product-2] [adversarial-7]
7	Pricing-trust damage — June 2025 backlash, documented churn to Windsurf, fragile developer trust	Medium	High	[competitive-13] [quant-2] [traction- pricing]
8	Key-person / governance immaturity — Truell concentration, co-founder exit pre-Series D, thin board, investor- competitor conflicts	Medium	Medium	[team-3][team- 5][team-6] [team-12]

#	Risk	Severity	Confidence	Evidence
9	Conditional related-party / fairness-opinion risk — <i>if</i> SpaceX/Colossus compute was live pre-deal, the acquisition is a vendor buying a captive customer, undermining the arm's-length basis of the \$60B (currently unconfirmed)	Medium (conditional)	Low– Medium	[product-10] [team-13]
10	Composer IP / provenance & export-control — built on Alibaba-origin Qwen 2.5 Coder open-weight base; provenance/license/export questions under defense-adjacent ownership	Medium	Medium	[product-9] [product-10]
11	Regulatory / antitrust (HSR/CMA second-request) — \$60B all-stock takeout of the leading independent coding tool by an AI-footprint acquirer, where two suppliers are also competitors; the \$10B break-up fee is sized to this risk	Medium	Medium	[company-overview-8] [company-overview-9] [competitive-1]

No fraud/wrongdoing flag exists. The pricing controversy was an acknowledged, refunded communication-and-economics misstep — a governance signal, not misconduct [traction-pricing][team-14].

10. Key Questions for Management

- Gross margin by tier and usage percentile.** What is contribution margin on Pro vs. Business vs. Ultra at P50 and P90 usage? The aggregate "slight profitability" almost certainly masks a deeply bimodal distribution (see §8.A sensitivity).
- Proprietary-model substitution rate.** What % of *total inference tokens* run on Composer today vs. retail Anthropic/OpenAI APIs? This single number determines both the margin thesis *and* the timing of the Anthropic disintermediation trigger (§6) — they are the same variable.
- Model-vendor contracts.** Do you have contractual pricing protection, volume floors, or access guarantees from Anthropic and OpenAI — or do you pay rack-rate that can be changed or revoked unilaterally? Show the contracts.
- Net revenue retention, by cohort.** What was NRR in Q3 2025 (post-pricing-change) for the individual tier *separately* from enterprise? Quantify the documented churn to Windsurf. (We infer individual contraction was masked by enterprise expansion — confirm or refute.)

5. **Customer concentration.** What share of ARR is the top 10 logos? What share is held by entities that are *also investors* (NVIDIA, Google)? What is the contractual term and auto-renew structure on the top 10?
6. **Audited TTM GAAP revenue + AR aging.** Reconcile "ARR" against trailing-twelve-month GAAP revenue and an accounts-receivable aging schedule. How much "ARR" is front-loaded enterprise commitment vs. recurring consumption?
7. **Composer training compute & provenance.** Are SpaceX/Colossus GPU arrangements real, and were they live *before* the deal signed? What are the terms and the cost-per-token delta vs. retail API? What is the export-control and license posture on the Qwen 2.5 Coder base?
8. **Deal mechanics & waterfall.** Provide the full liquidation-preference stack by round (participation, multiple, seniority), the SpaceX-share exchange ratio, the \$10B break-up fee trigger conditions, lock-ups, regulatory closing conditions, and any MAC clauses.
9. **Founder & talent retention post-close.** What retention/vesting is in place for Truell, Asif, and Sanger, and for the SF engineering core, under SpaceX/xAI ownership?
10. **Vendor-investor information rights.** What data do Google and NVIDIA (investors *and* competitors/suppliers) receive on your unit economics and model-consumption?

11. Sources & Confidence Notes

Confidence framework. (*primary*) = company filing/official announcement/named-executive quote; (*reputable-secondary*) = TechCrunch, CNBC, Fortune, Bloomberg, The Information, The Next Web, JetBrains survey; (*market-research*) = third-party research firm; (*secondary*) = analyst/aggregator estimate; (*inference*) = analyst judgment from sourced facts. **All financial and valuation figures — Anysphere's and competitors' — are unaudited and carry a ±30% skepticism band, applied symmetrically** (Claude Code's ~\$2.5B run-rate and the Codex/Terminal-Bench leaderboard numbers get the same discount as Cursor's own figures). Several Cursor figures are produced by *interested parties* — company blogs (Series D, pricing) and investor-attributed round figures control their own information flow; the Fundamental and Newcomer margin claims are critical-but-single-source and should be corroborated. Round labeling and Windsurf deal terms conflict across sources and are flagged inline.

Key sources by lane: Cursor blog/Series D & pricing (*primary*) [company-overview-3] [traction-pricing]; TechCrunch/CNBC funding & acquisition [company-overview-7][company-overview-8][company-overview-9]; Sacra ARR/margin/enterprise-mix [traction-1][traction-5]; Fundamental/Newcomer margin estimates [adversarial-4][adversarial-5]; Fortune "Cursor's Crossroads" + JetBrains survey on competitive position [competitive-1][competitive-2]; Mordor/Research and Markets on TAM [market-1][market-2]; Contrary Research on retention gap [traction-9]; product/moat from Cursor engineering blogs + notoriousplg/Towards Data Science analysis [product-1][product-5][product-8].

Missing-data register — what a real buyer must still obtain before closing:

1. **Audited financial statements** — TTM GAAP revenue, COGS detail, cash-flow statement.
2. **Gross margin by tier and by P50/P90 usage cohort** — the single most important operating unknown.
3. **Net revenue retention**, split individual vs. enterprise cohort (currently undisclosed; we infer divergence).
4. **Model-vendor contracts** — pricing protection, access guarantees, exclusivity, term, for Anthropic and OpenAI.
5. **Proprietary-model substitution rate** (Composer % of total inference tokens) and its cost-per-token delta — drives both margin and the disintermediation trigger.
6. **Definitive merger agreement** — break-up fee triggers, lock-ups, regulatory conditions, MAC clauses.
7. **Cap table & liquidation-preference waterfall** — participation, multiples, seniority by round, SpaceX exchange ratio, and the implied common-holder step-up under the \$60B all-stock structure (§8.E).
8. **Cohort retention curves** and CAC/payback by channel (individual vs. enterprise).
9. **Top-N customer concentration** — top-10-logo share of ARR, and the share held by investor-customers (NVIDIA, Google).
10. **Composer IP/provenance file** — Qwen 2.5 Coder license compliance, export-control posture, and the reality/terms of any SpaceX/Colossus training-compute arrangement (arm's-length test).

This report is illustrative, for public sample purposes only. It constitutes analytical opinion on public information and is neither investment advice nor a statement of fact regarding any wrongdoing.