



Harnessing Enterprise Value – The ROI of AI

Enterprise AI budgets are surging and pilots are everywhere. And yet AI products face an acute struggle to convert from experiment to production.

Here is the latest research on the gap between AI ambition and adoption — and what founders can do about it.

By Rolanda Fu, Matt McIlwain, Vivek Ramaswami

Findings Drawn From Three Exclusive Sources:

150

enterprise decision-makers surveyed

48

engineering leaders surveyed at Madrona's Builders Summit

5 years

of IA40 winners

Overview

For the past several years the headline has been the same: enterprise AI adoption is rapidly growing. And by the numbers, it is. But in five years of tracking the [Intelligent Applications 40](#), and after hundreds of conversations with the founders building these companies and the enterprises buying them, we've learned that the headline obscures a more complex story.

The market for enterprise AI is indeed enormous and growing. Our survey of 150 senior enterprise decision-makers confirms it: **74% plan to expand AI budgets in the next 12 months, 45% have carved out dedicated AI line items**, and virtually nobody is cutting. On the supply side, 2024 and 2025 IA40 winners have raised \$366 billion in two years — nearly a third of the \$1.3 trillion in global venture capital invested. Both the capital and buyers are committed.

And yet, 83% of those same enterprises convert fewer than half their AI pilots to production; on the flip side, only 1% of enterprises actually convert >75% of their pilots. It's not a product shortcoming: “didn't work as promised” is only the sixth most common reason a pilot fails! What kills pilots is the startup's ability to navigate an enterprise environment. Integration complexity ranks at the top, then security and compliance, then proving financial value or ROI.

Bottlenecks within the enterprise keep shifting. The workflows delivering the most value today often aren't always the ones decision-makers are targeting for the next 12–24 months. In technology and product, code generation is by far the highest-value AI workflow today; yet enterprise priorities are elsewhere in where they're hunting to invest next.

Two gaps persist: between AI ambition and execution from pilot to production, and between where enterprise value lies today vs. tomorrow. These are the defining opportunities for founders building enterprise AI products, where there remain massive opportunities in enterprise adoption of AI agents and systems of execution.

What follows aggregates three views: 1) a proprietary survey of enterprise decision-makers, 2) a practitioner survey from our AI builder community, and 3) five years of historical IA40 data — to help founders understand not just what enterprises say they want, but how and what products actually get bought, which pilots can scale into production, and where enterprises are investing real dollars.

Key findings

1 Enterprise AI budgets keep growing, while capital consolidates around a few winners

Budgets are surging: 74% of enterprises plan to expand AI spending in the next 12 months, 15% by more than 50%, and only 1% plan cuts. 45% now carry a dedicated AI budget line, with roughly 20% each reallocating from hiring or from software spend. On the supply side, capital is increasingly concentrating: the top 10 venture-backed companies took over 50% of US venture deal value in 2025, up from 19% in 2022. Among the 2024 and 2025 IA40 Winners, nearly 90% of the \$366B raised over the last two years went to just three companies — OpenAI, Anthropic, and Databricks — while the remaining 58 winners that raised still averaged over \$800 million each.

2 Pilots that make it to production drive strong end-user adoption

Most pilots never make it out: 83% of enterprises convert fewer than half of their AI pilots to production, and 35% convert fewer than one in four. The product usually works — integration complexity and security/compliance are the top two pilot killers, well ahead of “didn’t work as promised.” Strong end-user adoption is the #1 factor tipping the decision to scale, slightly outweighing executive sponsorship. When you win the users first, the business case follows.

3 Companies need to align with how enterprises actually buy AI

Discovery is bottoms-up: 41% of enterprises primarily find AI tools through internal tech-team research, and only 7% cite vendor outreach as their primary channel. In terms of pricing, buyers want outcome-based pricing (33% prefer it) but usage-based pricing is what they get (45% primarily encounter it), suggesting that usage is not yet resonating as a proxy for value. Ease of integration and security are the through-lines of the entire purchase lifecycle, mattering more in tipping the decision to scale a pilot than features of the solution itself. And 52% of deals close within 6 months, but 47% re-evaluate their stack on the same cadence: winning the deal is only the start.

4 The workforce is transforming faster than it is shrinking

83% of enterprises say AI is reshaping their workforce. 42% are rewriting job descriptions for “AI-augmented” roles with stable headcount, and 27% are slowing hiring as AI increases capacity. Expected net reductions are largely offset by expected net additions. The story is an investment in the transformation and transition of the workforce instead of mass layoffs.

5 What defines top-performing AI companies continues to evolve

The 2025 IA40 saw 68% turnover, among the highest in the list’s history — only Databricks has made all five lists. Exits are materializing among a concentrated few: Figma’s IPO (\$19.8B), CoreWeave’s IPO (\$23B), Google’s acquisition of Wiz (\$32B), and both OpenAI and Anthropic filing confidentially to go public. SpaceX’s \$1.75T IPO in June validated the path for broader market demand and the potential TAM for enterprise applications (estimated by SpaceX at \$22.7T). Infrastructure has captured ~70% of the \$14 trillion in AI market cap created since ChatGPT; but as in every prior technology cycle, application-layer value follows infrastructure to close the ROI gap by bringing workflow depth and domain expertise.

BOTTOM LINE FOR FOUNDERS

The enterprise AI market is enormous, growing, and backed by committed budgets. We are now in the era of rapid enterprise experimentation, with pilots not guarantees of longer-term partnerships. The founders who win will be the ones who build products, go-to-market motions, and implementation playbooks designed to close the delivery gap in realized value, and scale platforms that address both pain points today and anticipate those of tomorrow.

The state of enterprise AI in 2026

The capital is committed. Value is still emerging in production.

Enterprises have committed capital across the board. 74% of enterprises plan to expand AI budgets in the next 12 months, and virtually nobody is cutting. 45% have carved out dedicated AI funding — AI has graduated to a distinct spending category.


The broader capital markets reinforce this conviction: The market is seeing record infrastructure investment, frontier model labs valued in the hundreds of billions, and an ever-growing share of venture dollars flowing to AI. AI contributed 65% of all US venture deal value in 2025, and H1 2026 venture-growth activity has already exceeded twice the full-year 2025 total. The SpaceX IPO in June 2026 at roughly \$1.75 trillion validated the path for technology companies at scale, and the expected Anthropic and OpenAI IPOs will serve as direct pricing benchmarks for the AI ecosystem. These mega-IPOs could unlock a broader liquidity cycle that has been constrained for four consecutive years.

74% of enterprises plan to expand AI budgets in the next 12 months	\$366B raised by 2024–2025 IA40 winners — nearly a third of all US VC (\$1.27T)	\$1.75T SpaceX’s June 2026 IPO valuation, validating the path for mega-scale tech IPOs
--	---	--

However, enterprise budget commitments and significant venture capital dollars do not yet yield business results consistently for (or within) every enterprise. Closing the gap between experimentation and realized value is the defining dynamic of enterprise AI in 2026.

Self-reported business impact of AI, to date

N=150, single select



Five years of the IA40: what the supply side tells us

Since 2021, Madrona's Intelligent Applications 40 has tracked the most promising enterprise AI companies. Five years of data show a supply side that is broadening, accelerating, and increasingly vertical. Capital concentration at the top coexists with a vibrant ecosystem across the stack, the landscape reshuffles dramatically every year, and the categories that matter are shifting from model-centric to workflow-centric.

The landscape of dominant companies in AI is anything but stable. The 2025 IA40 featured 27 new companies compared to the prior year, a 68% turnover rate, reflecting the pace of innovation in applied AI. Of the original companies that appeared on the 2021 list, only a couple remain, with the majority displaced on the list by faster-growing newcomers, and a few that have since IPO'd or gotten acquired. Only Databricks has made the list all five years. Abnormal, Hugging Face, and Runway are the only four-time winners, but notably, none of the three made the 2025 list. The bar for what defines a top intelligent application company has never been higher.

The IA40 survivorship curve

Share of each year's cohort still on the list, by year



The categories have evolved dramatically. Code generation did not exist on the 2021 list; by 2025, Cursor (valued at nearly \$10 billion) and Factory are among the top winners. The 2025 list reflects agentic interfaces (Perplexity, HeyGen, ElevenLabs), vertical AI in legal (Legora),

healthcare (Abridge, OpenEvidence), and defense (Anduril), and a new wave of AI-native infrastructure (CrewAI, Browserbase, Langchain) enabled by protocols like MCP and A2A. The shift from model-centric to workflow-centric AI is visible in how the list composition has changed.

The IA40 has produced meaningful exits across both IPO and M&A channels:

Select IA40 exits, 2024–2025 cohorts

Company	Type	Value
 CURSOR → 	Acquisition	\$60B
 WIZ → 	Acquisition	\$32B
 CoreWeave	IPO	\$23B
 Figma	IPO	\$20B
 dbt Labs → 	Acquisition	\$2B
 STATSIG → 	Acquisition	\$1B
 NEON → 	Acquisition	\$1B

Both OpenAI and Anthropic have since filed confidentially for IPOs. Earlier cohorts produced Samsara (IPO), Moveworks (acq. ServiceNow), Weights & Biases (acq. CoreWeave), and Loom (acq. Atlassian).

WHAT IT MEANS

Enterprise capital is still flowing primarily to infrastructure, but the application layer is where deployment, workflow integration, and measurable ROI converge. As infrastructure consolidates, the application layer has more room to scale into established functions like healthcare, CRM, supply chain, and manufacturing.

Opportunities for builders: where enterprises are actually investing

Adoption across the stack

The enterprise AI stack is maturing across all layers: model, infrastructure, and application. 53% of enterprises are investing across all three simultaneously, while the remainder generally focus on two layers. This is a multi-layer market; enterprises are investing dollars in multi-layer AI architectures.

The vendor landscape: distribution vs. innovation

The broader organizational vendor landscape reveals a classic tension between the staying power of incumbents and the disruptive nature of new entrants. Microsoft dominates on distribution with high penetration among enterprises. That’s the power of bundling, existing enterprise relationships, and a salesforce that already has the buyer’s ear. But **satisfaction tells a different story: Anthropic leads NPS at 8.4, meaningfully ahead of Microsoft and OpenAI, winning disproportionate loyalty relative to its primary-vendor share.** The power of AI-native players is their agility to adapt and build to what customers demand, and lead with innovation.

Vendor	Paid customer	Primary AI vendor	NPS
Microsoft	78%	48%	7.6
OpenAI	63%	15%	7.5
Anthropic	60%	22%	8.4
Google	45%	8%	7.8
Salesforce	35%	5%	6.9
Amazon	19%	1%	7.2
Glean	8%	1%	7.0

In terms of vendor type preference: 48% of enterprises deliberately prioritize a mix of incumbents and AI-native vendors. 29% of enterprises are specifically prioritizing incumbent vendors adding AI, while 23% now prefer AI-native vendors. The market isn’t settled — there is

nearly equally strong appetite today for prioritizing riskier new AI-native vendors and AI-enhanced solutions.

Workflows by function: where AI has traction, where it’s headed

What matters most is the gap between where AI is most commonly leveraged today, which specific workflows are driving the highest value, and where enterprises plan to invest next. For simplicity, we examined workflows across three primary parent functions in an organization: technology-related functions (software, product, data), back office functions (finance, HR, legal, IT, operations), and go-to-market functions (sales, marketing, customer success).

Tech Functions

Code generation is the clear value winner today (highest-value workflow). But DevOps/infrastructure/incident response — in use relatively less today at 54%, and not sufficiently delivering value yet — is the #1 workflow teams plan to invest in next. Other workflows that are high priorities for investment include test generation & QA and code review.

Workflow	In use today	Highest value today	To invest
Data analysis & cleaning	74%	12%	52%
Product requirements drafting	70%	12%	42%
Code generation	68%	40%	50%
Prototyping / design / UI	62%	14%	52%
DevOps / infra / incident response	54%	6%	52%
Test generation & QA	60%	8%	44%
Code review	58%	2%	44%
User research analysis	42%	2%	40%
Anomaly detection	42%	4%	40%

Back Office

Usage and value are already aligning: knowledge management and reporting/forecasting lead on both axes. Interestingly, these workflows remain the biggest investment prioritizes as well, suggesting there is still significant opportunity for new capabilities. The clearest gap is internal

employee support — 66% adoption, 50% planned investment, but only 12% highest-value; support may not be used by everyone daily, but represents a clear opening for AI companies to innovate.

Workflow	In use today	Highest value today	To invest
Knowledge management / search	78%	28%	76%
Reporting / analysis / forecasting	76%	24%	54%
Internal employee support	66%	12%	50%
Content generation	64%	14%	44%
Research	60%	10%	38%
Invoice / expense management	34%	2%	38%
Candidate sourcing	26%	4%	28%
Candidate assessment	22%	2%	30%

GTM

Content creation anchors the GTM function on every axis as the most demanding workflow and a continued area for investment. Where some of the gap in value still exists today include personalized outreach and CRM data capture and automation, which draw heavy usage and investment interest despite thin realized value today — today’s tools aren’t delivering.

Workflow	In use today	Highest value today	To invest
Content creation	76%	38%	56%
Personalized outreach	66%	10%	46%
Customer segmentation	58%	24%	52%
CRM data capture / automation	58%	8%	48%
Lead generation / enrichment	50%	4%	40%
Forecasting / pipeline analysis	44%	4%	32%
AI SDRs / prospecting assistants	34%	8%	34%


Note that for "in use today" and "to invest," respondents were allowed to select all that apply, while for "highest value today," they could only select one workflow.

The builder's view: what practitioners see on the ground

Our annual Builders Summit survey of startup technology and product leaders tells the practitioner side of a similar story, with two key findings standing out:


First, AI usage across the SDLC has broadened dramatically: when asked which workflows AI is meaningfully leveraged, every category saw an increase from 2025 to 2026, with code reviews seeing the largest jump. Front-end and back-end code generation are now near-universal.

AI usage across the SDLC, year over year



Second, the bottleneck has shifted: AI is no longer concentrated in code generation, but has spread across the full development lifecycle, with active investment in new bottlenecks such as drafting specifications, code understanding, and code review.

Where bottlenecks have shifted to, by org size



This corroborates the enterprise decision maker perspective where the value frontier is moving to adjacent, harder-to-automate workflows across organizations of all sizes.


Navigating the enterprise: how the AI buying machine actually works

Discovery: the practitioner leads, the vendor follows

The enterprise AI discovery process is fundamentally bottoms-up. The “land” doesn’t happen in a boardroom; it happens when a practitioner finds your tool, proves it works, and creates internal demand. This has profound implications for how founders should allocate GTM resources, where product quality, community and brand development, and word of mouth are the new enterprise sales motions.


Primary channel for discovering AI tools

N=150, single select



Follow the budget: a new line item


AI has graduated from “innovation experiment” to “distinct budget category.” Nearly half of enterprises now have dedicated AI budget lines. But AI is drawing about equally from existing labor and software budgets today, disputing narratives around labor replacement.



In terms of who to sell to and who is making budget decisions, we see a wide mix — but most commonly, there is still a centrally held budget that rolls up to the CIO, but functional leads still play a significant role in ownership around what specific AI solutions get adopted. Builders need to have clarity around the decision maker, and ensure both their product champion and their economic buyer are bought in.

Budget owner vs. decision-maker for AI

N=150, single select




The pricing-value mismatch

Buyers want to pay for outcomes. Vendors are charging for consumption. Outcome or value-based pricing is the #1 preferred model but only the 4th most commonly encountered, while usage-based pricing is overrepresented. Buyers are now realizing that usage is not necessarily translating to outcomes, while also lacking the predictability of hybrid and seat-based models.

Pricing models — most preferred vs. common

N=150, single select



For founders, this gap is an exciting opportunity to align usage and value when designing their business models for enterprises. While value-based pricing is harder to implement operationally, there is longer-term durability in actively shaping how customers should think about the ROI you're delivering.


Fast in, fast out: the re-evaluation treadmill

AI sales cycles are compressing relative to traditional enterprise software — 52% close in under 6 months from first meeting to signed contract, 31% take 6-9 months, and 17% longer. For founders, this is good news on the front end.

The challenge is the back end: 77% of enterprises re-evaluate their AI vendors at least every 6 months or on a rolling basis. This creates a “fast in, fast out” dynamic that is fundamentally different from traditional enterprise SaaS, where multi-year contracts provided a moat of inertia. In enterprise AI, switching costs are lower and the re-evaluation cadence is relentless.

How often organizations re-evaluate their AI tools

N=150, single select



Landing the deal is the beginning; founders need to build for continuous value demonstration that survives the re-evaluation treadmill.

Implementation: the product alone isn’t enough

Even “buy” decisions require significant implementation effort. The product is not sufficient in and of itself. Implementation support — services, forward-deployed engineers, self-serve tooling, partner ecosystems — is a critical layer of the GTM motion. Founders who treat implementation as someone else’s problem will lose to those who own the end-to-end deployment experience.

How AI solutions are being implemented

N=150, select all that apply






The gap to value in production

83% of respondents report that fewer than half of their AI pilots became scaled deployments in the last 12 months. It's not that enterprises aren't experimenting — 85% are already in production or broadly deployed — and it's not that they're losing faith, since 74% are expanding budgets. The problem is structural: the path from “promising demo” to “signed annual contract” is where most AI solutions fail to convert. The product usually works; what pilots can't survive is the gauntlet of integration, compliance, procurement, and organizational buy-in.

Share of AI pilots converted to production in the last 12 months

N=150, single select




What makes a pilot scale: top-down buy-in

By the time the formal ROI analysis happens, the decision has often already been made informally. Bottom-up user love creates the internal demand; an executive champion provides political cover; clean integration removes objections. The ROI case ratifies what people already believe.

The pilot-to-production playbook is perhaps easier said than done: (1) make users love it, (2) cultivate an executive champion, (3) integrate cleanly, (4) build the ROI case in parallel.

Factors that tipped the decision to scale (pilot → full deployment)

N=150, rank up to 3




What kills a pilot: the integration-security gauntlet

The reasons pilots fail to convert point to where the execution gap actually lives. "Didn't work as promised" ranks only 6th at 27%. The product usually works — it's the lack of ability to fit to the surrounding enterprise environment that kills the pilot.

Why AI pilots did not convert to full contracts

N=150, rank up to 3



The ROI paradox: measuring everything, transforming little

On the surface, enterprises look sophisticated about AI measurement: the majority of enterprises already use quantitative metrics.

63% measure AI value with quantitative metrics — time saved, headcount cost, sales uplift	19% have expectations of a return with AI but don't yet have a formal framework	13% measure qualitatively based on user satisfaction and perceived impact on productivity	5% do not measure the value of AI today
---	---	---	---

But the self-assessed impact of AI suggest the lag between experimentation and value delivery, where the latter has not yet fully translated across organizations and functions. 61% describe their AI initiatives’ impact as “significant but contained”, while only 17% say it's already truly “transformational.” Enterprises are increasingly refining how they measure the value of AI, but the opportunity is in helping define and compound these gains.

What’s coming next

Budget trajectory: nearly unanimous expansion

74% expect AI budget growth in the next 12 months — 15% by more than 50%, 59% by 11–50%. Only a single respondent plans cuts. Even organizations still piloting are expanding (59%). Budget commitment is running ahead of proven impact.

The workforce shift: job descriptions are being rewritten

83% say AI is reshaping their workforce, and the dominant response is to transform roles rather than eliminate them.



The builder’s view: engineering hiring is already changing

Roughly half of engineering organizations now require AI tool use during interviews, and 78%+ report at least a moderate shift in hiring criteria — toward system design and architecture, away from raw coding output. The bar for new engineers is moving up the stack.

Closing the execution and value gap: a playbook for founders

The founders who win will build products, GTM motions, and implementation playbooks that specifically align their solution with how enterprises purchase and realize value.

- 1 Features only matter when you align with their existing systems.**
Integration threads through the entire buyer lifecycle — #2 key purchase criterion, #3 decision factor for scaling a pilot, complexity here the #1 pilot killer. Build deep integrations, not just features.
- 2 Win the users AND the sponsors.**
41% of enterprises discover AI tools through internal tech teams. Strong end-user adoption and feedback is the #1 factor that tips decisions to scale. Make the practitioner your first customer and give them the tools to champion you internally.
- 3 Price for value, not only consumption.**
A gap in demand for pricing more aligned with outcomes is a structural GTM advantage over competitors still charging only per token.
- 4 Make security and compliance a centerpiece.**
Meeting high standards for security, privacy, and compliance is the #1 purchase criterion, with the lack of it the #2 pilot killer. Not investing here is a non-starter.
- 5 Build for the re-evaluation treadmill.**
77% of enterprises re-evaluate their stack at least every 6 months. Build continuous value demonstration into the product: usage analytics, proactive health checks, expansion triggers.
- 6 The budget exists — find the right door.**
Budget ownership is fragmented across IT, functional leads, and split models. Map the specific org structure at each account across both your product champion and your economic buyer.

Looking ahead: the next phase of enterprise AI

From discovery to deployment

The first act was discovery — that phase is largely over. 85% of respondents are already in production or broadly deployed. Enterprises have decided to buy AI; the open question is whether they can operationalize it. In the deployment phase, the product is the workflow: the integration, the security posture, the implementation support, and the pricing model that aligns with the buyer’s definition of value. And even once deployed, renewals are not guaranteed; a vast majority of enterprises are re-evaluating their AI tools at least every 6 months. Consistent value delivery is a must.

The application layer’s moment

Every major platform shift has followed the same pattern: value concentrates in infrastructure first, then migrates to the application layer as the platform matures. Cloud followed this arc; mobile followed it; AI is following it now. Inference costs have fallen nearly 100x in two years, and protocols like MCP and A2A are standardizing interfaces between agents, tools, and data — the application layer has more room to scale as these layers stabilize.

Three structural shifts to watch

From copilots to agents to multi-agent systems Each step up the autonomy ladder requires exponential investment in governance, observability, and coordination infrastructure.	From seat-based to outcome-based value capture As AI moves from augmenting work to completing it, the unit of value shifts from “access” to “result.”	From technology investment to operational transformation The bottleneck is shifting from building AI to operationalizing it — change management, data quality, and process redesign.
--	---	--

The exit landscape is warming

For the first time in several years, the exit environment for AI companies is showing real momentum. SpaceX’s June 2026 IPO at roughly \$1.75 trillion validated the path for technology companies at unprecedented scale. Anthropic and OpenAI have both filed

confidentially for IPOs that could create the first AI-native trillion-dollar public companies — and set the pricing benchmark for the entire private AI ecosystem.

Beyond the mega-IPOs, AI-native M&A is trending toward earlier-stage, faster exits: roughly half of GenAI acquisitions in 2026 involve companies less than five years old, up from near zero in 2022. Hyperscalers are acquiring teams and capabilities as much as products.

FOR FOUNDERS

Building a product that survives the enterprise re-evaluation treadmill also builds one that commands a premium exit multiple. The window for building category-defining AI companies is open now, and it will narrow.



THE FOUNDER'S MANDATE

Enterprise buyers are spending, and the market opportunity is enormous. But the bar for what it takes to move from a successful demo to a scaled production contract has never been higher.

We are in the early chapters of what Madrona has framed as the Reasoning Revolution: the emergence of intelligent systems that go beyond curated data and models to enable true collaboration between humans and AI. The founders who prove out and harness durable value will define how the next generation of enterprise software works.

Madrona invests in visionary founders and teams from pre-seed to Series C. Come and talk to us when you're ready to go big.

Methodology & appendix

Enterprise decision-maker survey

N=150. Seniority: Director through C-Suite (58% Director/Sr. Director, 33% VP/SVP, 9% C-Suite). Organization size: \$500M+ annual revenue (61% at \$10B+). Functions: Tech (n=50), Back Office (n=50), GTM (n=50). Geography: North America and Europe. Fielded May 2026.

Madrona Builders Summit GenAI survey

n=48 engineering leaders from the Madrona Builders Community. Year-over-year comparison (2025: n=44, 2026: n=48). Org size split: ≤50 engineers (n=26), >50 engineers (n=23). Fielded at Builders Summit, May 2026.

Madrona IA40 list longitudinal analysis

Funding, valuation, and exit histories across five annual historical lists, 2021–2025. Figures reflect publicly reported data as of mid-2026 with a June 1st cutoff.

Madrona invests in visionary founders from pre-seed to Series C. Founded in 1995 with a first check into Amazon, the firm invests across AI, developer tools, security, and consumer apps, from its offices in Seattle and the Bay Area. Madrona was an early investor in companies such as Snowflake, UiPath, Smartsheet, Rover, Runway, Statsig, and Temporal.