

JULY 2026

2026 State of AI

The Builder's Economy

Private and Strictly Confidential
Copyright © 2026 ICONIQ Capital, LLC. All Rights Reserved

For Professional Clients Only. ICONIQ Partners (UK) LLP (973080) is an appointed representative of Kroll Securities Ltd (466458) which is authorized and regulated by the Financial Conduct Authority

Introduction

Building and operationalizing AI products has moved from competitive differentiation to table stakes. In January 2026, we published “The State of AI: Bi-Annual Snapshot” to elevate the voices of architects, engineers, and product leaders driving the disciplined execution of delivering and scaling AI products.

What began as a race to ship AI features has matured into a more consequential challenge: turning AI capabilities into durable, economically sound businesses. Many of builders who are pulling ahead have made deliberate tradeoffs across the full stack – from model selection and data infrastructure to org design, pricing architecture, and internal deployment – and are managing the cost, trust, and go-to-market implications that come with each.

This report is ICONIQ's third bi-annual report of the AI builder landscape, grounded in our proprietary Q4 2025 and Q2 2026 surveys of executives at software companies building AI products, alongside perspectives from our ICONIQ Community. Across both waves, one shift stands out: **the market has moved from proving AI works to proving AI pays**. AI products are approaching half of revenue at the companies we surveyed, margins are expanding, and the operators pulling ahead are the ones treating pricing, cost, and org design as product decisions rather than afterthoughts.

[Explore Our AI Perspectives](#)

Table of Contents

Section	Topics Covered	Page
AI Models & Infrastructure	• AI Product Landscape • Models Powering AI Products and Providers • Performance, Quality, and Security for AI Models • Data Readiness	6-14
AI Go-to-Market & Economics	• GTM Motions & Pricing Models • Revenue and Costs of AI Products • Gross Margin Optimization	15-21
Talent & Organization	• Organizational Structure of AI Companies • How AI is Impacting Org Structures • Spotlight: Forward Deployed Engineering	22-31
AI for Internal Productivity	• Spend on AI Tooling • Adoption and Impact of AI Tools and Agents • Case Studies: How Functional Teams Are Implementing AI	32-43

Data Sources & Methodology¹


This study summarizes data from a **Q2 2026 survey of ~305 executives² at software companies building AI products**, including CEOs, Heads of Engineering, Heads of AI, and Heads of Product.

Throughout this report, we compare insights to our prior State of AI report, published in Q4 2025³, “[2026 State of AI: Bi-Annual Snapshot](#)”, where applicable. Where necessary, longitudinal data has been normalized to account for differences in firmographics to ensure trends are representative of the data.

We also weave in insights and what we believe to be best practices from AI leaders from the ICONIQ community.

All industry perspectives shared in this report have been anonymized to protect company-level information.

Respondent Firmographics



1 – This data was collected anonymously by an external survey. Survey responses include some but not all ICONIQ Venture and Growth portfolio companies as well as companies not part of ICONIQ Venture and Growth’s portfolio.
2 – Certain questions in the survey were optional or routed based on persona. Accordingly, some n-Size numbers in this presentation are less than 305.
3 – The Q4 2025 report summarizes data from a December 2025 survey of ~300 executives at software companies, including CEOs, Heads of Engineering, Heads of AI, Heads of Product, Chief Revenue Officers, and Chief Financial Officers.

Four shifts we see defining the AI market today

01 Builders have converged on the application layer and are betting on agentic depth

Nearly **two-thirds of builders** are building **horizontal or vertical apps** (vertical alone at 43% of companies analyzed) and they're moving into **high-complexity domains**, with **financial services and healthcare use cases** rising in popularity. **Agentic capabilities and agent reliability/orchestration** are the **top 2 customer-facing priorities** for the coming year.

02 Multi-model appears to be the default with the provider mix reshuffled

To power AI products, most builders now run **~3.3 providers on average** (up from 3.1 in Q4'25). Among survey respondents, **Anthropic moved to the top spot** (51% Q4 '25 → 81% in Q2 '26) as the field has broadened. Selection criteria held steady (reliability, then cost), but **security/privacy and SOC2/SLAs both climbed**, signaling enterprise-readiness pressure as products mature.

03 AI is now more of a revenue and margin story rather than an experiment

AI products grew from **32% of revenue (2025 actuals)** to a projected **42% this year**, with gross margins moving from **45% to 53% and on track for ~59% by 2027**. Part of rising margins is driving by pricing model changes: **consumption- and outcome-based models are rising** (42% and 23%, with companies now blending 1.7 models on average vs. 1.5 in Q4'25), as **builders align price to the cost and value of usage**.

04 AI is reshaping the org chart, not just the product


AI-driven companies¹ **run flatter and more cross-functional than their peers**. **45% of companies are planning a different mix of roles** and **33% are actively planning a smaller team**. While companies continue to **see gains on internal productivity**, it is **difficult to predict internal AI's true cost**, with overruns emerging from tokens, data infrastructure, and enablement

Source: Perspectives from the ICONIQ GenAI Surveys (December 2025 & May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network


1 – Companies with 50%+ revenue driven by AI products

01

AI Models & Infrastructure



Builders have converged on the application layer, with nearly two-thirds of the market focusing on horizontal and vertical AI applications. Within this, data analytics and productivity tools are the most targeted use cases



★ Which use case does your primary AI product target?

% of Respondents, Horizontal & Vertical AI Builders Only, Single-Select, N=192


<u>Q2'26 Rank</u>	<u>Use Case</u>	<u>Q4'25 Rank</u>	<u>Q2'26 Rank</u>	<u>Use Case</u>	<u>Q4'25 Rank</u>
#1	Data Analytics And Business Intelligence	— 1	#7	IT & Security	▼ 4
#2	Productivity & Collaboration	— 2	#8	HR & Recruiting Tools	— 10
#3	Financial Services	▲ 8	#9	Legal & Compliance	— 9
#4	Marketing & Sales	— 5	#10	Product & Design	▼ 7
#5	Healthcare & Life Sciences	▲ 11	#11	Developer Tools	▼ 6
#6	Customer Support	▼ 3	#12	Other	— 12

Major Shifts

Companies are increasingly focusing on **use cases with highly complex and specialized workflows**, e.g., financial services and healthcare, compared to 6 months ago. The workflow-embedded nature of vertical products appears to give the product a natural moat that is harder for general-purpose AI models to displace overnight.

Source: Perspectives from the ICONIQ GenAI Surveys (December 2025 & May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Deep workflow integration is the priority. Many builders are designing their products for habitual use and prioritizing investing in agentic capabilities over the next 12 months




Where are you prioritizing customer-facing AI product investment over the next 12 months?
Ranked by share of respondents placing each category in their top 3, N=305

1	Agentic AI capabilities	66%
2	Agent reliability & orchestration infrastructure	35%
3	Deeper vertical or domain-specific functionality	35%
4	Application layer differentiation	34%
5	Enterprise readiness	27%
6	Context engineering	24%
7	Multimodal capabilities	20%
8	Integrations and ecosystem	18%
9	Cost and inference optimization	14%
10	Evaluation and reliability infrastructure	14%
11	Model layer differentiation	13%


Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Licensed third-party APIs are the most popular model type to power customer-facing AI products, but adoption of open-source models is increasing. We expect this shift to continue as cost and control advantages compound



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network; 1 – N = 288; 2 – N=292

Similarly, builders are choosing a wider variety of model providers (~3.3 models on average). Anthropic has grown to the top provider among survey respondents over the last 6 months



Source: Perspectives from the ICONIQ GenAI Surveys (December 2025 & May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network; 1 – N = 194; 2 – N = 301

Top model selection criteria have largely remained consistent over the last 6 months, with model reliability & accuracy remaining in the top ranked priority

Top Considerations When Choosing a Foundational Model for Customer-Facing Use Cases

Ranked by share of respondents placing each in their top 3, N=305

Q2'26 Rank	Consideration	Q4'25 Rank
#1	Model Reliability & Accuracy	— 1
#2	Cost	— 2
#3	Security / Privacy	▲ 4
#4	Ability to fine-tune / Customize	▼ 3
#5	Latency	— 5
#6	Speed of Model Advancement	▲ 7
#7	Model Transparency / Explainability	▲ 8
#8	SOC2 / Enterprise SLAs	▲ 10
#9	Competitiveness with Peers	▼ 6
#10	Open Source	▼ 9

Perspectives from the ICONIQ Network

Security/privacy, explainability, and SOC2/SLAs have climbed in importance over the last 6 months, signaling enterprise-readiness pressure as products mature.

Security is becoming a switching cost.
An enterprise software CISO noted that the first provider to lead on enterprise controls (account blocking, execution governance, governance hooks, etc.) becomes the default and drives consolidation.

Observability is the new blocker.
A CISO at a major data platform flagged a lack of telemetry and visibility into what agentic tools are actually doing as an acute and growing concern.

Regulatory explainability is a hard gate.
A CISO at a global insurer can't yet describe to regulators what deployed agents are doing or how they're monitored which instead of capability, is the core blocker.

The gaps are live, not hypothetical.
A global bank's security team frames five production domains (identity, shadow AI, data provenance, supply chain, insider risk) as the lens through which model providers are now evaluated.


Source: Perspectives from the ICONIQ GenAI Surveys (December 2025 & May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Builders are widely adopting a range of approaches to help ensure model output quality and safety, including user feedback-driven evals and audit logging

QUALITY & RELIABILITY

Quality assurance is mostly reactive today. Builders catch issues through feedback, monitoring, and human review rather than pre-deployment testing. Proactive adversarial testing remains a minority practice at 21%.


How builders ensure output quality
% of Respondents, Select All That Apply, N=302



SAFETY & SECURITY

Data protection guardrails, including audit logging, scope restrictions, and PII controls are widely adopted. Adoption is lower for safeguards aimed at AI-specific risks (e.g., prompt injection, data exfiltration, autonomous model actions).

Guardrails in place
% of Respondents, Select All That Apply, N=302




Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Context engineering is becoming increasingly important to AI product development as builders combine RAG, memory, agent-to-agent context, among other techniques, to manage and provide context to models at runtime

How does your AI product manage and provide context to models at runtime?


% of Respondents, Select All That Apply, N=294



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Increased investment in AI infrastructure appears to be paying off, with more companies stating their data foundation is fully ready and optimized to support accurate AI workflows


How would you rate your data foundation’s readiness to support accurate AI workflows?¹
% of Respondents, Single-Select




Source: Perspectives from the ICONIQ GenAI Surveys (December 2025 & May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network. N-Sizes: Q4 2025 – N=198; Q2 2026 – N=304
1 – Fully ready = Governed, high-quality, and labeled data optimized for AI workflows; Mostly Ready = Solid data pipelines, minor gaps; Somewhat ready = Basic data systems in place, limited for AI; Not ready = Significant data gaps or quality issues

02

AI Go-to-Market Strategy & Economics




GTM strategies for AI builders are diversified across different motions, but hybrid motions (combination of PLG and SLG) has gained the most traction over the last 6 months



Source: Perspectives from the ICONIQ GenAI Surveys (December 2025 & May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network; 1– N=288; 2– N=289


Subscription/platform components are still the most common pricing model; however, consumption and outcome-based pricing has grown in prevalence over the last 6 months to better align pricing to product value



Source: Perspectives from the ICONIQ GenAI Surveys (December 2025 & May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network; ; 1– N=297; 2– N=305

As AI becomes a more significant driver of revenue and customer base for companies building AI products, the economics of building and selling AI products is becoming increasingly important

% of Revenue from AI Products¹
Average, Excludes AI-Native Companies, N=265



Perspectives from the ICONIQ Network

1. AI adoption with customer bases frequently requires a Trojan horse

- Top-down "end-to-end automation" pitches reliably trigger pushback from non-AI-native enterprise users
- What works: match the AI workflow to what users already do, just faster, and let them discover the compounding effects themselves. Adoption follows organically once two accelerated steps visibly connect
- Customer literacy programs pay off even with no production output. Getting users hands-on creates the internal advocates needed to scale later

"We tried the full end-to-end automated pitch and everyone had that knee-jerk reaction. Trojan-horsing it in, the same thing just faster, got them to organically form their own view of what could be automated." – CTO @ Series A, AI-Native Company

2. Feature distribution and customer education remain the bottleneck



- Shipping fast creates its own problem: the humans who talk to customers need continuous education on new features to translate them into adoption
- The more layers between product and customer, the more it compounds across product marketing, sales enablement, and account management

"I haven't found an AI-powered way to drive customer adoption automatically. The humans who talk to customers are still the main channel, and they need constant education." – Field CTO @ <\$50M AI-Enabled Company

Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network


1 – Includes companies where AI products drive >0% of revenue and excludes companies that have 95%+ of revenue driven by AI products, as these respondents are likely AI-native

Talent costs tend to scale down, and inference costs tend to increase as a product scales. For companies using consumption-based pricing, token/inference costs are generally a shared expense between providers and customers



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Overall, companies are seeing unit economics improve moderately on a per-query or per-task basis by focusing primarily on reducing inference costs, implementing routing strategies, and growing revenue for cost leverage




Primary Levers For Improving AI Margins
Ranked by share of respondents placing each in their top 3

- #1 Reducing inference costs
- #2 Implementing routing strategies
- #3 Growing revenue for cost leverage
- #4 Switching to lower-cost / OSS models
- #5 Raising prices for AI models
- #6 Negotiating better provider pricing
- #7 Reducing talent costs vs revenue
- #8 Other

Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network


Revenue expansion and utilizing optimization levers has led to an increase in gross margin; margins are expected to increase to ~60% on AI products in 2027



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

03

Talent & Organization




AI-driven¹ and high-growth companies tend to operate with more cross-functional structures, fewer layers, and wider spans of control

ORGANIZATIONAL STRUCTURE

AI-driven companies tend to have more cross-functional organizational structures

CURRENT ORGANIZATIONAL STRUCTURE BY AI REVENUE MIX

COMPANIES WITH CROSS-FUNCTIONAL STRUCTURE
Single-Select, N=297




We explore more themes around the new ways that AI companies are structuring their organizations and re-defining the workplace here:
[Org Charts to Outcomes: The New Operating Model for Talent](#)

MANAGEMENT LAYERS

AI-driven companies are run on fewer layers between CEOs and most junior ICs

MANAGEMENT LAYERS BETWEEN CEO/FOUNDER TO MOST JUNIOR IC BY REVENUE MIX

COMPANIES RUN ON 1-4 LAYERS
\$100M+ Companies², Single-Select, N=156



SPANS OF CONTROL


High-growth companies are widening first-line manager spans of control in 2026

FIRST-LINE MANAGERS WITH 7+ DIRECT REPORTS

R&D MANAGERS WITH 7+ DIRECT REPORTS
High-Growth Companies, Single-Select, N=56




GTM MANAGERS WITH 7+ DIRECT REPORTS
High-Growth Companies, Single-Select, N=43



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network
1 – Companies with 50%+ revenue driven by AI products; 2 – We expect earlier stage companies to be run on fewer layers and therefore, have focused this view on later stage companies

AI is reshaping workforce planning with 78% of companies planning for either a different role mix or smaller teams

Has AI-driven productivity changed how you think about headcount planning over the next 12-24 months at the organization level?
% of Respondents, Single-Select, N=302



Perspectives from the ICONIQ Network

The "right" engineering team is being redesigned in real time


Across the network, org design has shifted from a future planning question to a live operating challenge. **The debate is less about how many engineers and more about which roles still create value.**

The traditional software engineer who converts a design doc into a feature is the role most exposed; AI already does much of that work well. What holds value is infrastructure and security ownership and the product-engineering hybrid who can talk to customers and ship.

Companies are pushing ship authority as far down the chain as trust allows, typically to the level above new joiners, paired with a coherent strategy so leads can decide without escalating. One company caps teams at one lead per eight engineers, with the lead holding full ship authority inside their product area.

Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

As a result of headcount changes and AI-enabled efficiencies, ARR or Revenue per FTE is projected to grow steadily into 2027, with high-growth companies expected to outperform peers by ~11% in 2027




Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network
1 – Median Revenue = ~275M; 2 – Median Revenue = ~200M

R&D, sales, and product & design teams are generally expected to grow over the next 12-24 months, while customer support and G&A teams are projected to shrink

What is your organization’s plan for headcount in the next 12-24 months by function?

% of Respondents, Single-Select, N=303



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Many companies are actively hiring for forward-deployed engineers (FDEs) and planning to hire more engineers focused on AI safety / trust & reliability

Has your organization hired / started to evaluate candidates for these AI-specific roles?

% of Respondents, Single-Select, N=303

MAINSTREAMED *over half of orgs already staff it*

AI product manager



AI solutions engineers



SCALING UP *real footing, still expanding*

Forward-deployed engineers



AI prompt engineer / workflow specialists



CONTESTED *more orgs are sitting it out than buying in*

Chief AI Officer / Head of AI



AI safety / trust & reliability engineer



AI agent strategist



■ We have this role today ■ Actively hiring for this role now ■ Planning to hire in the next 6–12 months ■ Not planning to hire

Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

~50% of companies are scaling up their FDE model as part of their permanent GTM motion; FDEs are primarily used as a revenue driver and expected to materially contribute to expansion, retention, and strategic account outcomes

What is your organization's intended trajectory for the FDE model over the next 2-3 years?

% of Respondents, Single-Select, N=203¹

50%


Of respondents plan to scale up their forward deployed engineer model as a permanent part of their GTM motion

Full distribution



Which best describes the primary role forward deployed engineers play in your business today?


% of Respondents, Single-Select, N=200¹



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network
1- Excludes respondents that primarily target consumers

Forward deployed engineers are scaling up across customer segments, with enterprise customers expecting the highest integration of FDEs with customers in 2027

Across what percentage of customers does your company use forward deployed engineers (FDEs)?
Average by Customer Size, Companies with Customers Using FDEs, N=203¹




Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network
¹ – Excludes respondents that primarily target consumers

Most FDEs receive variable compensation with an average split of ~70/30 base vs. variable, usually tied to customer retention or renewal

PAY MIX

How are organizations incentivizing forward deployed engineers?

% of respondents, single-select, N=201¹




65% of FDEs receive some form of variable compensation

INCENTIVE DRIVERS

What incentives are tied to variable compensation?


% of respondents, multi-select, N=131¹



BASE VS. VARIABLE


What is the average split of base vs. variable compensation?

average across respondents, N=132¹



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network
1 – Excludes respondents that primarily target consumers


Monetization of FDE models remains fragmented, with companies split across bundled subscriptions, separate professional services fees, and hybrid models



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network
1 – Excludes respondents that primarily target consumers

04


AI For Internal Productivity



Internal AI systems spend is projected to rise from 11% to 16% of revenue in 2026, with further growth expected in 2027


As a percentage of total revenue, approximately what is your organization’s annual spend on AI systems for internal productivity?

Includes both direct spend (e.g., purchasing AI tools, salaries for AI-related roles targeted at internal productivity, etc.) and indirect spend (change management and upskilling, data governance, etc.); Average, N= 303



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Token spend for internal-facing AI tools tends to rise with company scale; however, tracking token spend is often less effective at measuring success of AI tools



Perspectives from the ICONIQ Network

Token counts fail fast as a metric: once consumption becomes the target, the practice undercuts the original goal of the metric. Charging AI usage back to cost centers works better, making spend visible and attributable without distorting behavior. Traditional engineering metrics, cycle time and deploys per week, are holding up as velocity proxies, and revenue per engineer cut by team is a useful signal for where to direct attention. **The most sophisticated approach in the room linked CSAT and product telemetry directly to specific features shipped, closing the loop between output and customer outcome.**


One early-stage company takes the aggressive version of this. **Unlimited tokens, but managers owe 3 to 5x productivity and own delivery accountability; top engineers run up to \$8K/day.** They pushed senior engineers onto AI code review, collapsed PM and designer into single-person product ownership, and removed the red tape between engineering and shipping.

Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Most companies struggle to predict internal AI’s true cost, with overruns emerging from tokens, data infrastructure, and enablement

Has your organization found it difficult to predict the “true cost” of implementing AI for internal productivity, or has it been easy enough to predict "true cost" in advance?

% of Respondents, Single-Select, N=305



Where Builders Saw Unexpected Costs


Token spend was the most common source of budget overruns. Once organizations moved from single-turn LLM calls to multi-step agentic pipelines, API costs scaled faster than initial estimates. One respondent noted a workflow projected at \$0.10 per run reached \$1.50 or more when agents needed to retry or self-correct across complex inputs. Without token guardrails and cost budgeting built into the architecture from day one, inference spend rapidly accumulated.

Data infrastructure was the second underestimated cost. Structuring data, defining access permissions, and building production-grade RAG pipelines proved significantly more resource-intensive than anticipated.

The third was **organizational enablement**. Moving from manual workflows to agentic processes requires governance frameworks, usage standards, and a sustained investment in training – costs that rarely appear in initial business cases but consistently ranked among the largest sources of budget variance.


Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

While AI tool access is broad across the organization, R&D teams tend to convert that access into deeper usage and higher maturity scores than G&A and GTM teams




Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network
1 – Employees that see 50%+ productivity gains in functional-specific workflows
2 – Survey respondents self-reported AI adoption maturity at their organization on a scale of 1 to 5, with 1 representing the lowest maturity and 5 representing the highest

High-growth companies are generally seeing a larger productivity impact across many AI use cases, with coding assistance in the lead with an average productivity gain of 48%





Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Scaled companies are more frequently using AI agents; however, the productivity impact for agents remains low, with average productivity gains of <30% across revenue bands



Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Agents often require human intervention to complete a task, particularly for multi-step tasks or where additional context / memory is required




Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Most organizations take a hybrid or centralized approach to internal AI deployment to help manage security concerns, AI tool sprawl, and cost visibility

How does your organization manage internal AI deployment (e.g., driving adoption, measuring usage / ROI, upskilling)?

% of Respondents, Single-Select, N= 304



How Ramp Achieved 99% AI Tool Adoption Across the Company¹

Ramp reached near-universal AI adoption internally, then confronted an unexpected problem: **most employees had no idea how to get more out of the tools they were using.** Complex setup requirements and fragmented configurations meant productivity gains were concentrated among a technical few, with no mechanism to scale what those individuals learned.

Their solution centered on removing configuration friction and letting the product itself do the enablement. over 350 reusable workflows were shared company-wide, Git-backed, versioned, and reviewed like code, helping the impact from AI enablement compound. Every skill shared helps raise the floor for everyone.

Ramp built this infrastructure in-house deliberately. **Internal productivity is a moat – the companies that make every employee effective with AI will move faster and compound advantages competitors cannot match.** Solving hard AI product problems internally also gave them direct conviction about what works before applying those lessons to customer-facing products.




Source: Perspectives from the ICONIQ GenAI Surveys (May 2026) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

1– Source: Public perspective from Seb Goddijn, Internal AI @ Ramp, ([Link](#), 4/9/2026)

AI adoption spotlight: How teams are implementing R&D AI use cases

01 — THE SHIFT


From copilot assistance to autonomous pipelines



Running **active agentic projects in parallel** to core work has emerged as the forcing function for adoption. Non-engineers now ship PRs too — **the PR is becoming a cross-functional artifact**, not an engineering-only output.

02 — WHERE GAINS LAND


Code velocity sees the highest productivity gains



Lead times compressing	-50-70%
PR throughput up	+25-40%
Concept-to-prototype	<2 days

03 — ORG IMPACT

AI absorbs entry-level capacity



Companies that are more mature in AI adoption are **reallocating budget toward senior and AI-specific roles** as juniors' work is automated. The most aggressive operators have **tied AI adoption to performance reviews and compensation**.

EMERGING
SIGNAL

- **Speed without verification infrastructure often fails.** One company reported ~10 autonomous agents produced ~180 PRs with a 0% CI pass rate, underscoring that autonomous output requires verification infrastructure before it touches production
- Companies that are more mature in AI adoption in regulated environments have **paired tooling rollouts with formal AI acceptable use policies and governance** is becoming a competitive input

Source: Illustrative examples based on practices observed at select ICONIQ portfolio companies; examples above are not intended to represent benchmarks or survey findings

AI adoption spotlight: How teams are implementing **GTM** AI use cases

01 — PIPELINE GENERATION

Prospecting that drives real revenue uplift

Enriched outbound + research

AI voice agent qualification

Qualified meetings

2–3×
meetings / rep

Some AI voice agents now handle 50–100%+ of inbound qualification at conversion rates comparable to human reps. A frontier move is replacing the SDR role entirely with in-house AI SDRs.

02 — CONVERSATION INTELLIGENCE

A closed coaching loop on sales calls

Live call

Transcribe + intel

Weekly AI coaching

Rep improves

Closed loop

Win rates up	+15–30%
New-rep ramp time down	–10–20%
Inbound SQOs up	3x per rep

03 — ORG IMPACT

CS and marketing headcount leverage

40–75%

AI deflection
no CSAT degradation

CSM portfolio / rep
+40–60%

Content output / writer
2–3×

Leading CS teams run high AI deflection while expanding portfolios per CSM. On the marketing side, output and lead volume are climbing while headcount stays flat or declines at several companies. Across GTM, growth is decoupling from headcount, with new hires gated on what AI cannot yet cover.

EMERGING
SIGNAL


- AI search optimization is an emerging priority: **companies are beginning to manage visibility in AI-generated answers as a distinct channel from traditional search**, and the operators moving early are building this into their marketing infrastructure now rather than retrofitting later
- The companies seeing the biggest efficiency gains are **also investing in RevOps infrastructure** (e.g. many are replacing non-AI native RevOps hires with AI-native ones) to make sure teams can access and query their own data

Source: Illustrative examples based on practices observed at select ICONIQ portfolio companies; examples above are not intended to represent benchmarks or survey findings

AI adoption spotlight: How teams are implementing G&A AI use cases

01 — FP&A


Close and forecasting cycles compressing



Some of the most advanced teams run **lean, with no traditional FP&A headcount**. The frontier move is **autonomous agents handling billing, AP review, and flux analysis end-to-end**, with humans reviewing exceptions.

02 — HR


Recruiting cycles getting shorter and cheaper



Many recruiters using AI tools are closing **roughly 2× the hires per quarter**, absorbing higher req loads without expanding the team. Onboarding and HR admin are also increasingly automated.

03 — ORG IMPACT

AI-first backfill policies



Several operators have **removed finance-ops and order-management roles** as agentic workflows took over, redirecting budget to strategic and AI-specific functions. The bar for any new G&A hire is rising: **roles are added only where AI cannot yet cover the work**.

EMERGING
SIGNAL

- **Natural-language analytics is moving data access out of specialist teams:** non-technical staff across finance, HR, and ops are now querying company data directly, collapsing the reporting-request bottleneck
- **A connective knowledge-retrieval layer is emerging across G&A**, with the most advanced operators treating internal data access as shared infrastructure rather than a function-by-function tool

Source: Illustrative examples based on practices observed at select ICONIQ portfolio companies; examples above are not intended to represent benchmarks or survey findings

ICONIQ

Where visionaries define the **future of their industries**

San Francisco | Palo Alto | New York | London



Private & Strictly Confidential

Disclosures

Unless otherwise indicated, the views expressed in this presentation are those of ICONIQ (“ICONIQ” or the “Firm”), are the result of proprietary research, may be subjective, and may not be relied upon in making an investment decision. Information used in this presentation was obtained from numerous sources. Certain of these companies are portfolio companies of ICONIQ. ICONIQ does not make any representations or warranties as to the accuracy of the information obtained from these sources.

This presentation is for educational purposes only and does not constitute investment advice or an offer to sell or a solicitation of an offer to buy any securities in connection with any investment fund or investment product that ICONIQ sponsors. Any such offer or solicitation will only be made pursuant to definitive offering documents and subscription agreements.

Any reproduction or distribution of this presentation in whole or in part, or the disclosure of any of its contents, without the prior consent of ICONIQ, is prohibited.

This presentation may contain forward-looking statements based on current plans, estimates and projections. The recipient of this presentation (“you”) is cautioned that a number of important factors could cause actual results or outcomes to differ materially from those expressed in, or implied by, the forward-looking statements. The numbers, figures and case studies contained in this presentation have been included for purposes of illustration only, and no assurance can be given that the actual results of any ICONIQ portfolio company will correspond with the information contained in this presentation. No information is included herein with respect to conflicts of interest, which may be significant. The portfolio companies and other parties mentioned herein may reflect a selective list of the prior investments made by ICONIQ.

Certain of the economic and market information contained herein may have been obtained from published sources and/or prepared by other parties. While such sources are believed to be reliable, none of ICONIQ or any of its affiliates and partners, employees and representatives assume any responsibility for the accuracy of such information.

All of the information herein is presented as of the date made available to you (except as otherwise specified), and is subject to change without notice, and may not be current or may have changed (possibly materially) between the date made available to you and the date actually received or reviewed by you. ICONIQ assumes no obligation to update or otherwise revise any information, projections, forecasts or estimates contained in this presentation, including any revisions to reflect changes in economic or market conditions or other circumstances arising after the date the items were made available to you or to reflect the occurrence of unanticipated events. Numbers or amounts herein may increase or decrease as a result of currency fluctuations.

































For avoidance of doubt, ICONIQ is not acting as an adviser or fiduciary in any respect in connection with providing this presentation and no relationship shall arise between you and ICONIQ as a result of this presentation being made available to you.

ICONIQ is a trading name of ICONIQ Partners (UK) LLP. ICONIQ Partners (UK) LLP (Registration Number: 973080) is an appointed representative of Kroll Securities Ltd. (Registration Number: 466588) which is authorised and regulated by the Financial Conduct Authority. ICONIQ Partners (UK) LLP is a limited liability partnership whose members are ICONIQ Capital (UK) Ltd, Seth Pierrepont and Lou Thorne, and it is registered in England and Wales and has its registered office at 27 Soho Square, London W1D 3QR. ICONIQ Partners (UK) LLP acts as an adviser to ICONIQ Capital LLC.

These materials are provided for general information and discussion purposes only and may not be relied upon. This material may be distributed to, or directed at, only the following persons: (i) persons who have professional experience in matters relating to investments falling within article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 (the “FP Order”), (ii) high-net-worth entities falling within Article 49(2) of the FP Order, and (iii) any other persons to whom it may otherwise lawfully be communicated (all such persons together being referred to as “FPO Relevant Persons”). Persons who are not FPO Relevant Persons must not act on or rely on this material or any of its contents. Any investment or investment activity to which this material relates is available only to FPO Relevant Persons and will be engaged in only with FPO Relevant Persons. Recipients must not distribute, publish, reproduce, or disclose this material, in whole or in part, to any other person.

Copyright © 2026 ICONIQ Capital, LLC. All rights reserved.

A global portfolio of category-defining businesses

These companies represent the full list of companies that ICONIQ Venture and Growth has invested in since inception through ICONIQ Strategic Partners funds as of the date these materials were published (except those subject to confidentiality obligations or companies for which the issuer has not provided permission for ICONIQ to disclose publicly). Further, the list of companies may not reflect the most recent ICONIQ Venture and Growth investments. Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.

The force behind every founder

