

ICONIQ

June 2025

The Builder's Playbook

2025 State of AI Report



Private and Strictly Confidential
Copyright © 2025 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

Follow our research



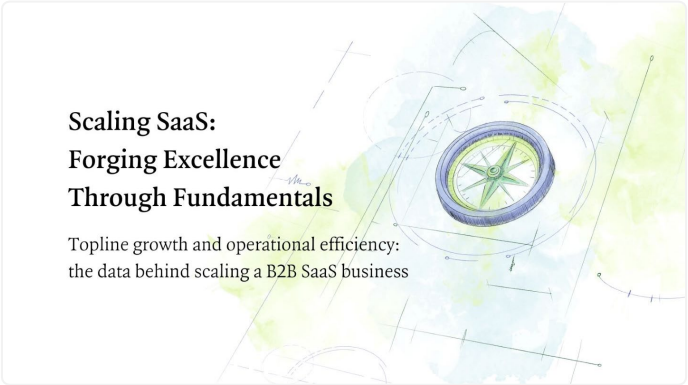
Go-To-Market Series

Guides to sales, customer success, marketing compensation – and more



Navigating Today's Public Markets

The metrics that matter and the market realities of 2025 and beyond



Growth & Efficiency

Explore our research on best-in-class SaaS growth and efficiency



The ICONIQ Enterprise Five

Key performance indicators of Enterprise SaaS companies



The SaaS Glossary

A guide to understanding and tracking key SaaS metrics



Engineering Series

Definitive guides to engineering excellence

Introduction

We believe that building and operationalizing AI products is the **new frontier of competitive advantage** – and that the voices of the architects, engineers, and product leaders driving this work deserve their own spotlight. While last year’s State of AI report centered on the buying journey and enterprise adoption dynamics, our 2025 report pivots squarely to the **“how-to”: what it takes to conceive, deliver, and scale AI-powered offerings end-to-end.**

This year’s report unpacks core dimensions of the builder’s playbook:

1. **Product Roadmap & Architecture:** The emerging best practices for balancing experimentation, speed to market, and performance at each stage of model evolution
2. **Go-to-Market Strategy:** How teams are aligning pricing models and go-to-market strategies to reflect AI’s unique value drivers
3. **People & Talent:** Building the right team to harness AI expertise, foster cross-functional collaboration, and sustain long-term innovation
4. **Cost Management & ROI:** Strategies and benchmarks for spend associated with building and launching AI products
5. **Internal Productivity & Operations:** How companies are embedding AI into everyday workflows and the biggest drivers of productivity unlock

Drawing on our proprietary survey results alongside in-depth interviews with AI leaders across the ICONIQ community, the 2025 State of AI report offers a blueprint for anyone tasked with turning generative intelligence from a promising concept into a dependable, revenue-driving asset.

[Explore Our AI Perspectives](#)

Table of Contents

Building Generative AI Products	Types of AI Products	9	Internal Productivity	Internal Productivity Budget	40
	Model Usage and Key Purchasing Considerations	11		Budget Sources	41
	Top Models Providers	13		AI Access and Usage	42
	Model Training Techniques	14		Key Purchasing Considerations	43
	AI Infrastructure	15		Deployment Challenges	44
	Model Deployment Challenges	16		Number of Use Cases	45
	AI Performance Monitoring	17		Top Use Cases	46
	Agentic Workflows	18		Attitude Towards Internal AI Adoption	48
<hr/>			<hr/>		
Go-to-Market Strategy & Compliance	AI Product Roadmap	20	Top AI Tools	LLM & AI Application Development	51
	Pricing	21		Model Training & Finetuning	52
	AI Explainability and Transparency	24		Monitoring & Observability	53
	AI Compliance and Governance	25		Inference Optimization	54
<hr/>				Model Hosting	55
Organization Structure	Dedicated AI/ML Leadership	27		Model Evaluation	56
	AI-Specific Roles and Hiring	28		Data Processing & Feature Engineering	57
	Pace of Hiring	29		Vector Databases	58
	% of Engineering team Focused on AI	30		Synthetic Data & Data Augmentation	59
<hr/>				Coding Assistance	60
AI Costs	AI Development Spend	32		DevOps & MLOps	61
	Budget Allocation	33		Product & Design	62
	Infrastructure Costs	34		Other Internal Productivity Use Cases	63
	Model Training Costs	36			
	Inference Costs	37			
	Data Storage & Processing Costs	38			

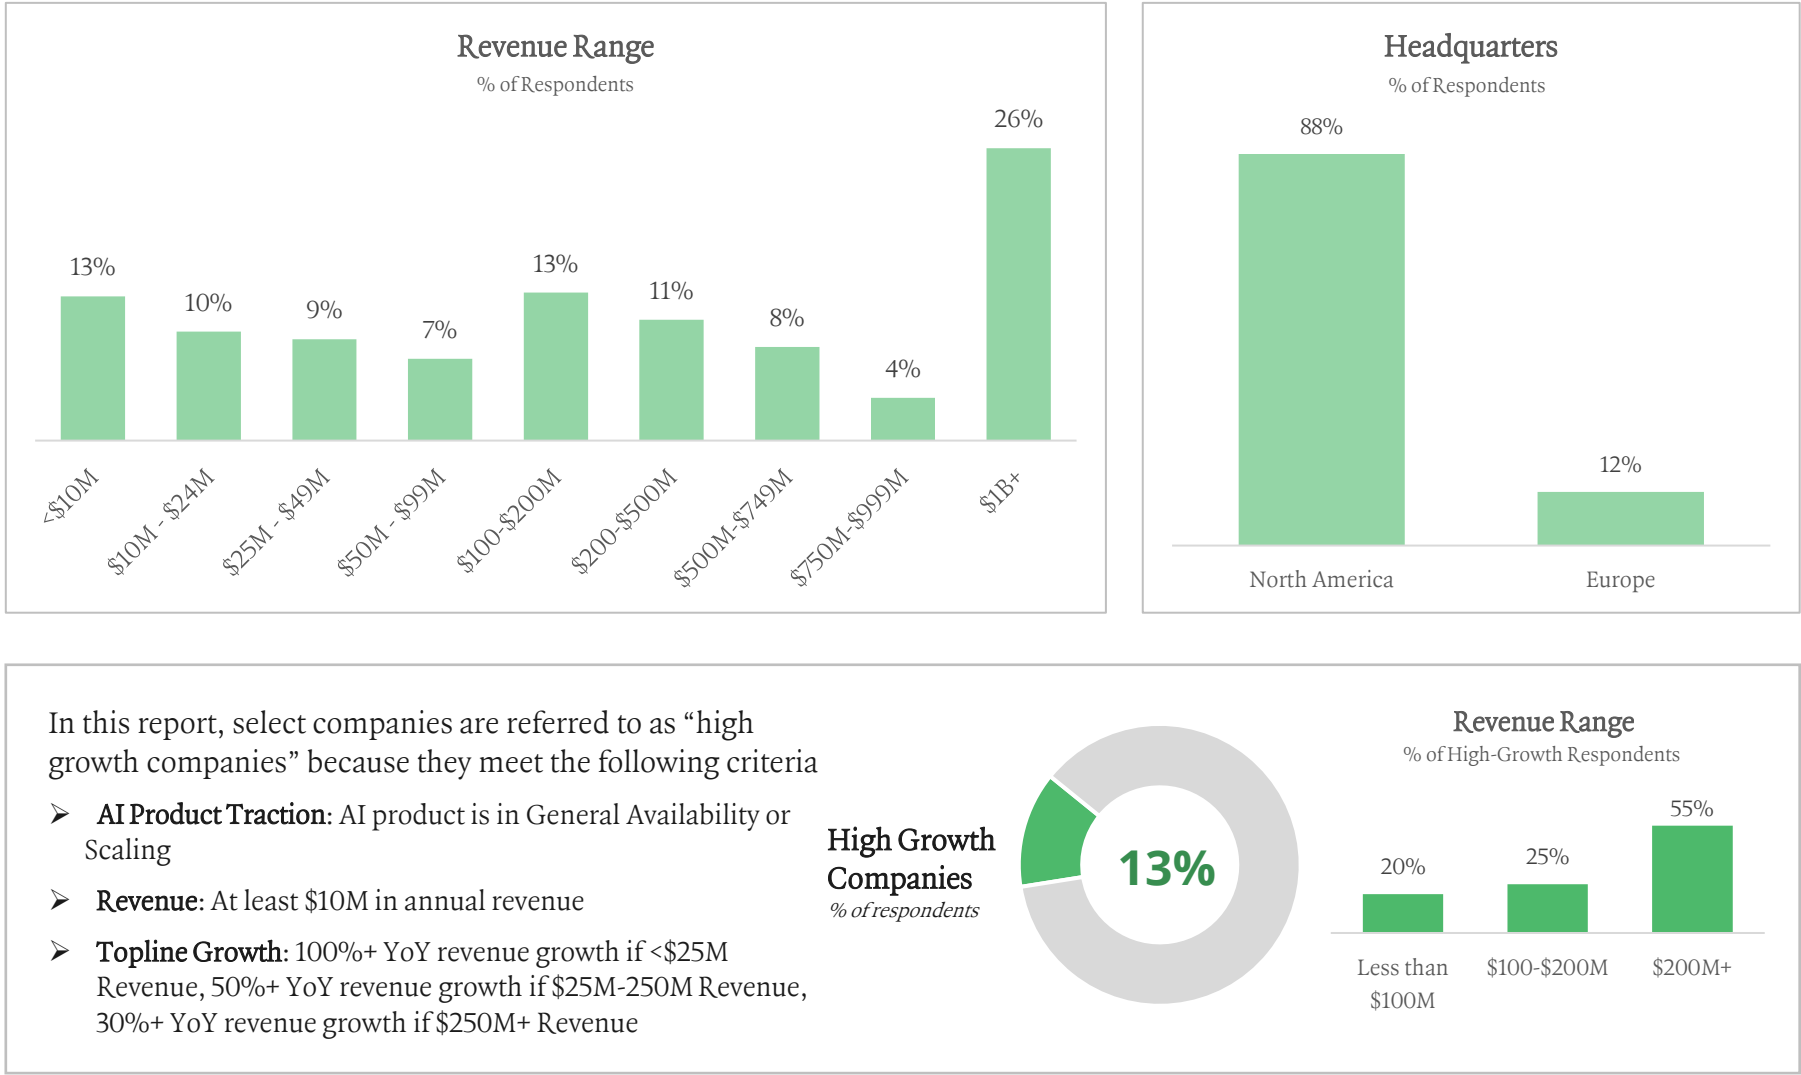
Data Sources & Methodology

This study summarizes data from an **April 2025 survey of 300 executives at software companies building AI products**, including CEOs, Heads of Engineering, Heads of AI, and Heads of Product.

Throughout this report, we also weave in perspectives, insights, and what we believe to be best practices from AI leaders from the ICONIQ community.

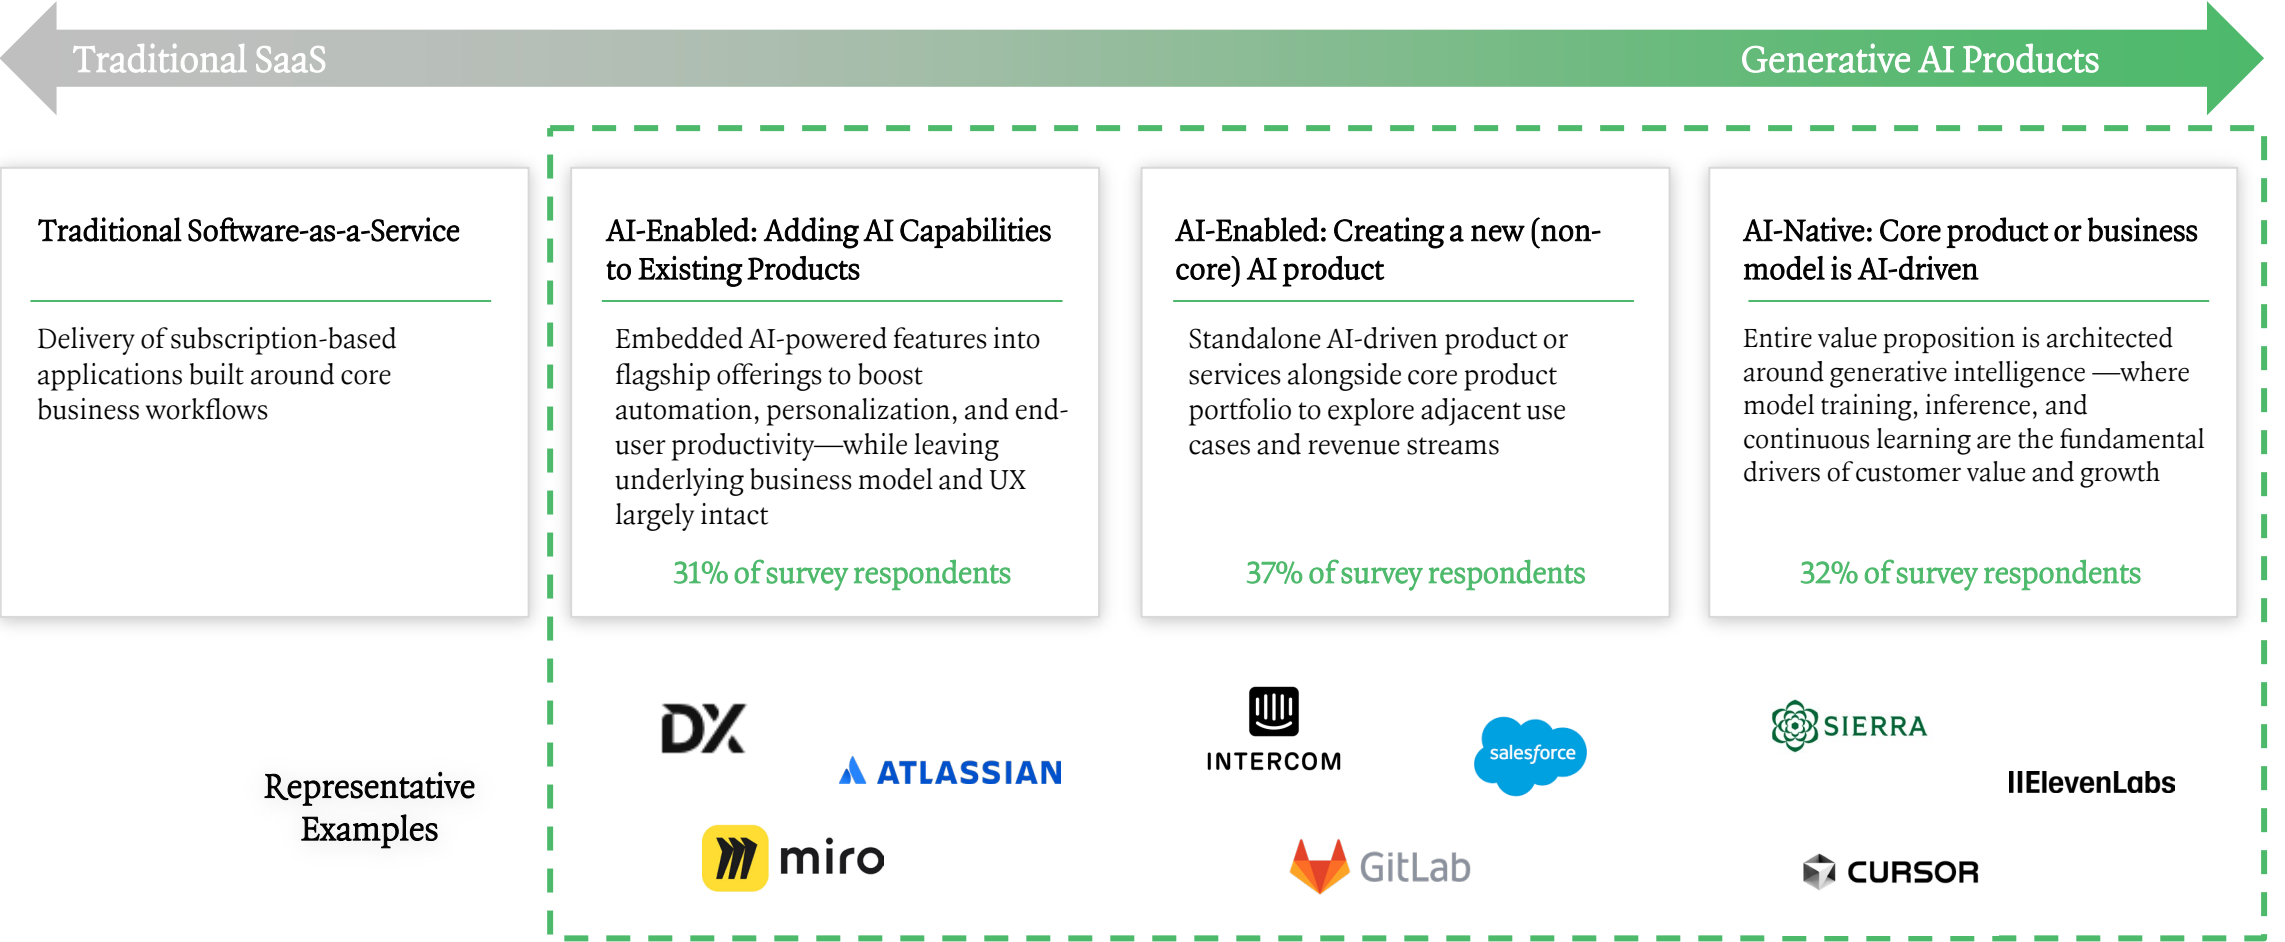
All industry perspectives shared in this report have been anonymized to protect company-level information.

Respondent Firmographics



Notes: (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. Accordingly, some N-Size numbers in this presentation are less than 300

Most SaaS companies have evolved to add new AI capabilities and products; the following pages will dive into how AI-enabled and AI-native companies are approaching product development



Focus of this report

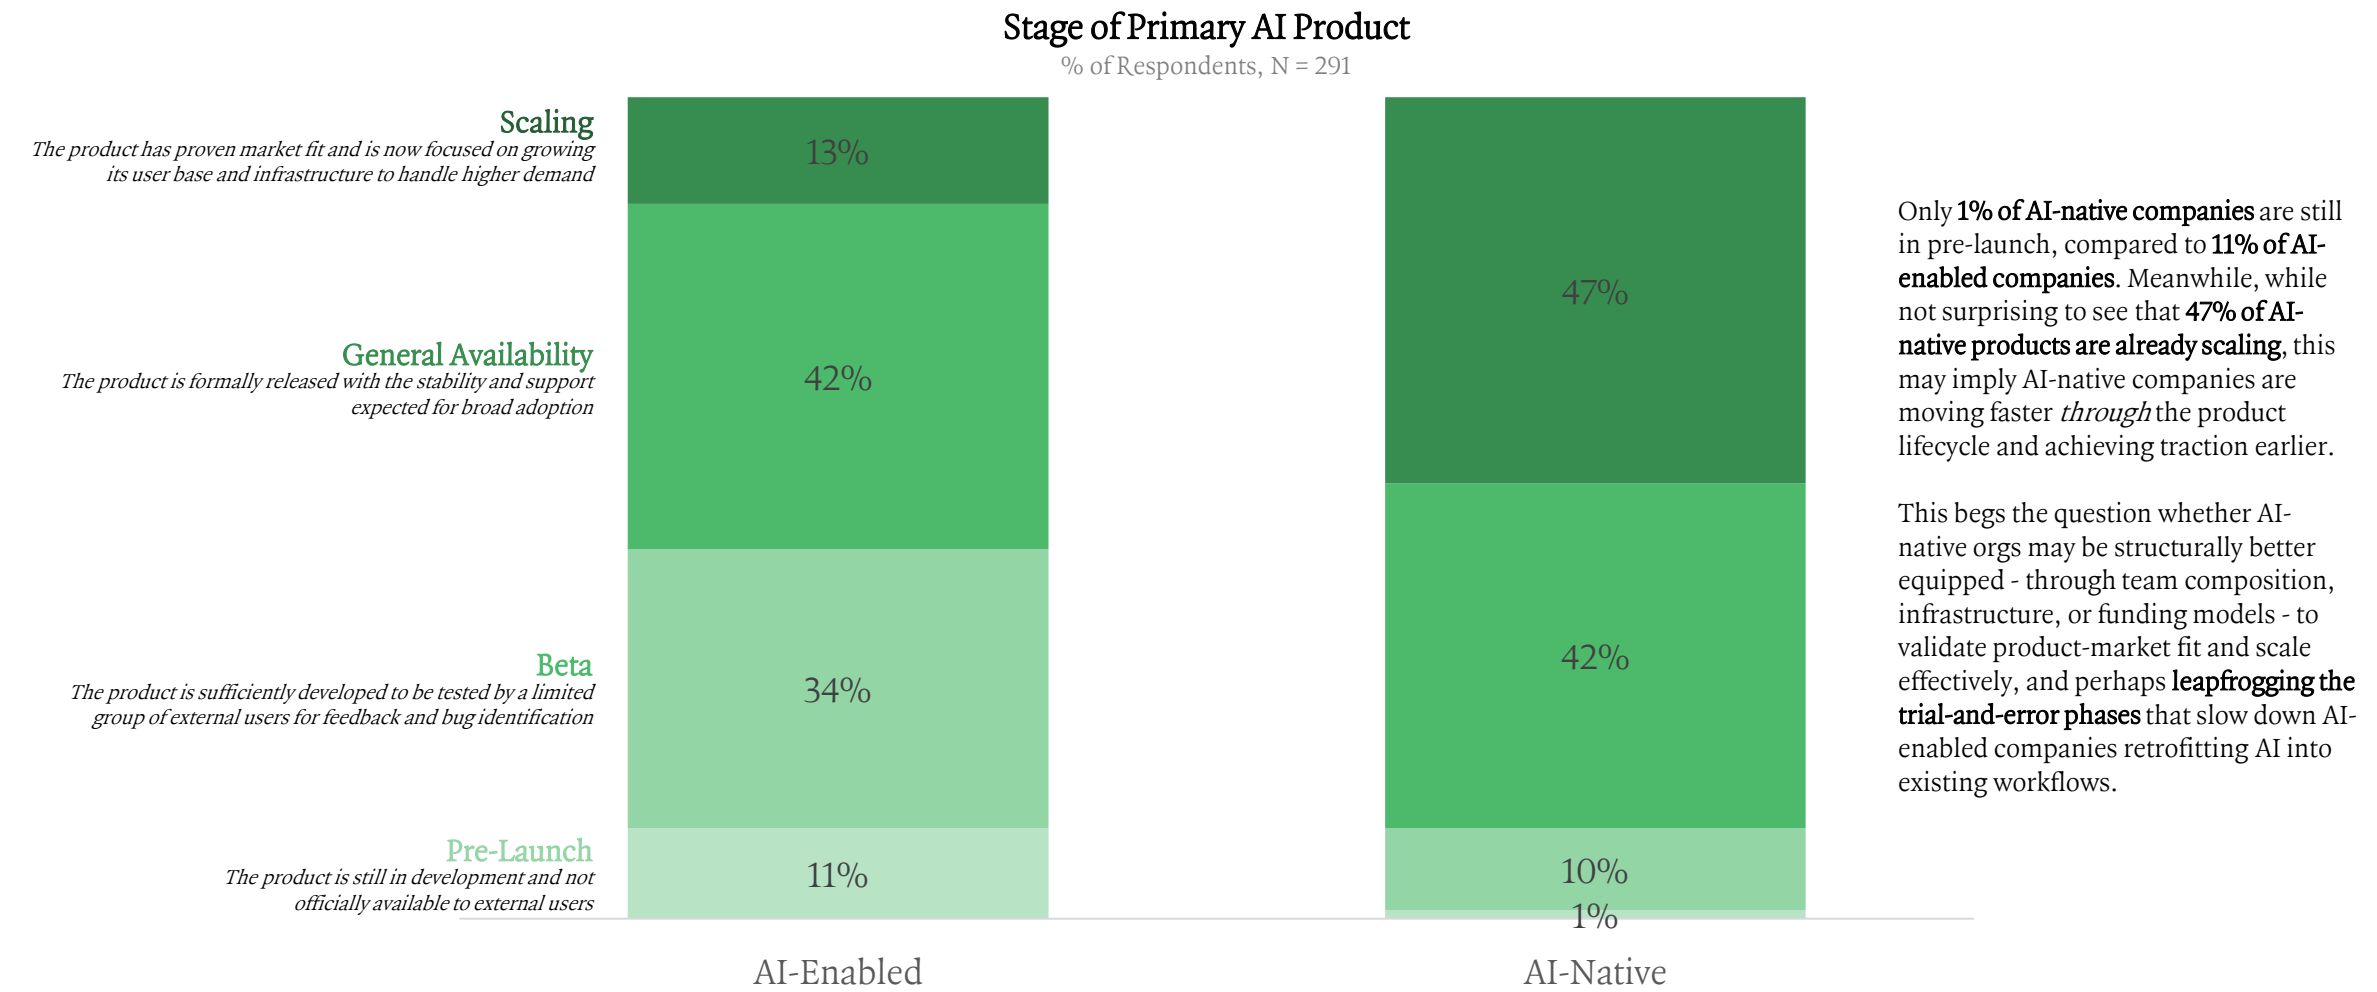
Notes: Representative Examples provided for illustrative purposes only. Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) and perspectives from the 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

Building GenAI Products



Stage of Primary AI Product

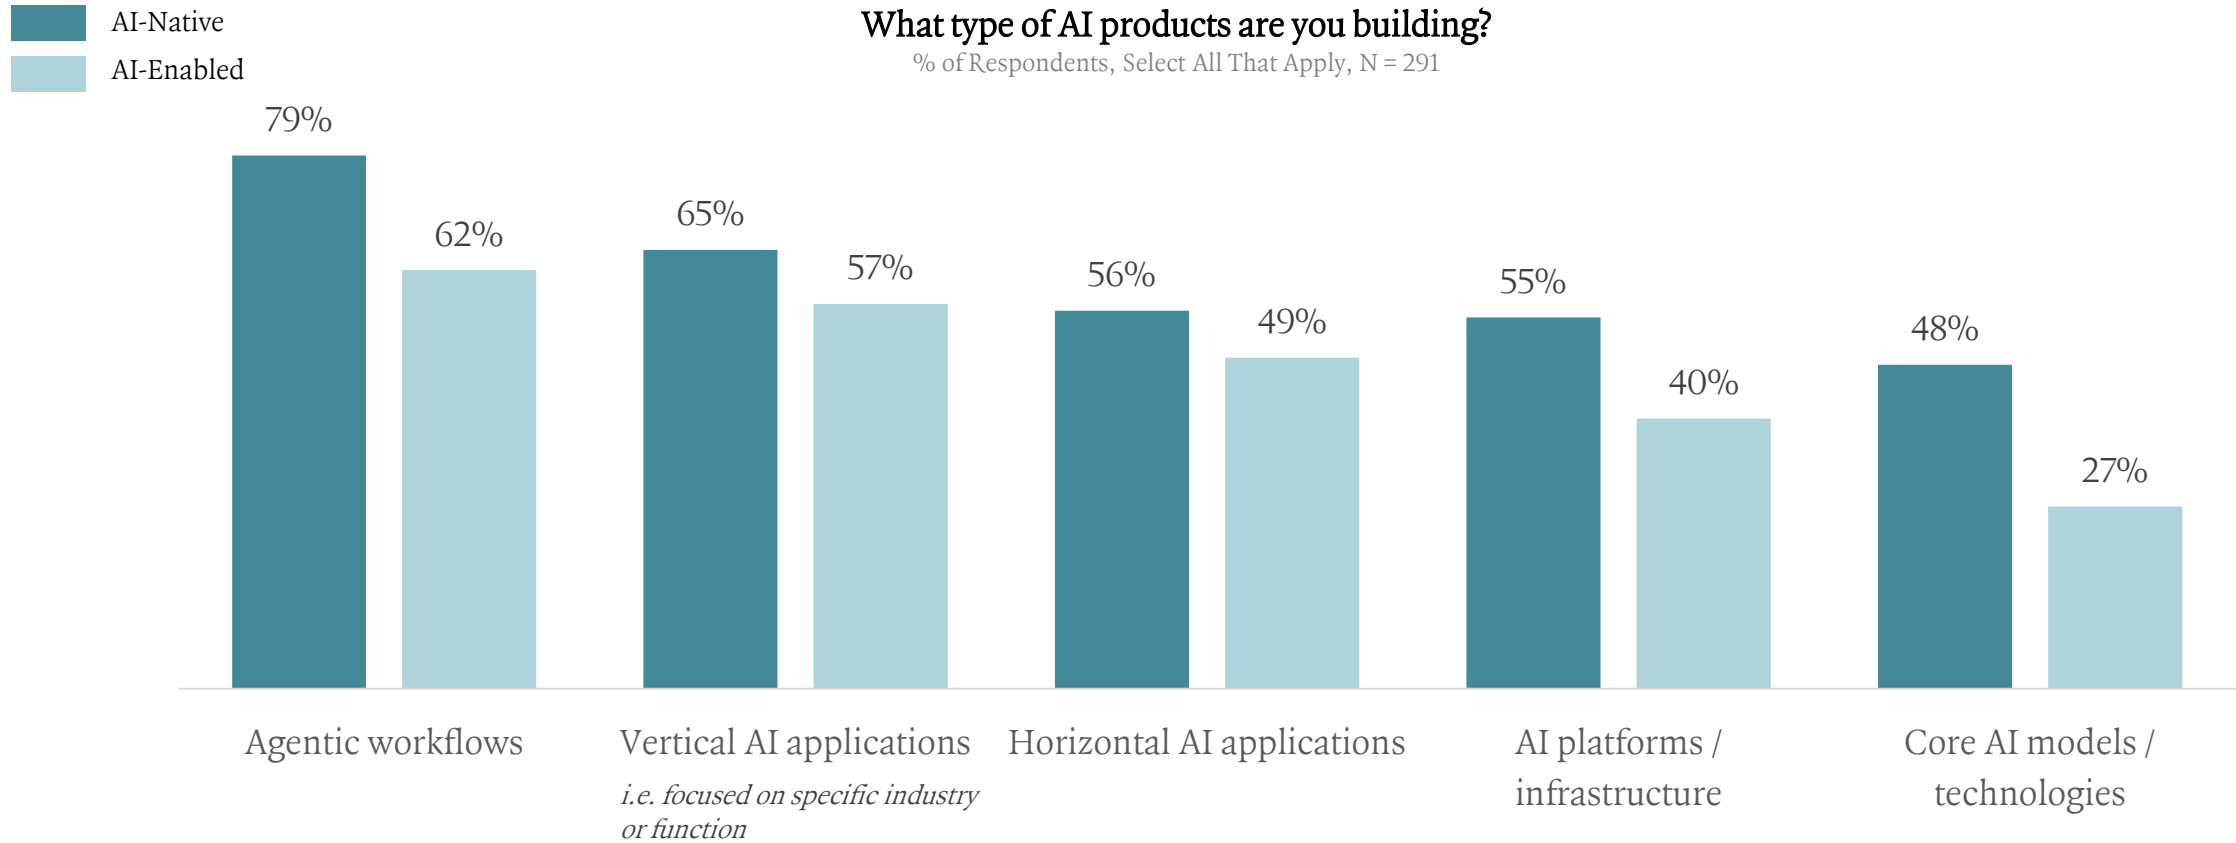
AI-native companies are further along in the development cycle compared to AI-enabled peers, with around 47% of products analyzed having reached critical scale and proven market fit



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Types of AI Products

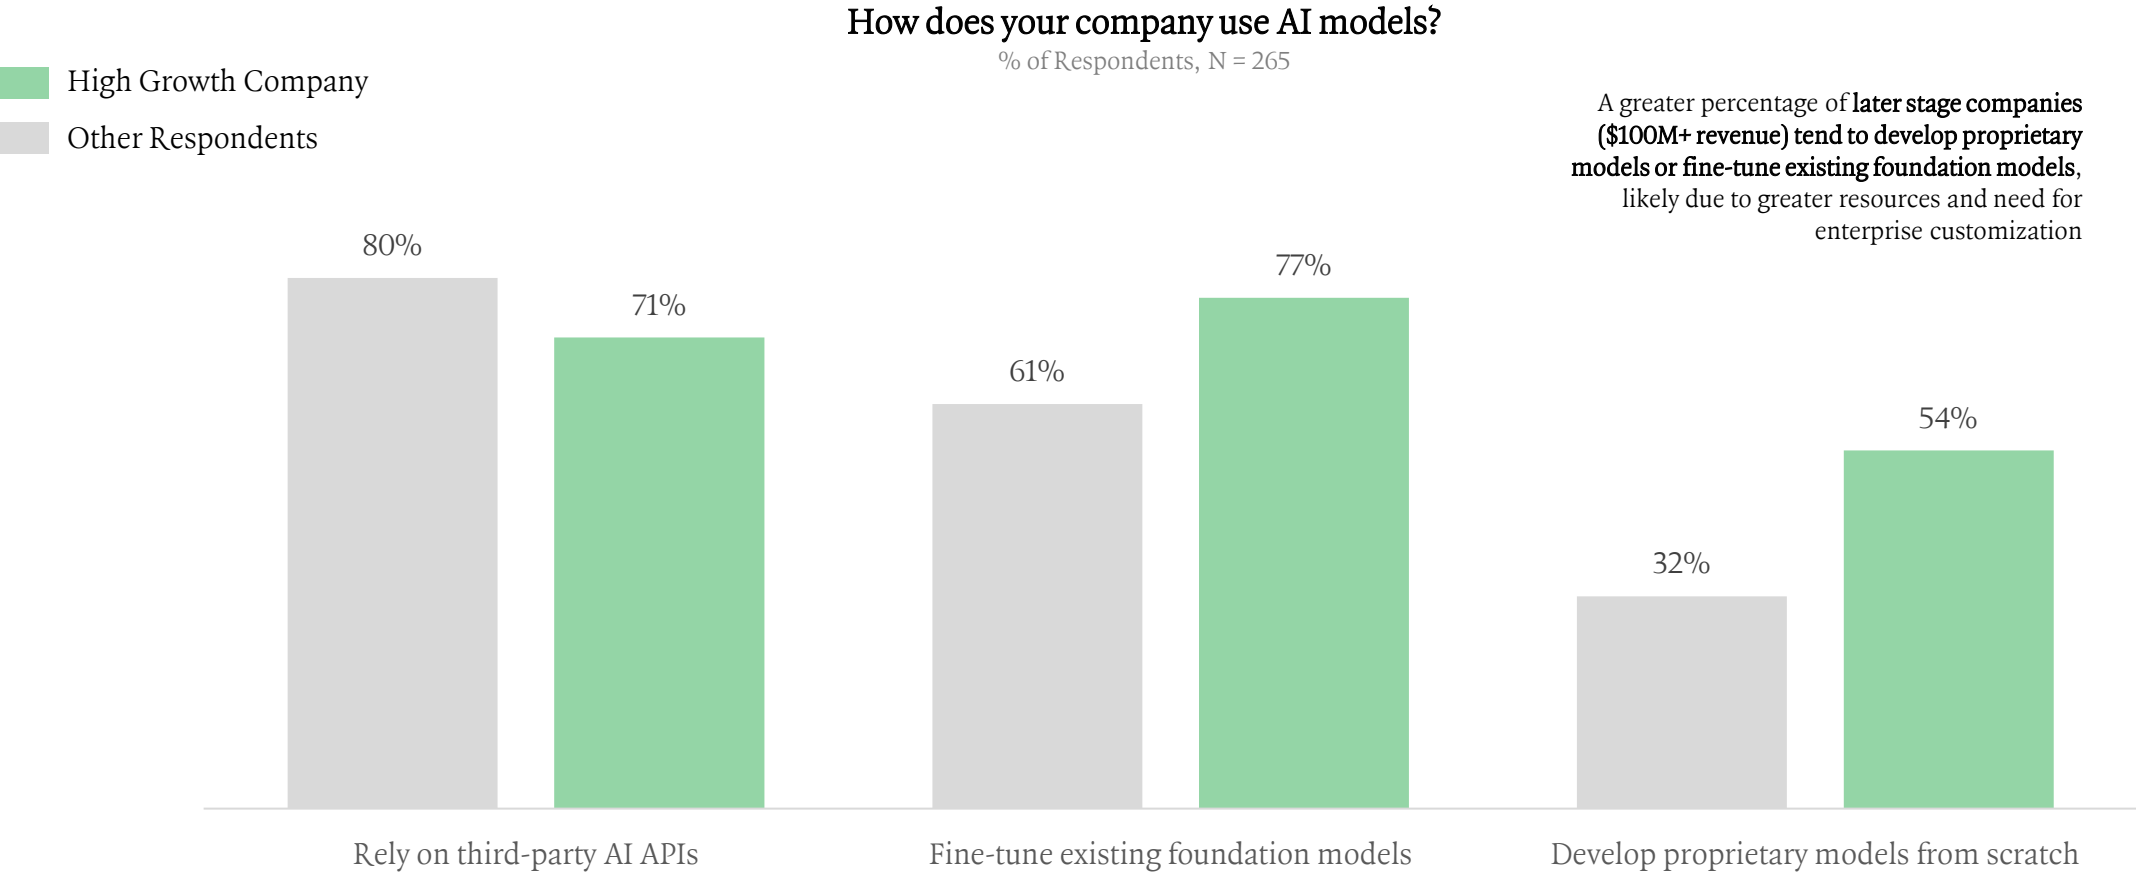
Agentic workflows and the application layer are the most common types of products being built across AI-native and AI-enabled companies; notably, around 80% of AI-native companies are currently building agentic workflows



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

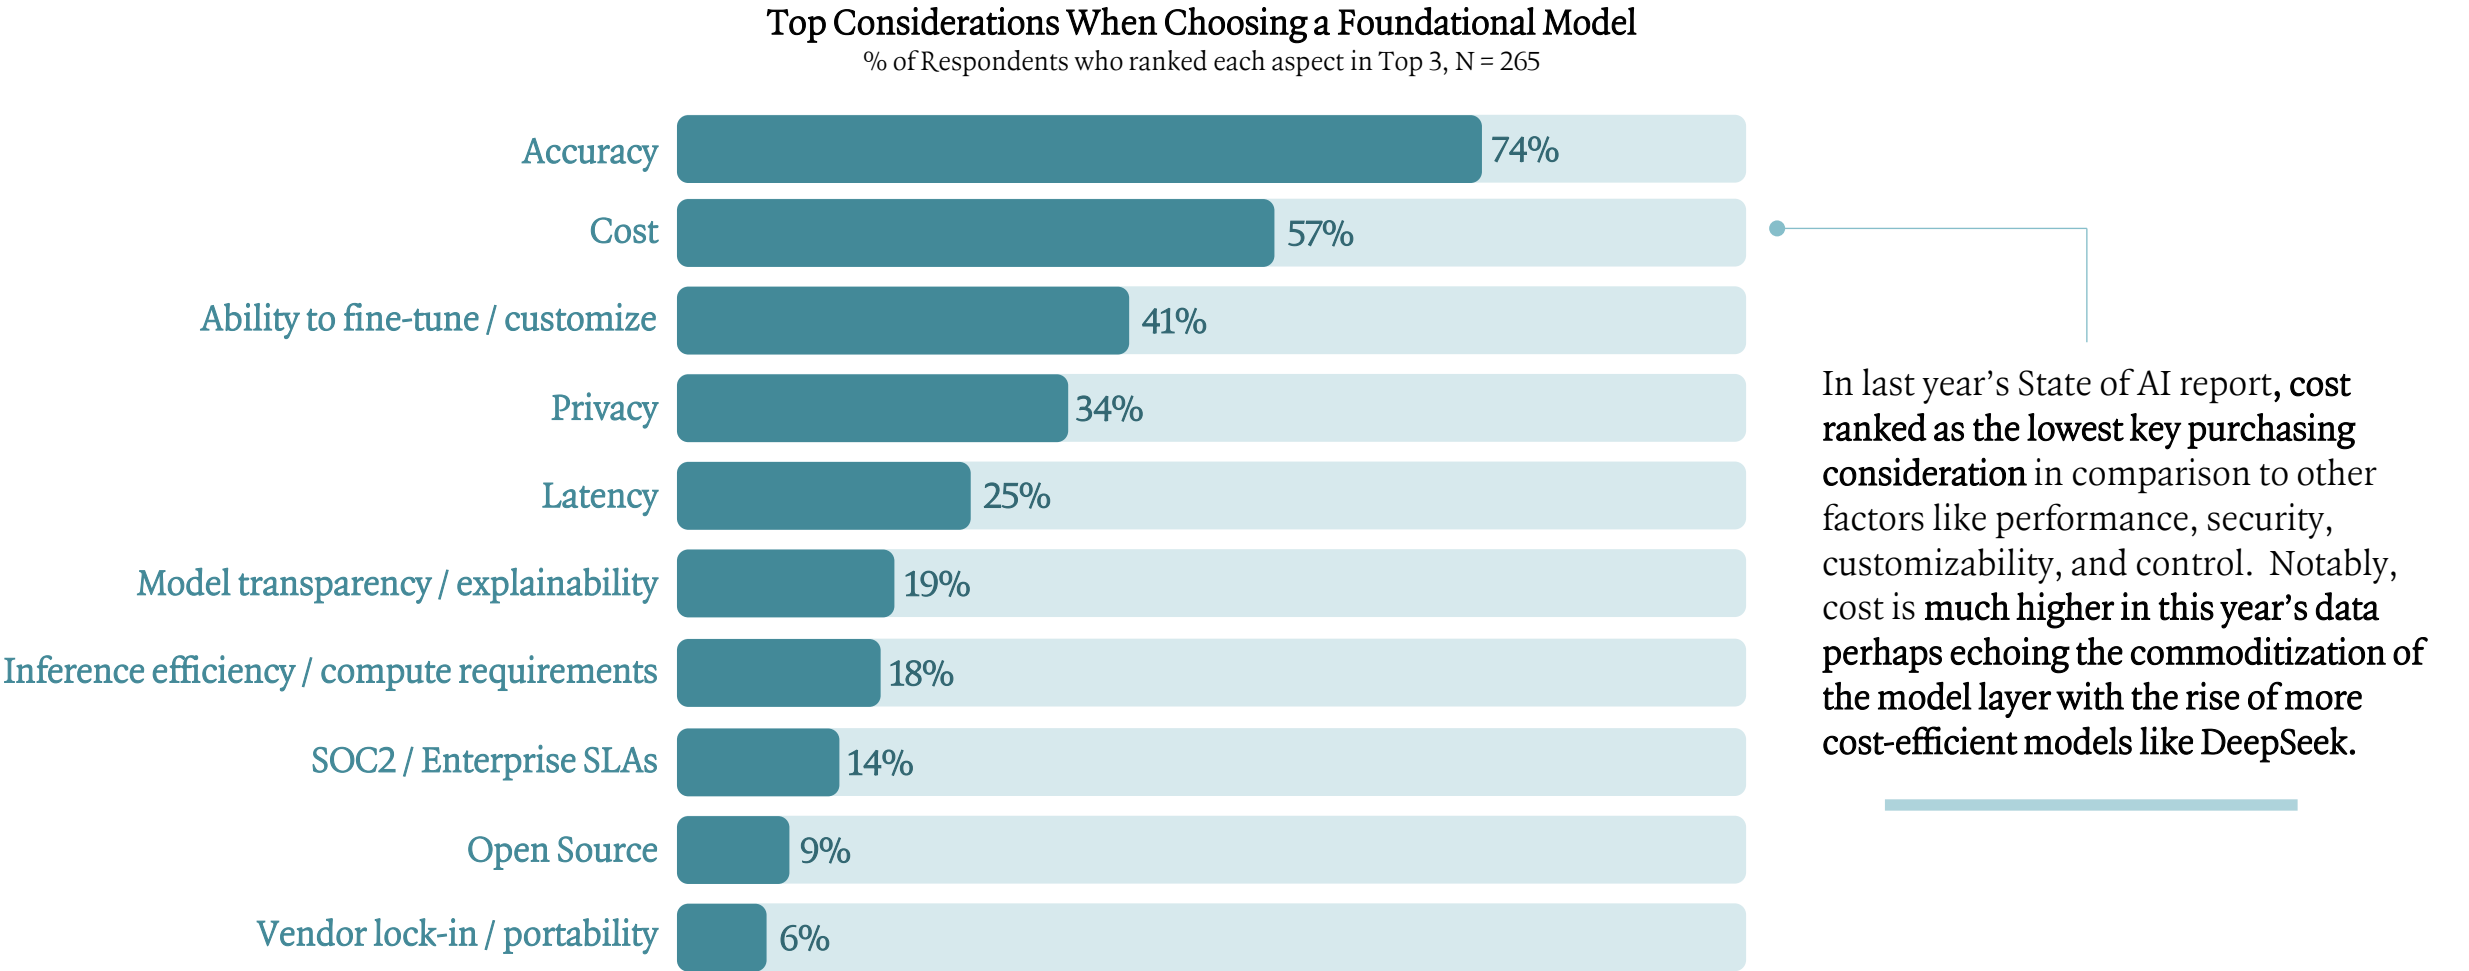
Model Usage

Most companies building AI applications are relying on third-party AI APIs; however, a larger proportion of high-growth companies are also finetuning existing foundation models and developing proprietary models from scratch



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

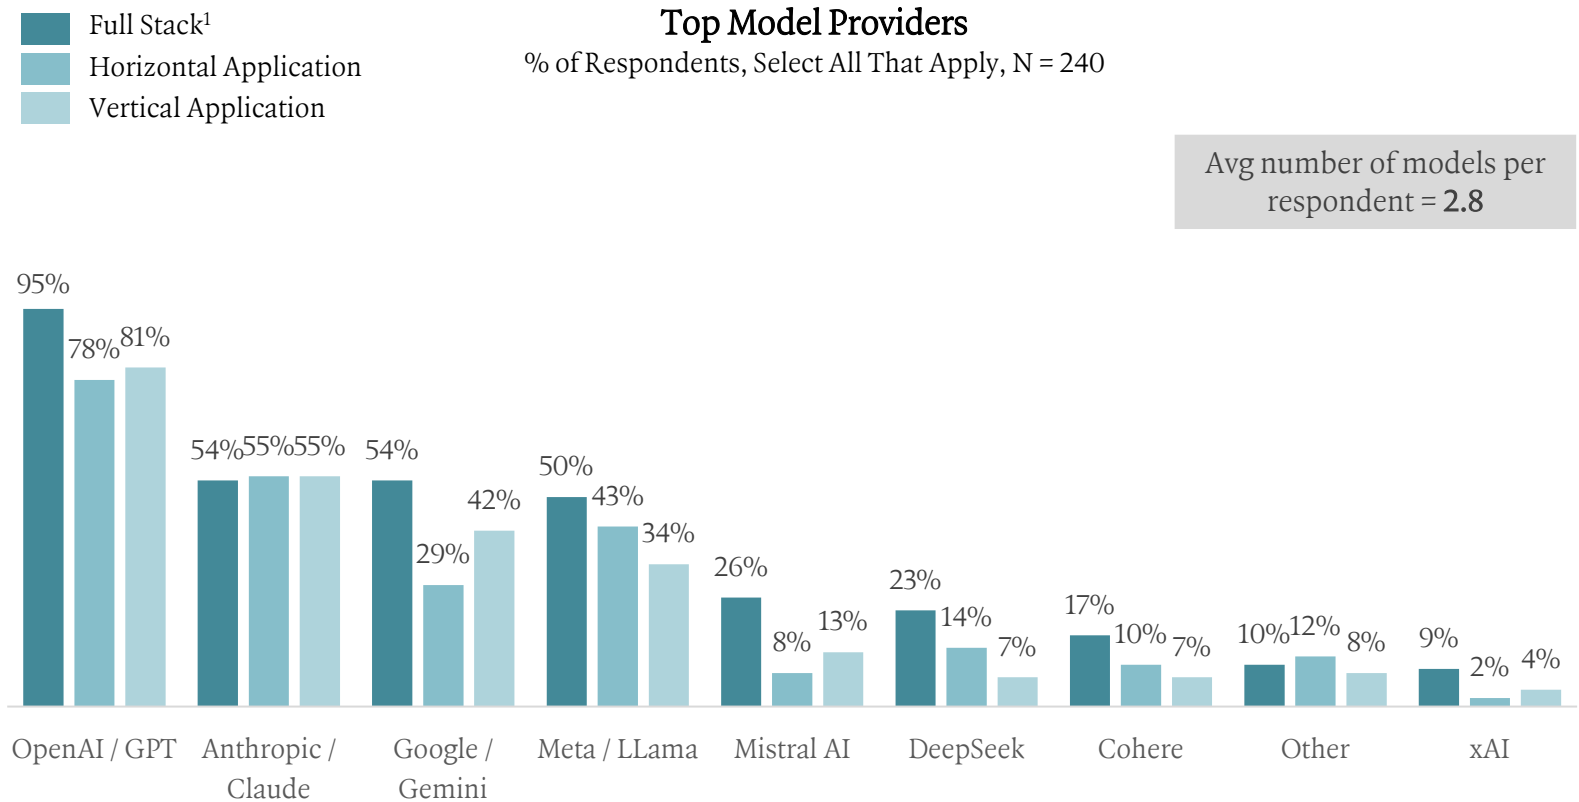
When choosing foundational models for customer-facing use cases, companies prioritize model accuracy above all other factors



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Top Model Providers

OpenAI’s GPT models continue to be the most popular model; however, many companies are increasingly adopting a multi-model approach to AI products across use cases



Notes: (1) Companies building both end user applications and core AI models/technologies

Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Companies are increasingly adopting a **multi-model approach** to AI products, leveraging different providers and models based on use case, performance, cost, and customer requirements.

This flexibility enables them to optimize for diverse applications like cybersecurity, sales automation, and customer service while **ensuring compliance and superior user experience across regions.**

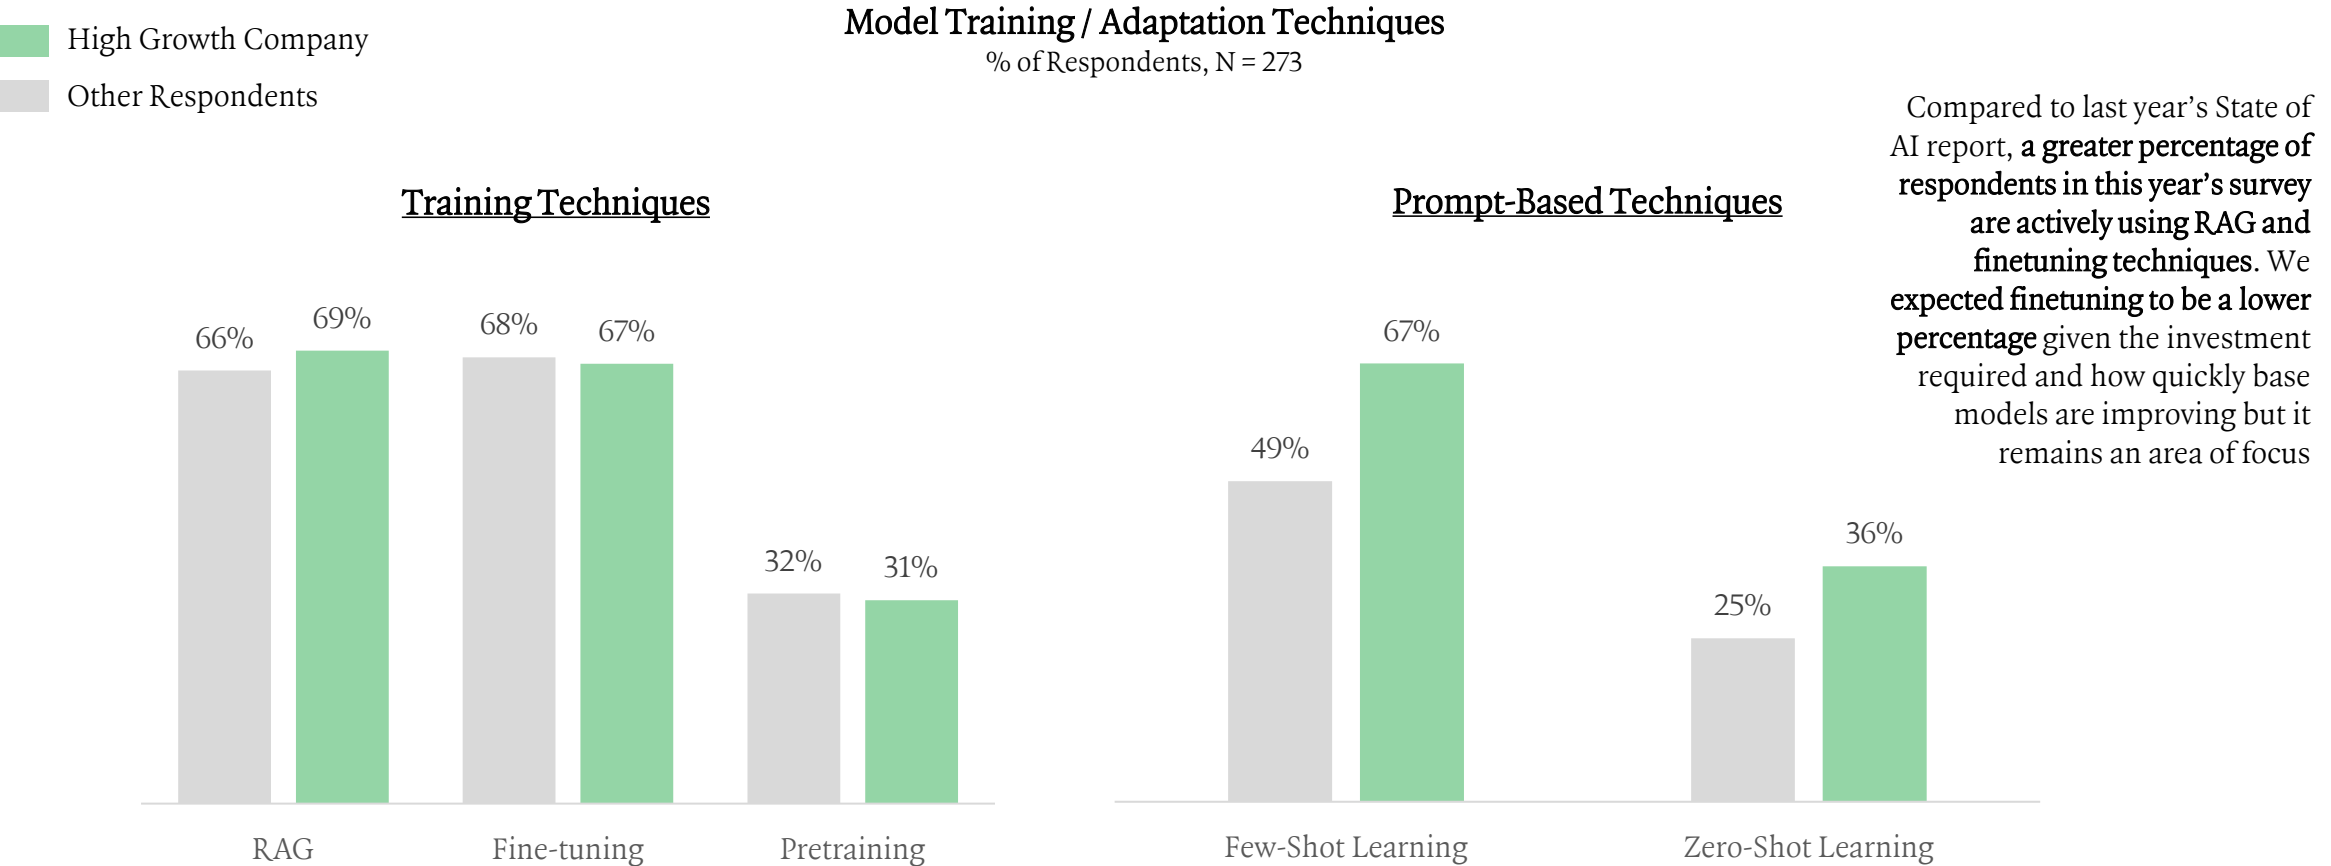
Architectures are being built to support quick **model swaps**, with some leaning toward open-source models for cost and inference speed advantages.

Generally, most respondents are using a combination of OpenAI models and 1-2 other models from the other providers.

“We use different proprietary and 3rd party models because our customers have diverse needs. **Specialized models allow us to better tailor the experiences for our customers** and their use case -- sales automation, agents for customer service and internal tools. In addition, we can **offer our customers more flexible price points and options**, as well as be constantly experimenting with new models and business opportunities.
VP Product, \$1B+ Revenue, Full Stack AI Company

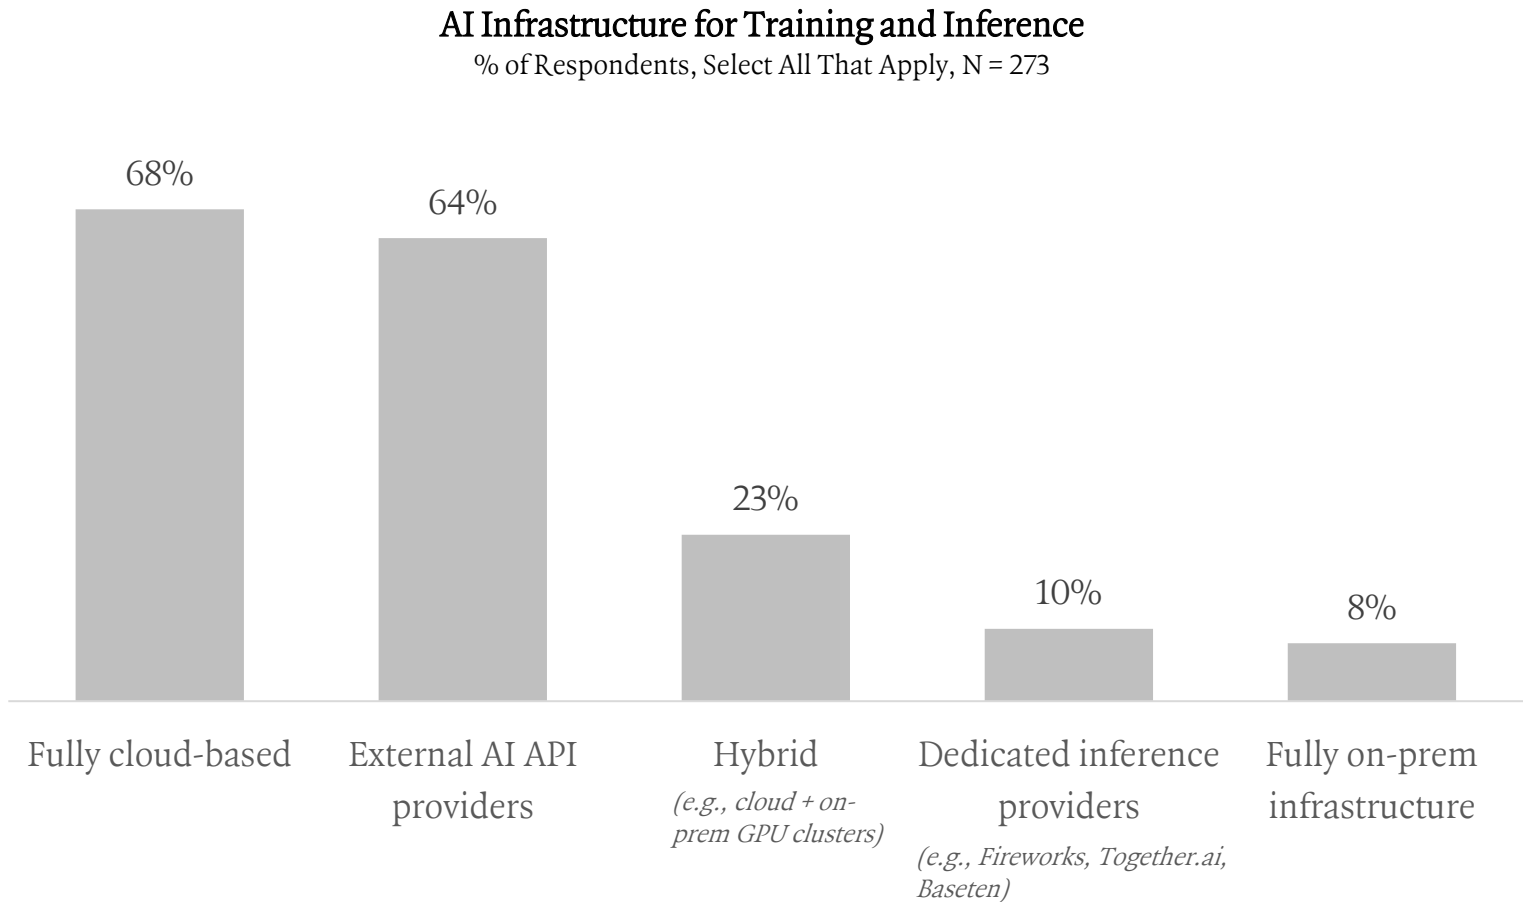
Model Training Techniques

Retrieval augmented generation (RAG) and fine-tuning are the most common model training techniques; high-growth companies tend to use a greater variety of prompt-based techniques



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 are using cloud-based solutions and AI API providers for training and inference



Most organizations are clearly leaning into **fully managed AI solutions** - 68% operate entirely in the cloud and 64% rely on external AI API providers - because this model **minimizes upfront capital outlay and operational complexity, while maximizing speed-to-market**. However, this reliance also means vendor selection, SLA negotiation, and cost-per-call management have become strategic priorities rather than just technical considerations.

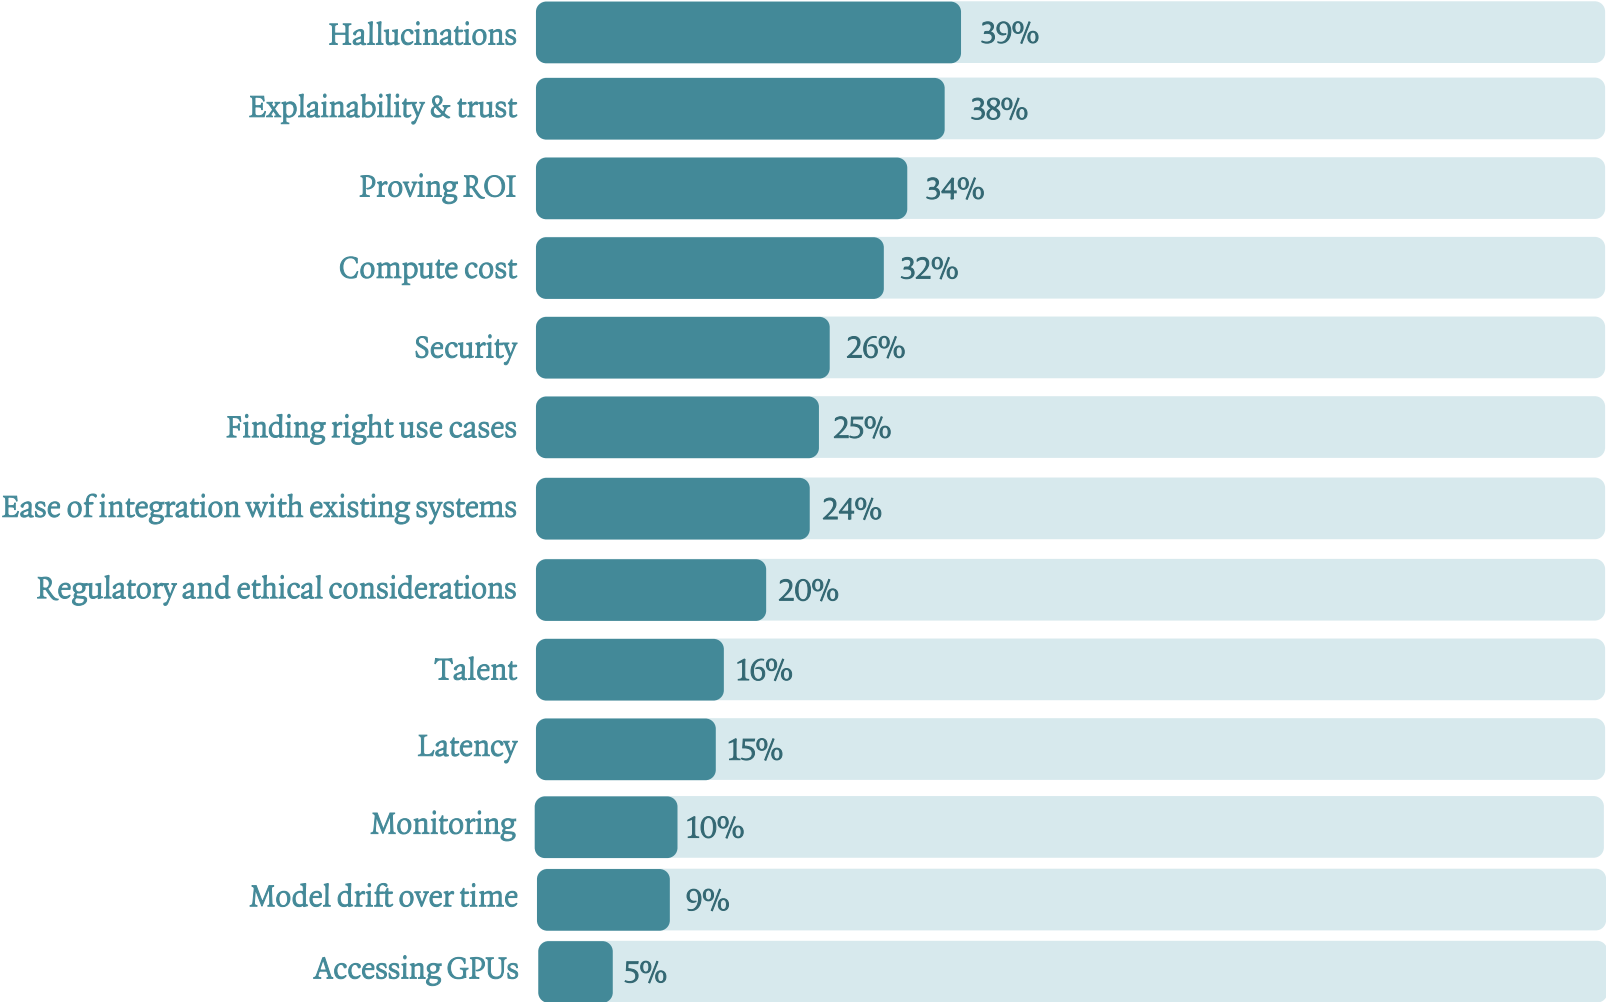
Meanwhile, **only 23% of teams use a hybrid approach** and fewer than 1 in 10 maintain on-prem or dedicated inference **infrastructure**, underscoring that these models remain niche, adopted primarily in scenarios where control, compliance, or specialized performance justify the extra overhead. As real-time AI use cases grow, there's an **emerging opportunity for turnkey inference platforms to capture more share**, but any move away from fully managed services will hinge on a clear business case or regulatory imperative.

Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Top challenges noted by companies when deploying models include hallucinations, explainability / trust, and proving ROI

Challenges in Model Deployment

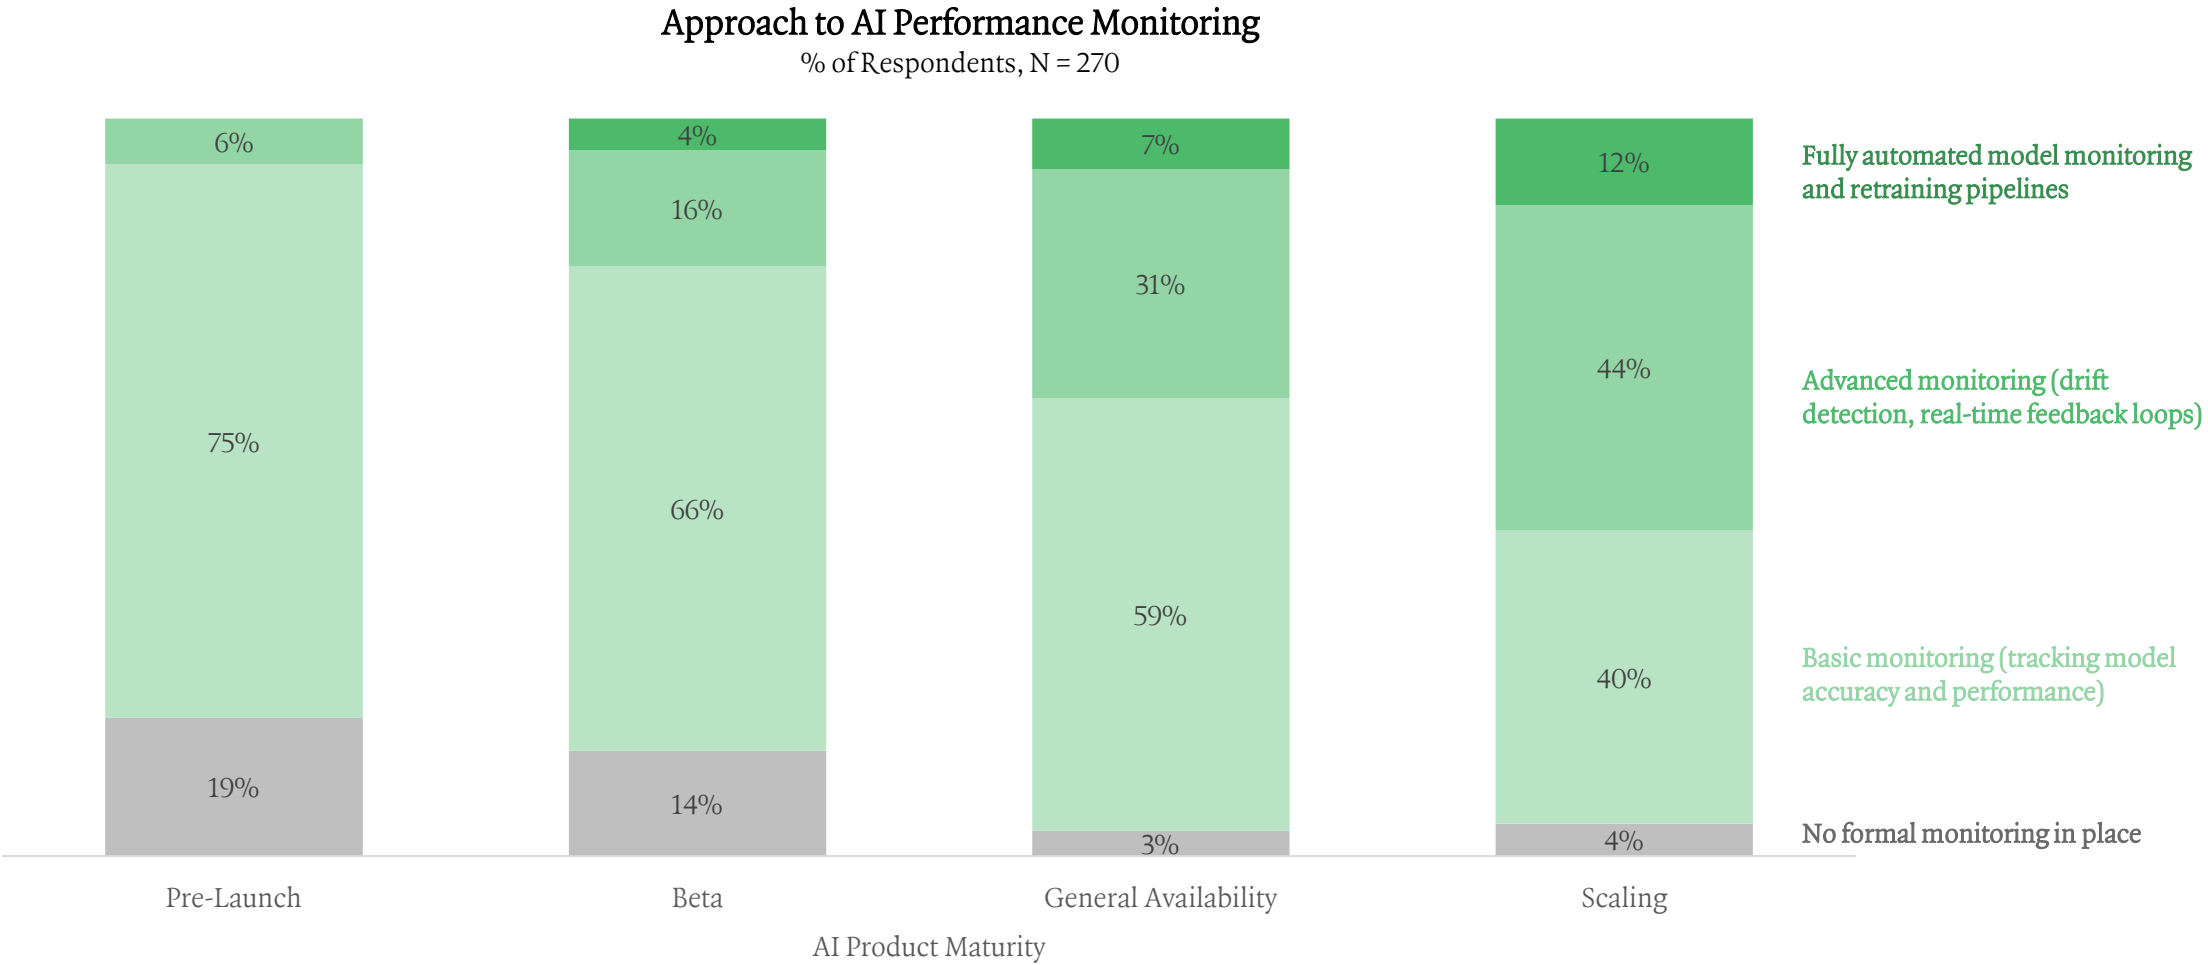
% of Respondents who ranked each aspect in Top 3, N = 273



Explainability and trust ranked higher for companies building vertical AI applications, who may deal with additional compliance and legal restrictions in regulated industries like healthcare

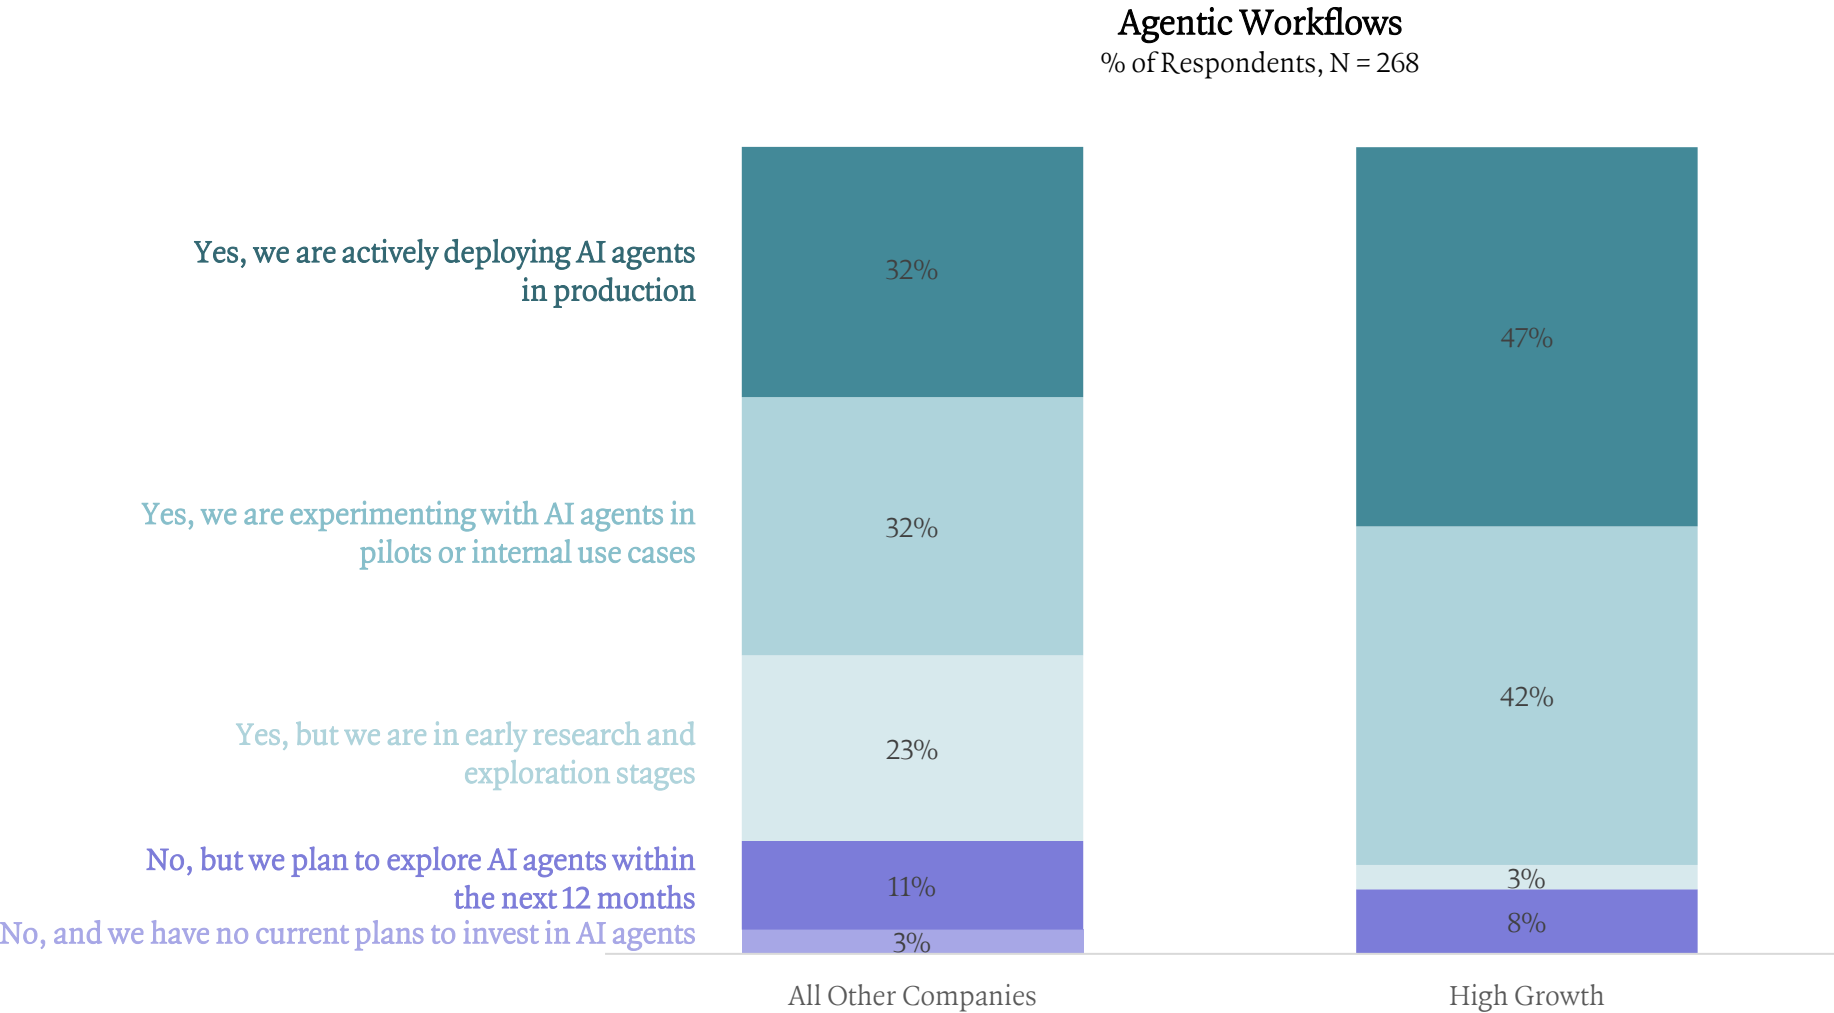
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

As AI products scale, performance monitoring becomes more important with many scaled AI products offering some kind of advanced performance monitoring



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

A significant number of companies are evaluating agentic workflows, with high growth AI companies more actively deploying AI agents in production



Many of our users like the insights and analytics we are surfacing but are unwilling to commit the time to fully explore the information housed in the product. We are looking to build out AI agents that **effectively use the product for the end-users to surface worthwhile user-journeys** and bring the end-user along for them.

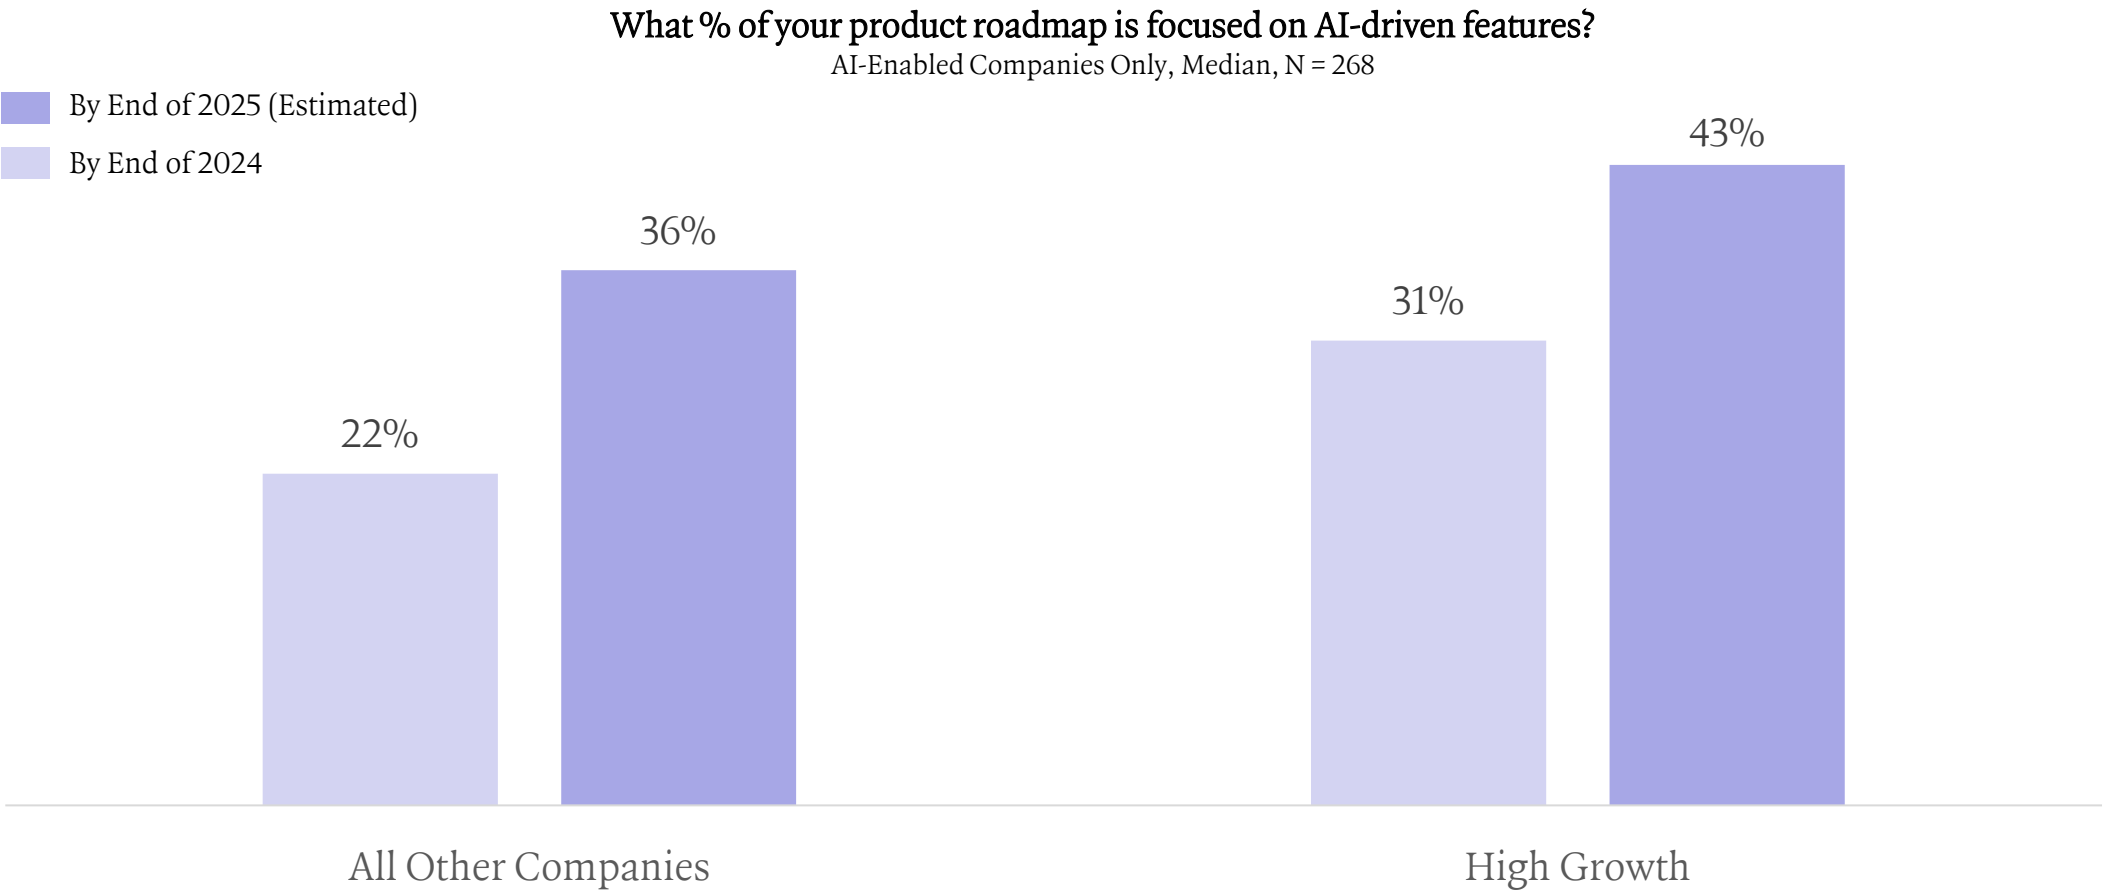
VP Product, \$10-25M Revenue, Full Stack AI Company

Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Go-to-Market Strategy & Compliance



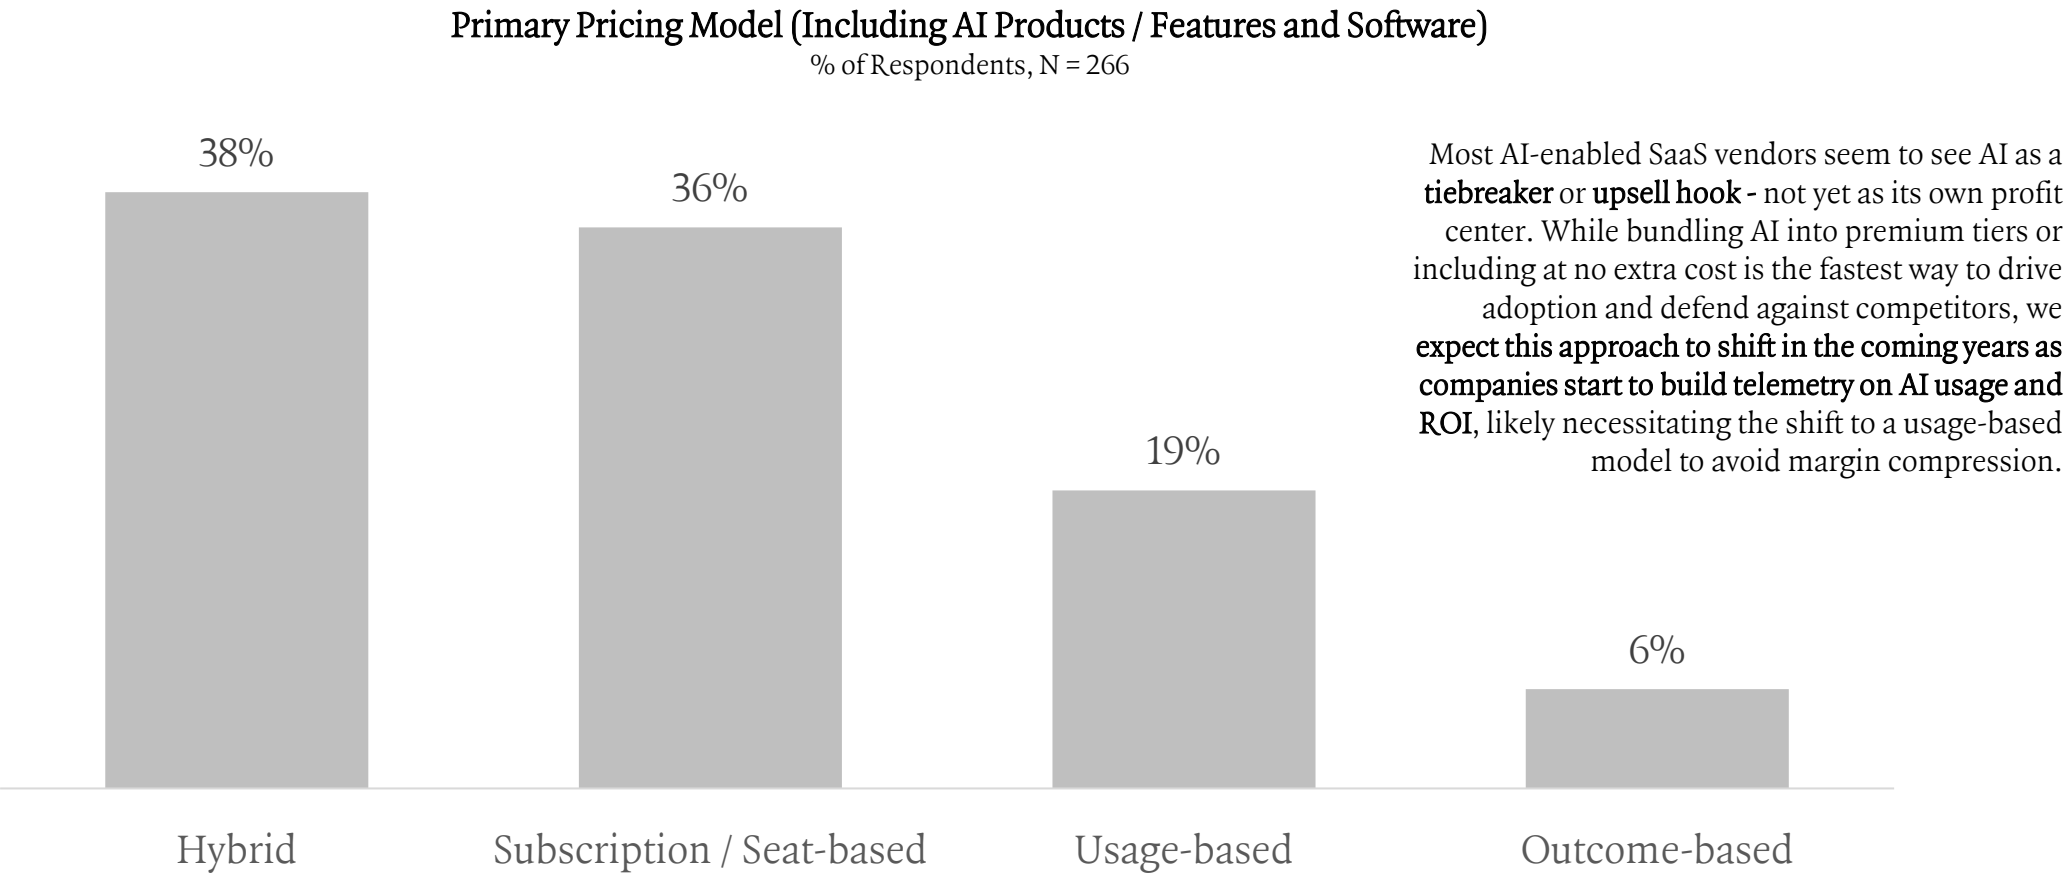
For AI-enabled companies, around 20-35% of their product roadmap has been focused on AI-driven features with high-growth companies dedicating closer to 30-45% of their roadmap to AI-driven features



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Primary Pricing Model

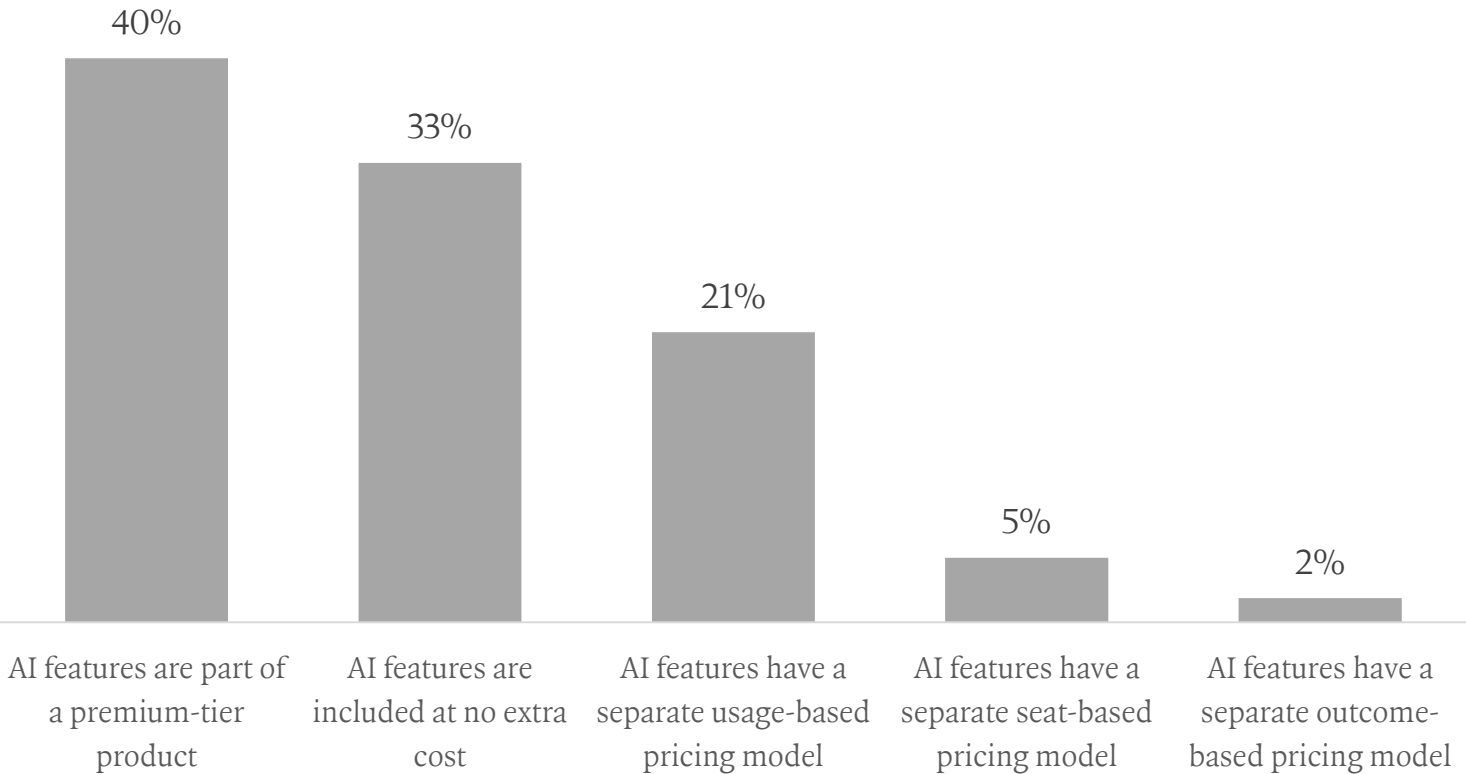
Many companies are using a hybrid pricing model which includes a combination of subscription / plan-based pricing along with either usage-based or outcome-based pricing



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

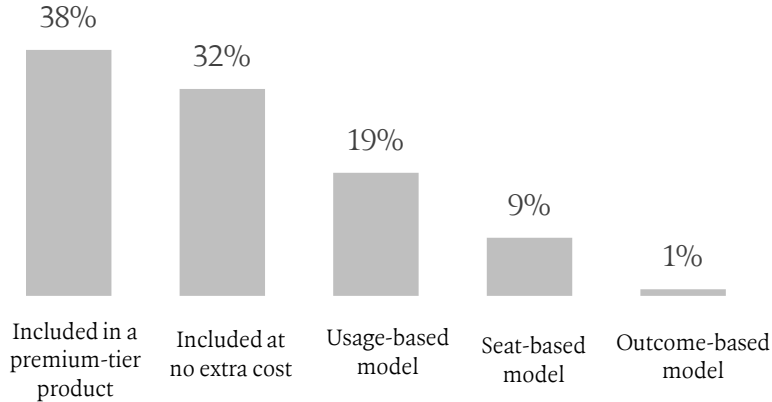
Currently, most AI-enabled companies are either including AI features as part of a premium-tier product or including them at no extra cost

Primary Pricing Model for AI Features / Products
AI-Enabled Companies Only, % of Respondents, N = 174



ICONIQ Cross-Functional Insight

In our [2025 State of GTM report](#), we asked this same question to GTM leaders, and their responses largely aligned with R&D leaders – further reinforcing the consistency of this trend across the market.



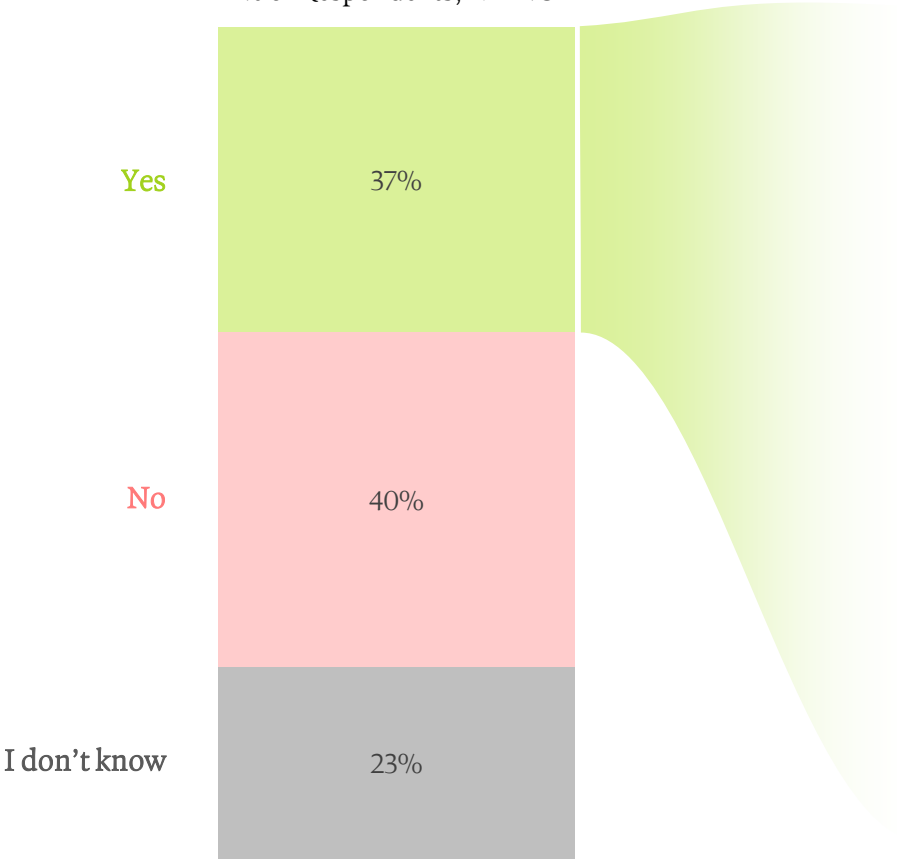
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Pricing Changes

40% of companies have no plans to change pricing, but 37% of respondents are exploring new pricing models based on consumption, ROI, and usage tiers

Plans to Change AI Pricing in Next Twelve Months

% of Respondents, N = 273



Factoring in ROI

“We would like to **integrate willingness to pay and clear connection to ROI** outcomes into our pricing model”
VP Product, \$100-150M Revenue, Full Stack AI Company

“We are observing if AI capabilities deliver extra value to customer. Once we have critical adoption and proof of added value, we might **segment the current tiers of our platform** (i.e. create a top tier with the full AI /agents, a limit on the basic, and enterprise tiers)”
VP Product, \$100-150M Revenue, Full Stack AI Company

Consumption and Outcome-Based Pricing

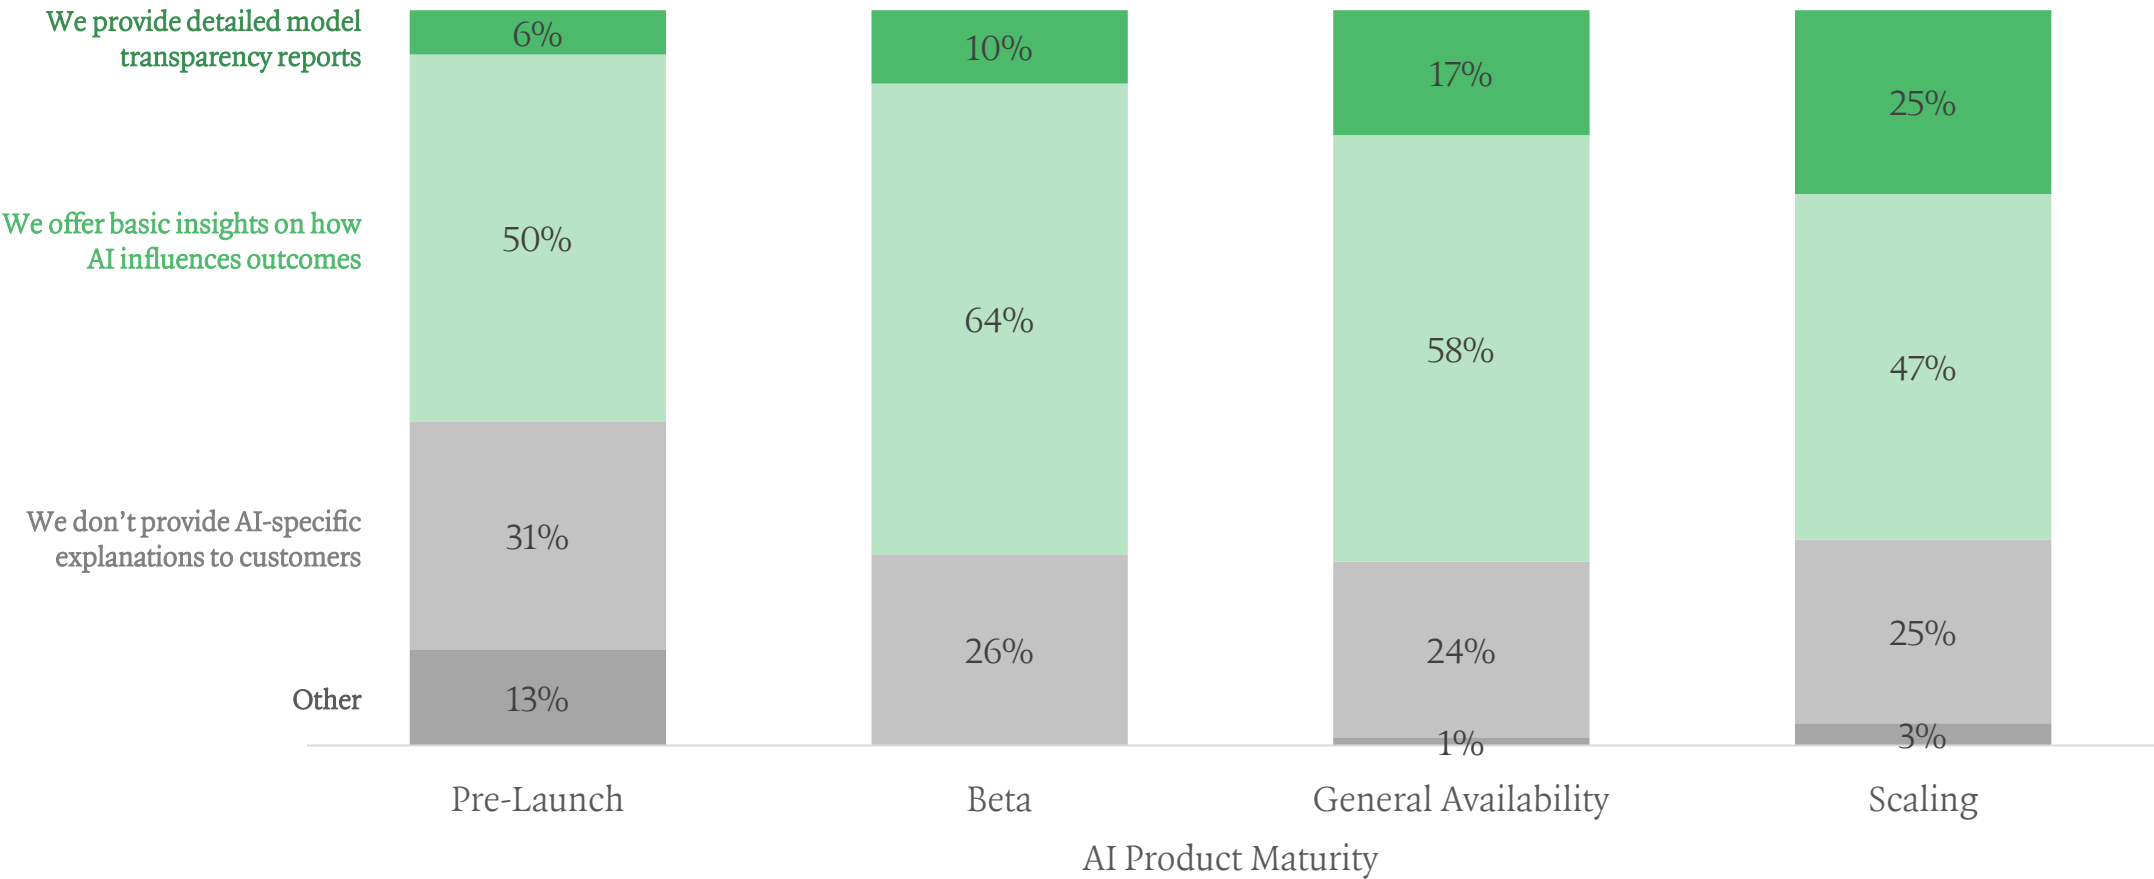
“We will **complement premium tier model pricing with pricing models centered around consumption**. I expect we will also experiment with outcome-based pricing but it is unclear how we will structure pricing in such a way that it **allows customers to accurately budget for these costs.**”
VP Product, \$100-150M Revenue, Full Stack AI Company

“The subscription model is not working for us. **Power users tend to use a lot resulting in negative margins** considering LLM API costs, while **users who aren't using are at risk of churn**. Considering the high variable cost we are **planning to move to usage based** but bundle usage as a subscription e.g., 10M token per year package”
VP Product, \$100-150M Revenue, Full Stack AI Company

Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

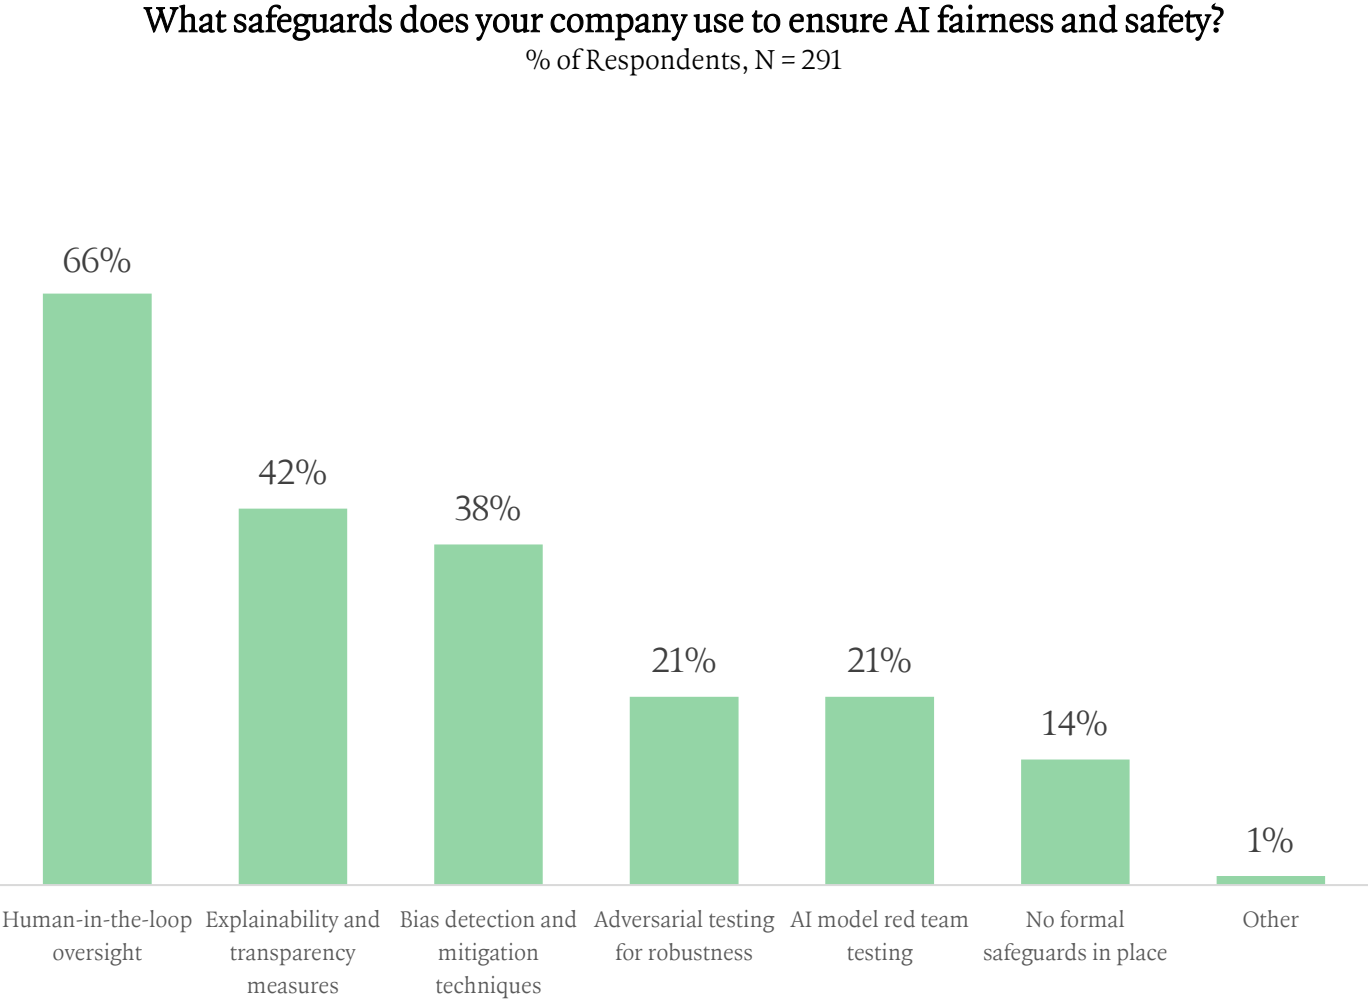
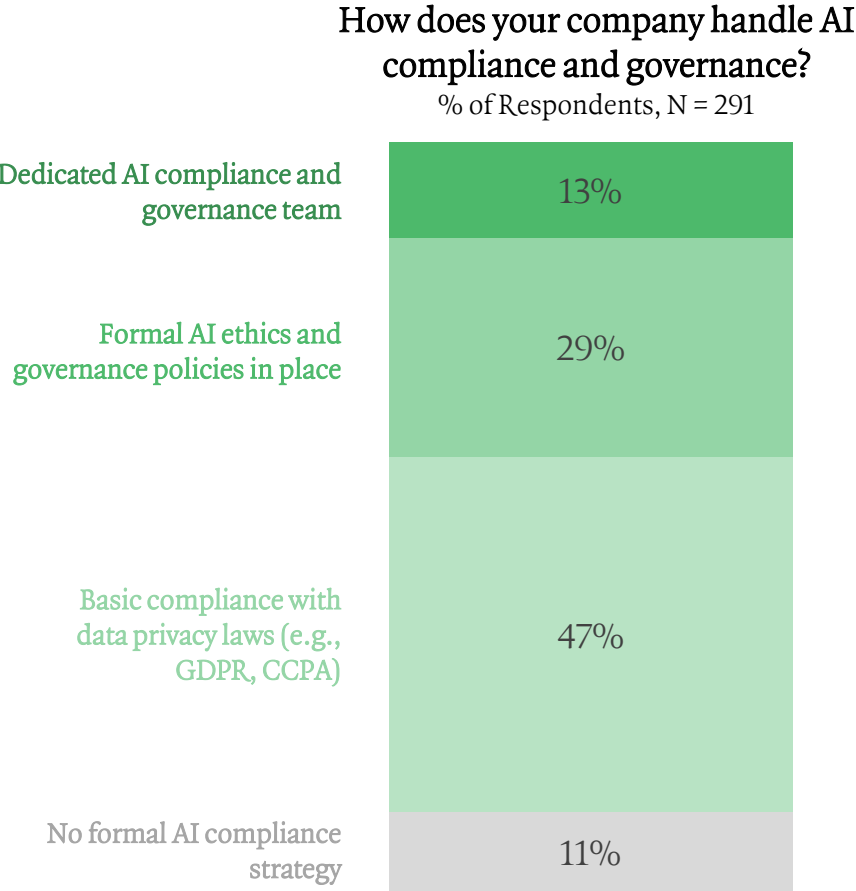
As AI products scale, providing detailed model transparency reports or basic insights on how AI influences outcomes becomes more critical

Strategy for AI Explainability and Transparency to Customers
% of Respondents, N = 266



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 have guardrails around AI ethics and governance policies, with the majority of respondents using human-in-the-loop oversight to ensure AI fairness and safety



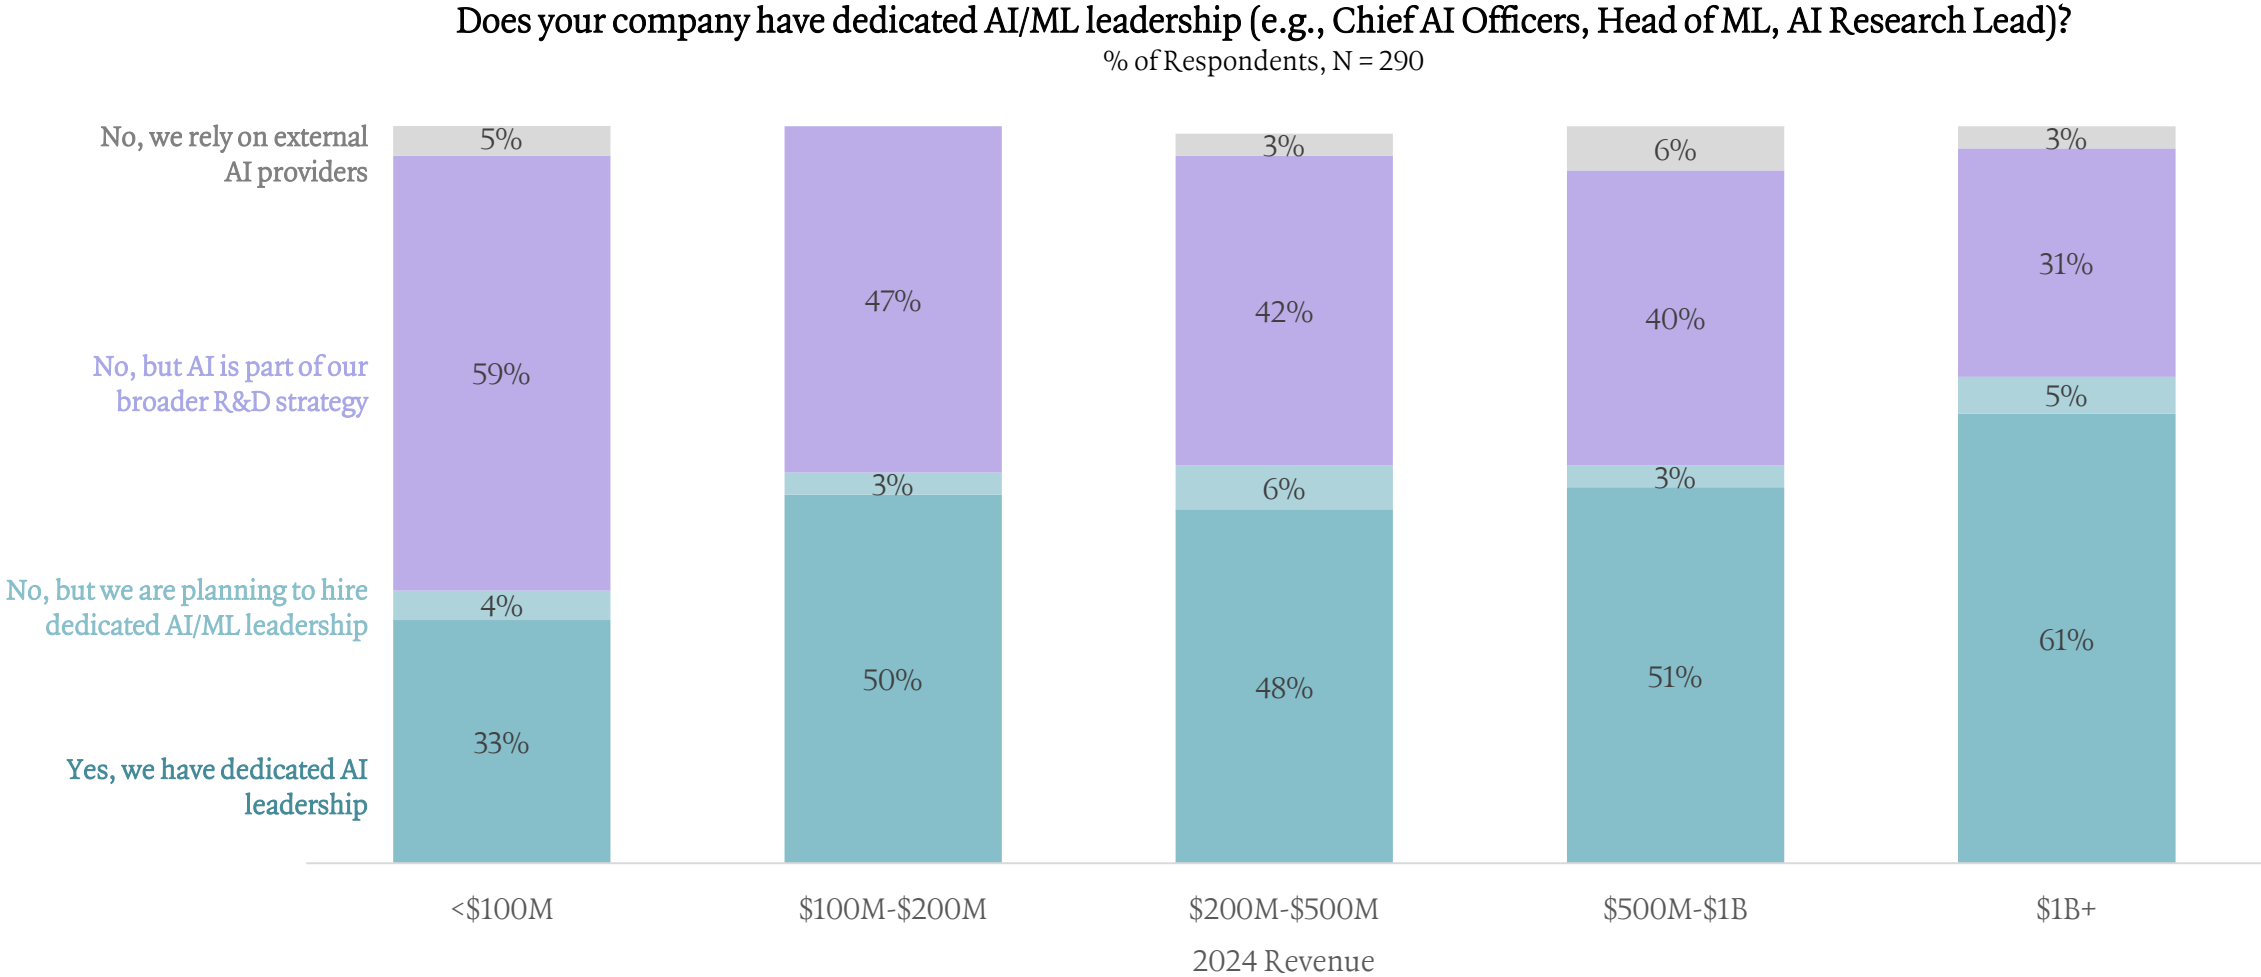
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Organization Structure



Dedicated AI/ML Leadership

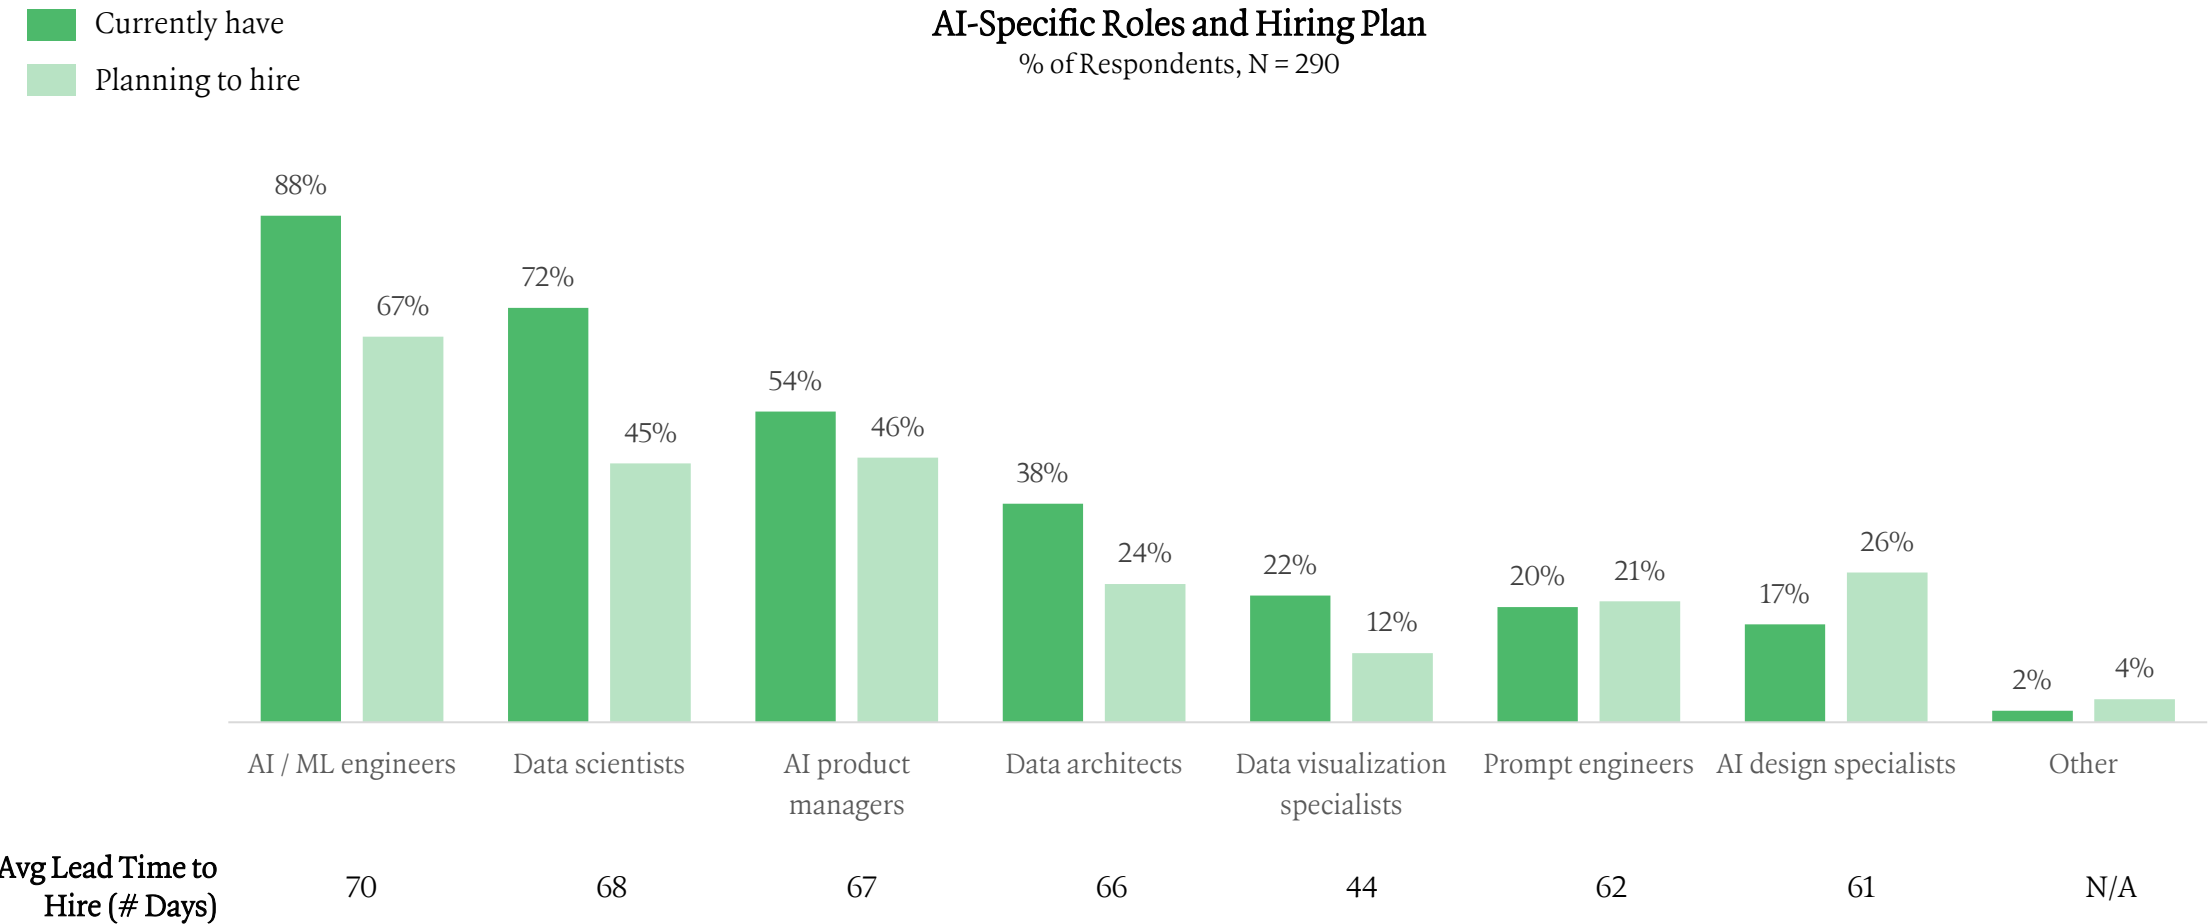
Many companies have dedicated AI leadership by the time they reach \$100M in revenue likely due to increasing operational complexity and the need to have a centralized owner for AI strategy



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

AI-Specific Roles

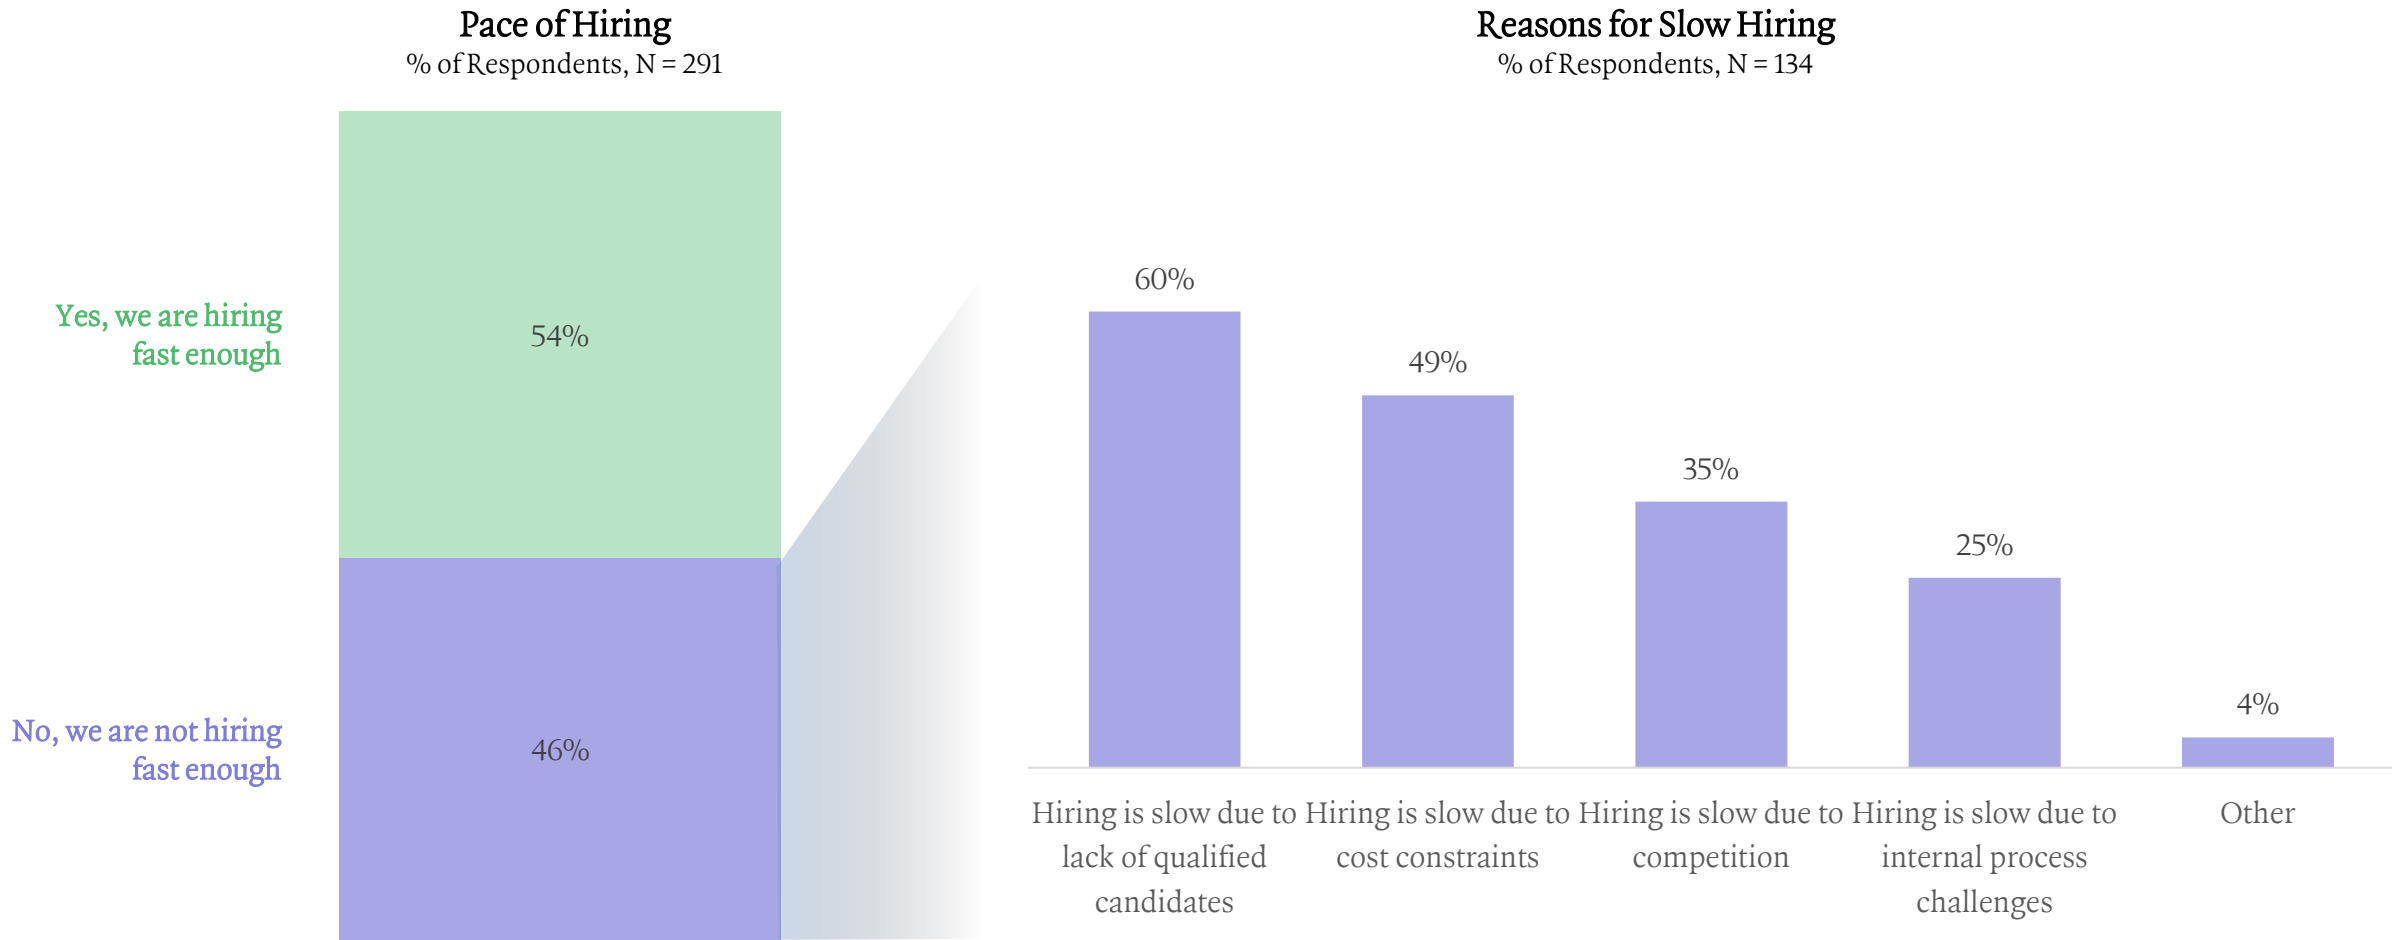
Most companies currently have dedicated AI/ML engineers, data scientists, and AI product managers, with AI/ML engineers taking the longest time on average to hire



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Pace of Hiring

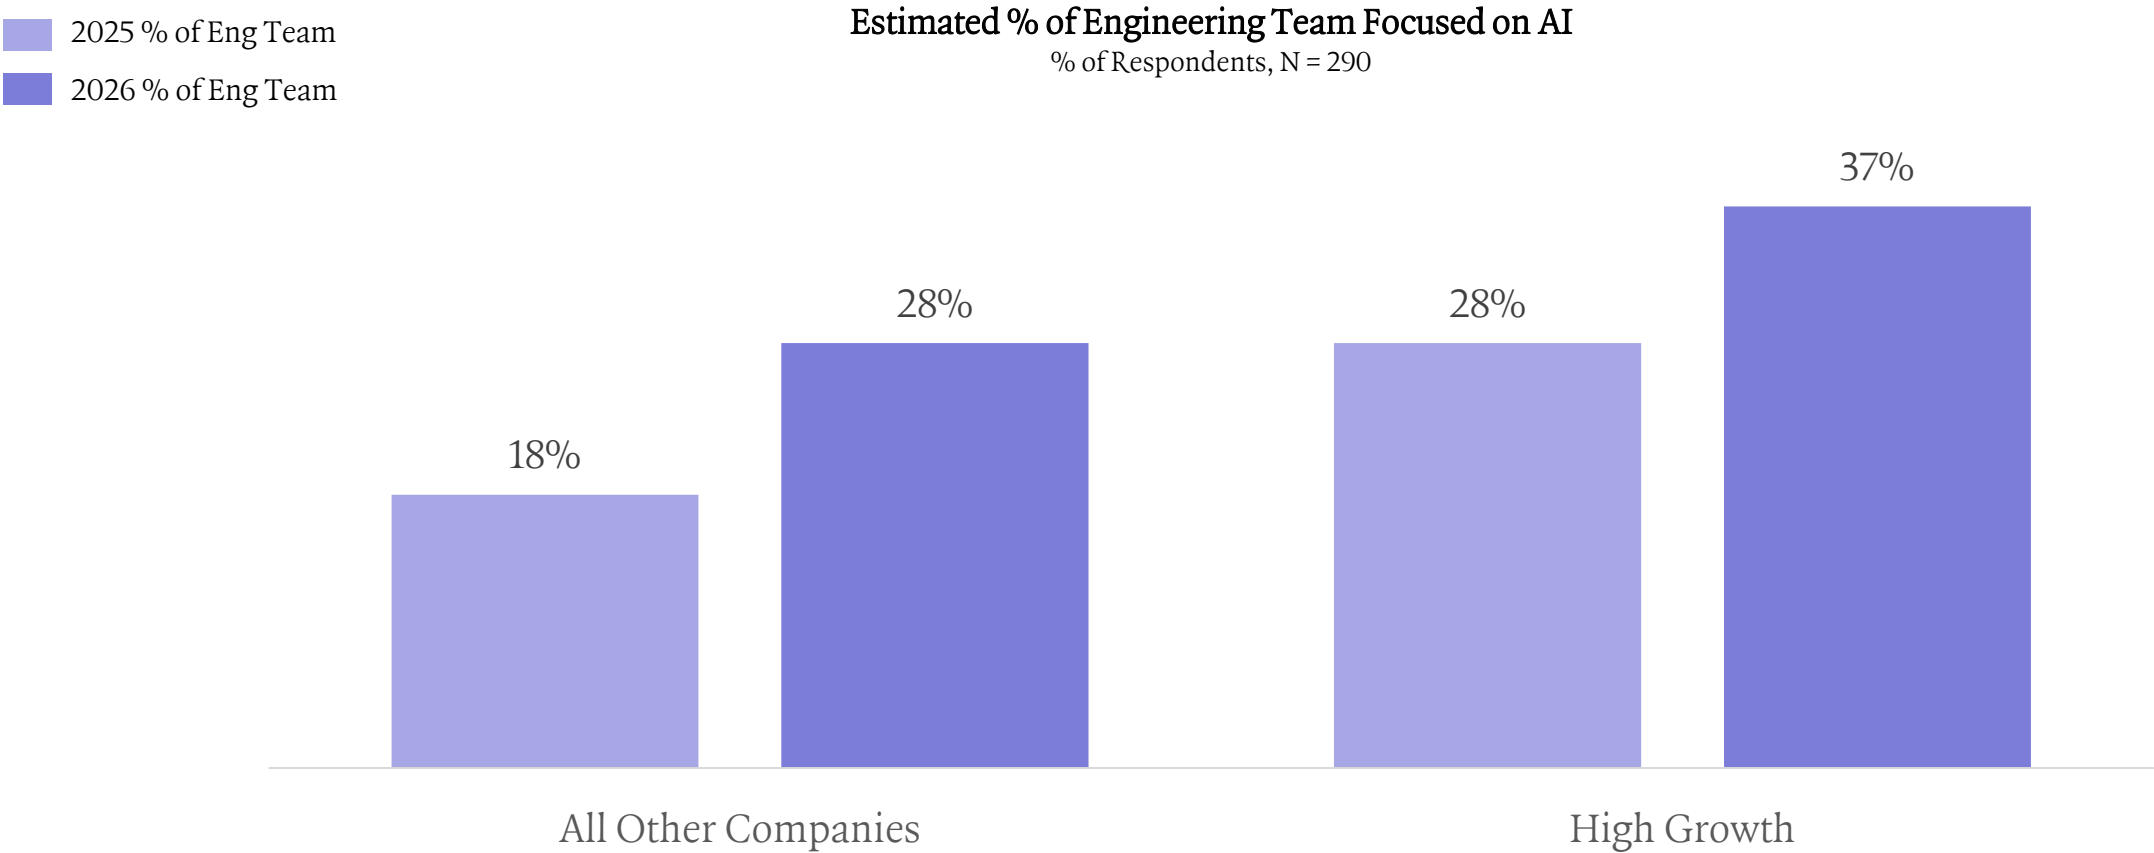
Across respondents, there was a relatively even split in sentiment towards the pace of hiring, with those who felt like they were not hiring fast enough primarily citing lack of qualified candidates as the main constraint



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

% of Engineering Team Focused on AI

On average, companies plan to have 20-30% of their engineering team focused on AI, with high growth companies having a higher proportion of their engineering team focused on AI



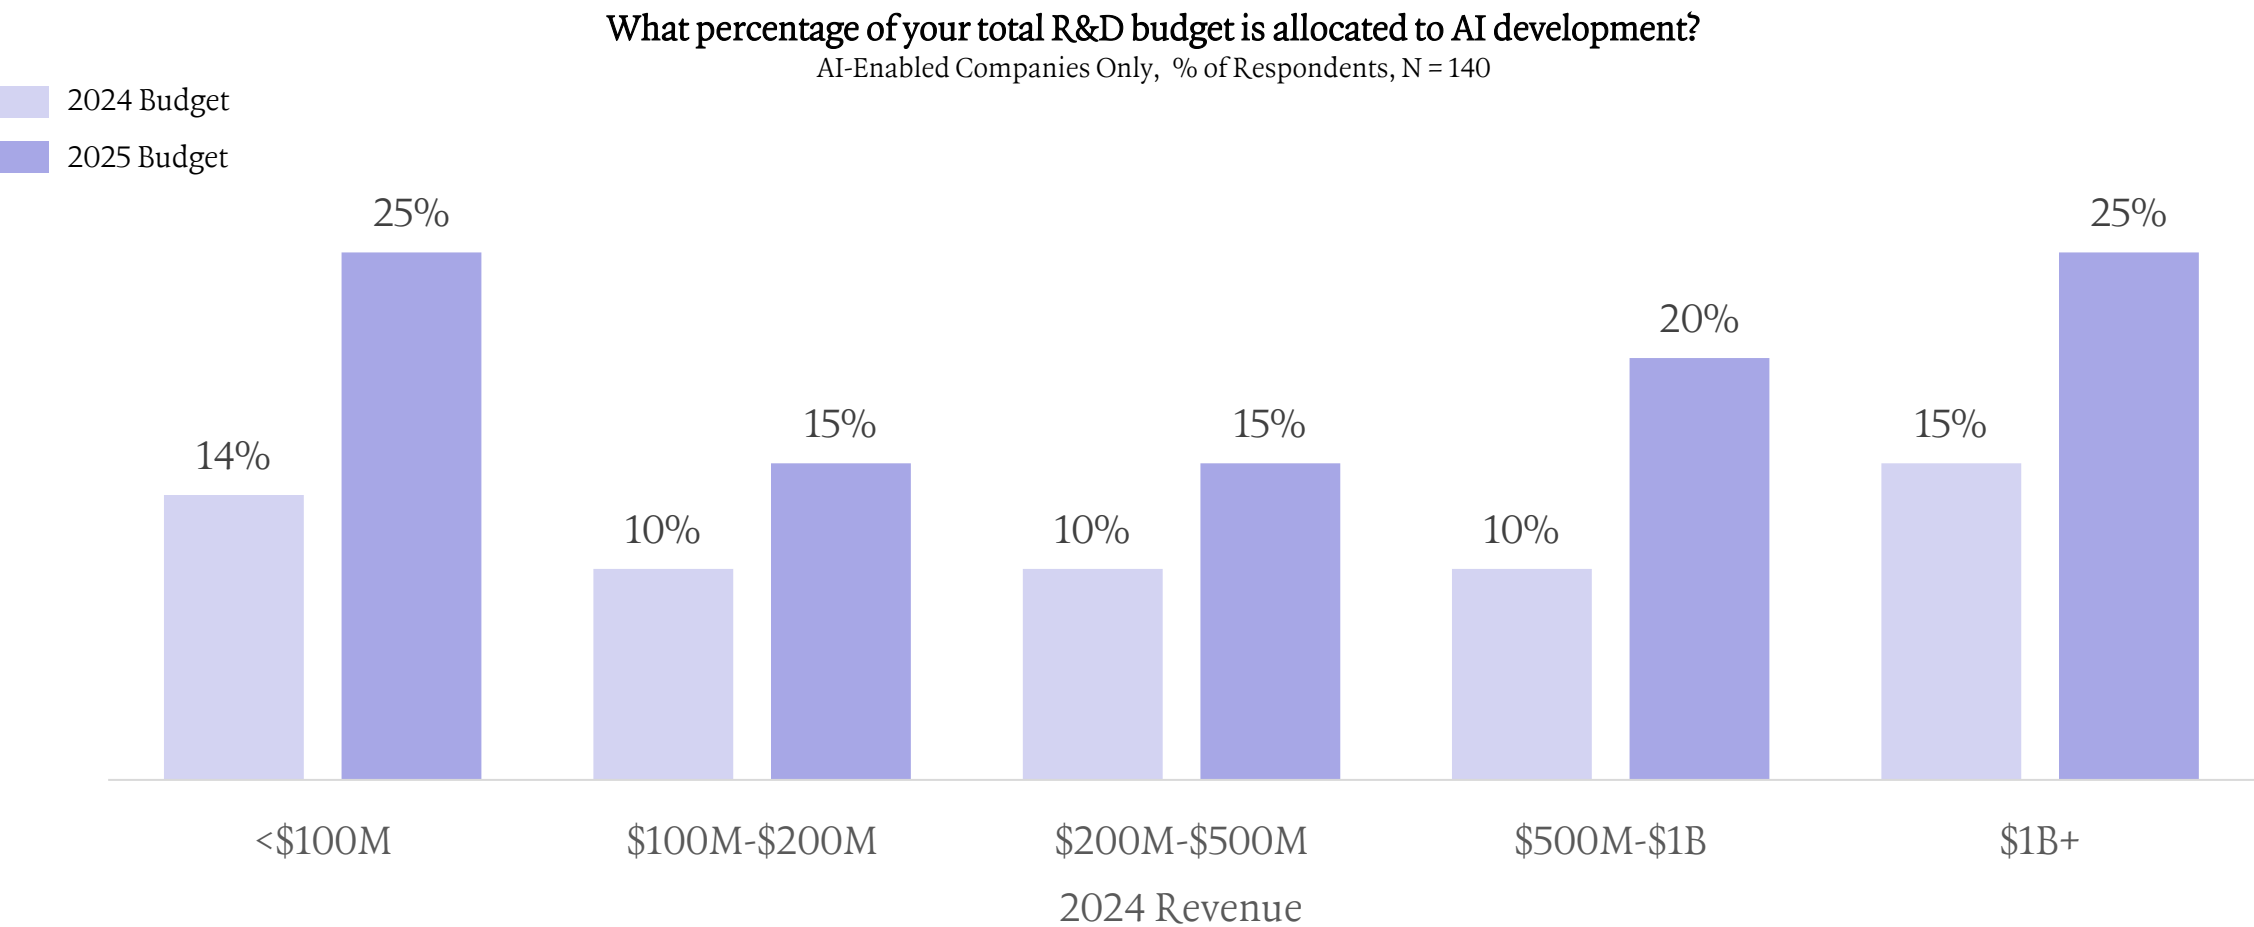
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

AI Costs



AI Development Spend

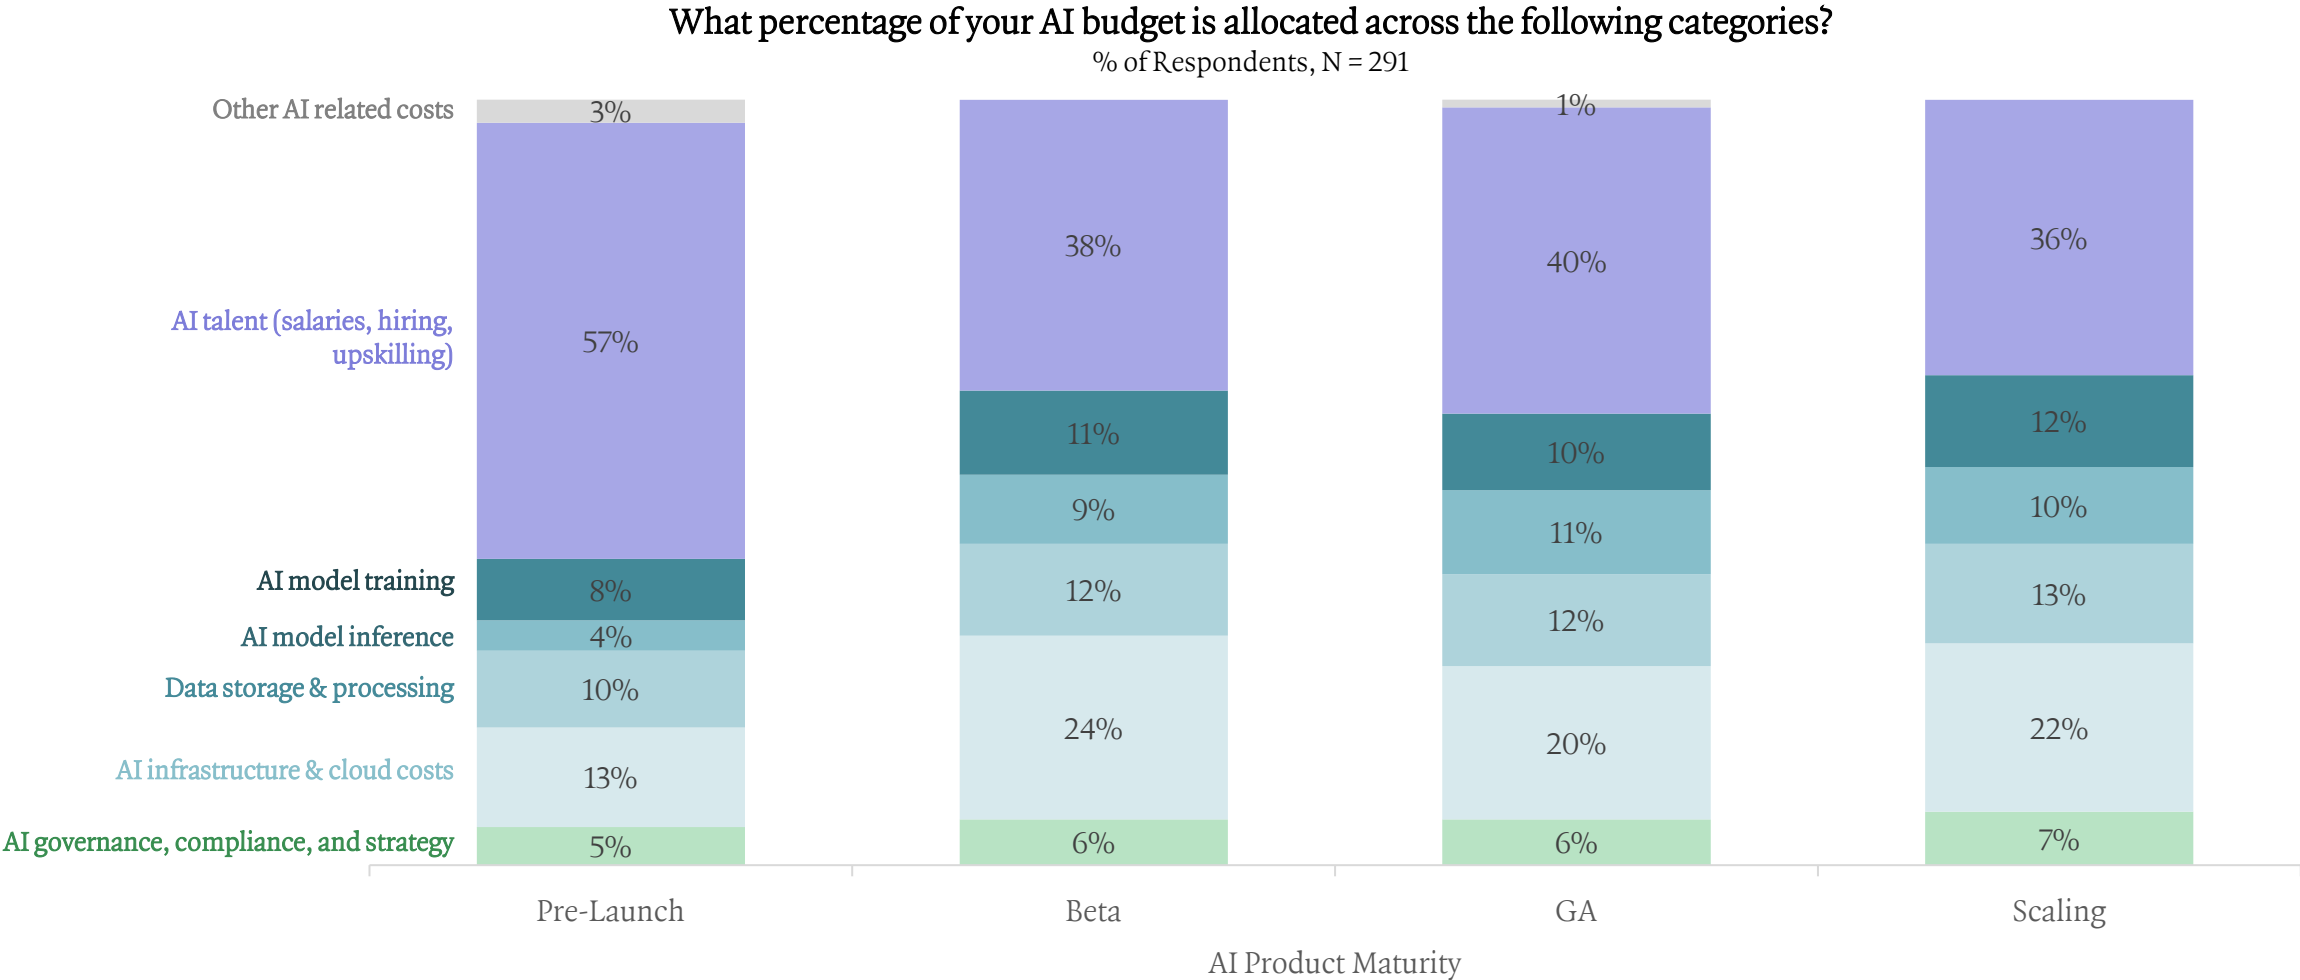
On average, companies are allocating ~10-20% of their R&D budget to AI development, with most companies planning to increase spend on AI in 2025



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

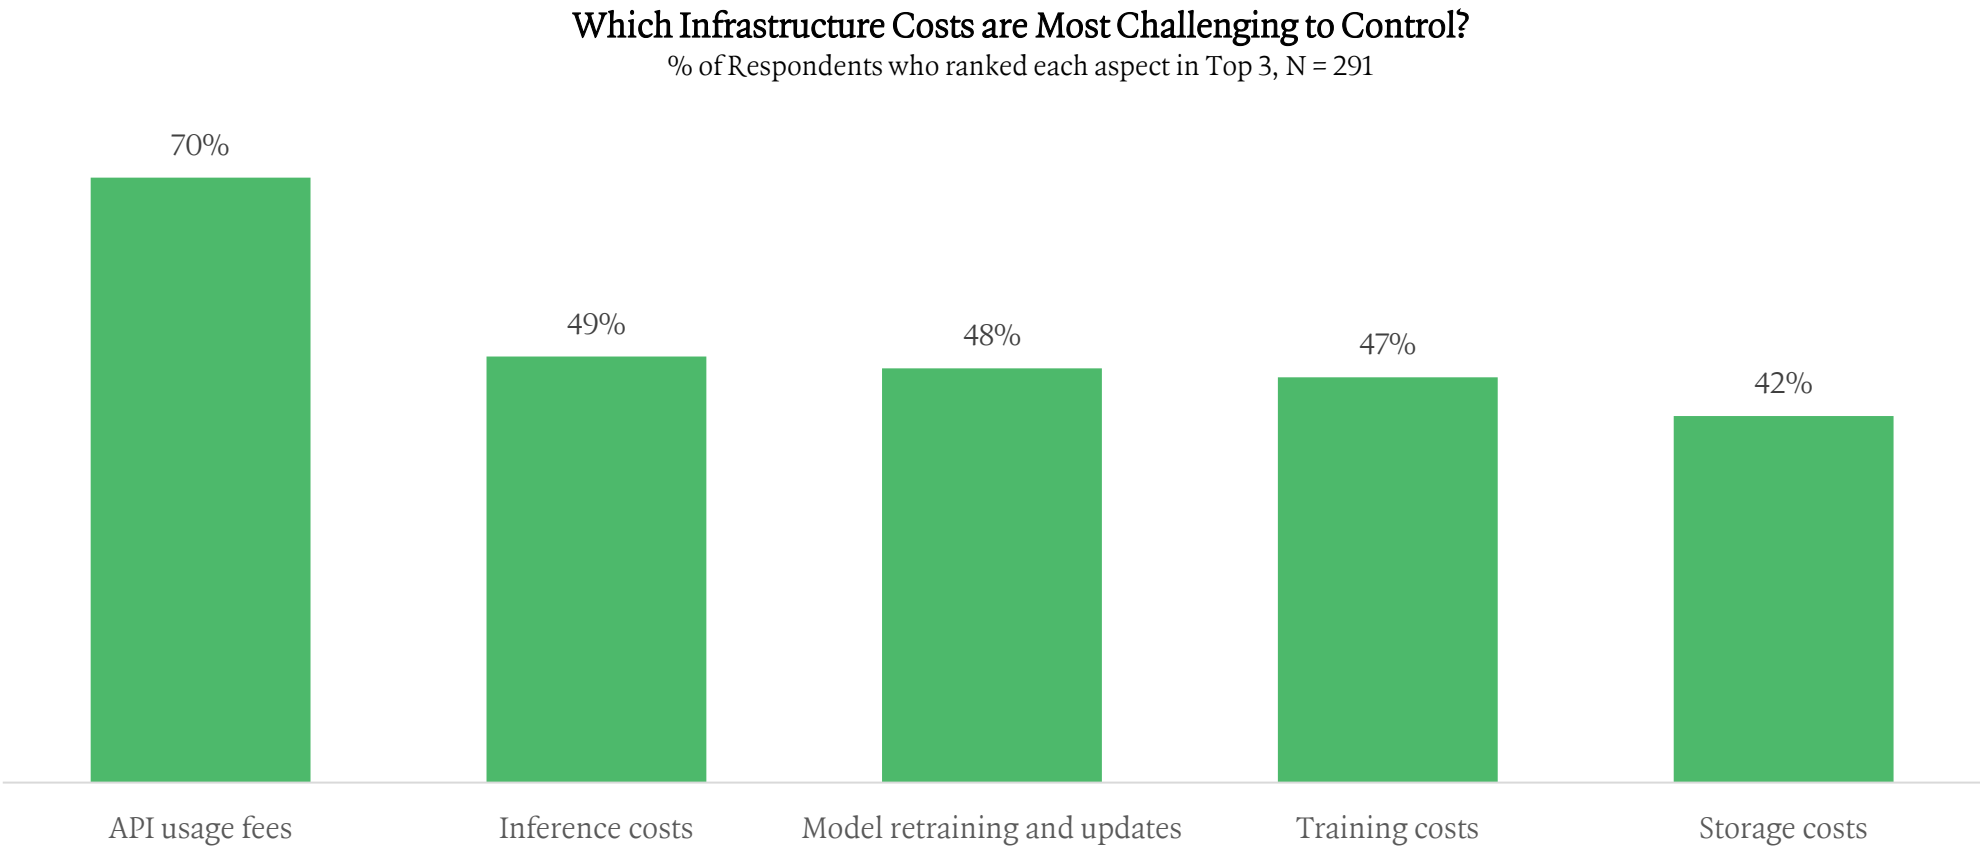
Budget Allocation

As AI products scale, the cost of talent tends to go down as a total proportion of spend; conversely, infrastructure and compute costs tend to increase as products start to see market traction



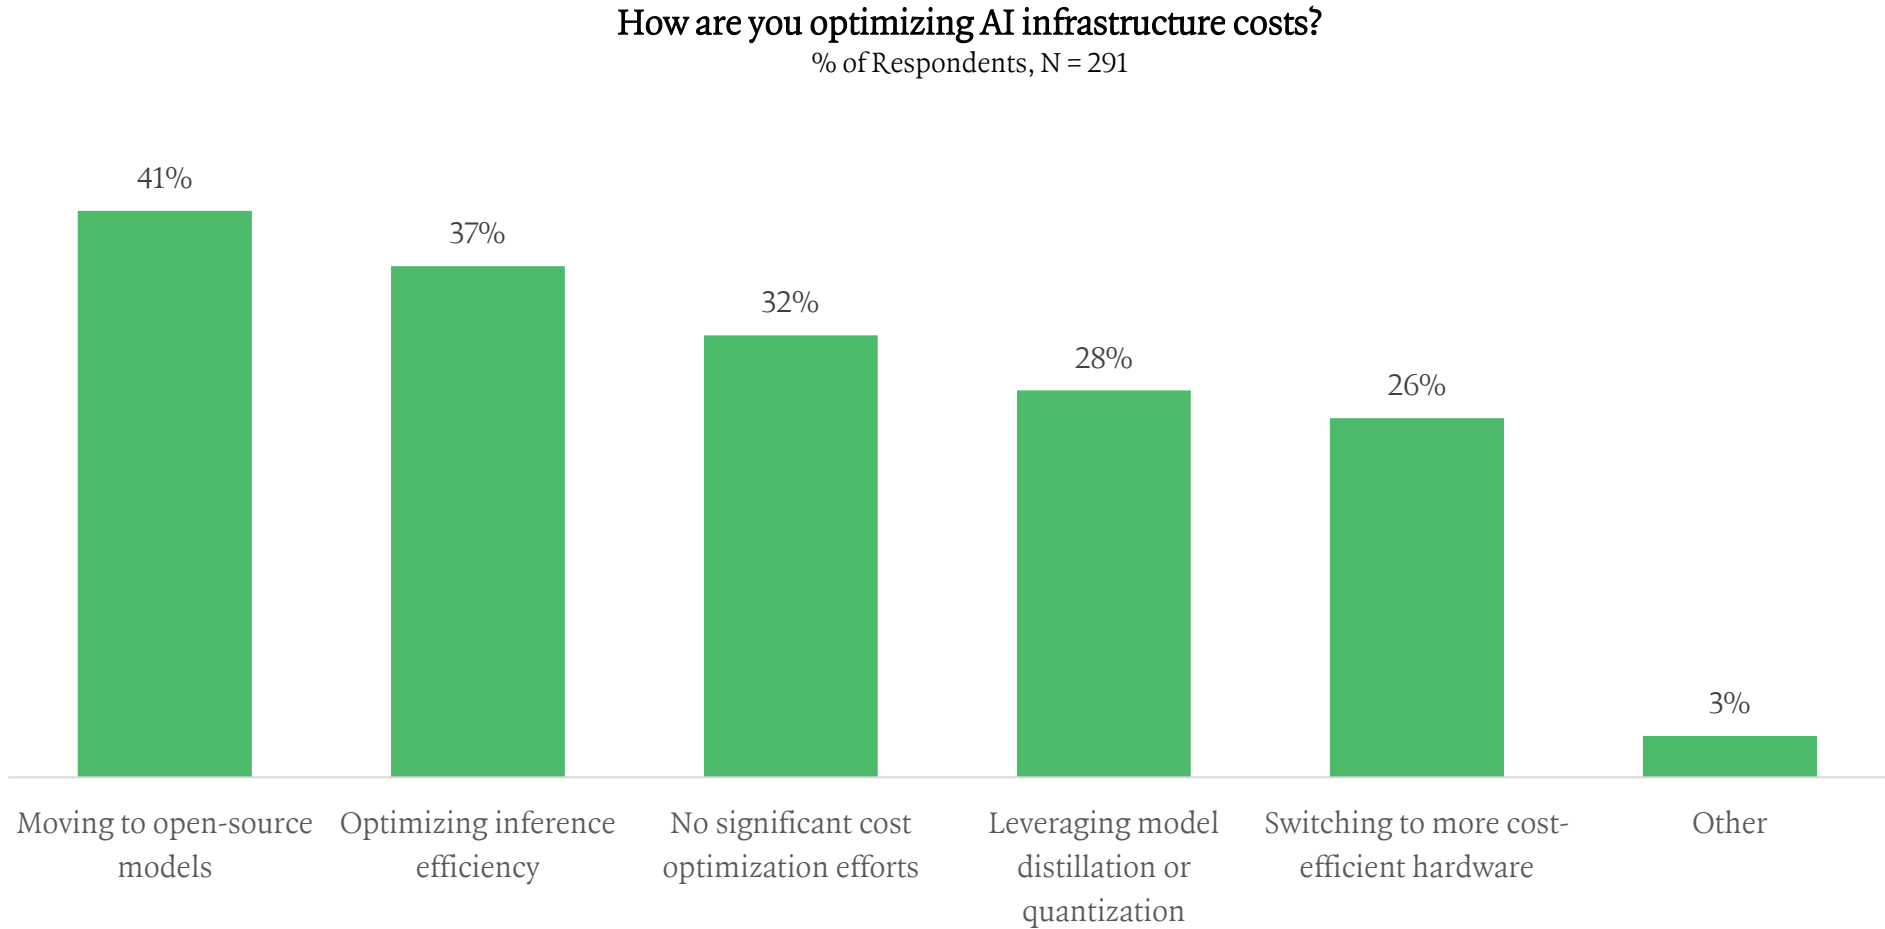
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Of the various infrastructure costs, respondents cited API usage fees as the cost most challenging to control, suggesting companies face the most unpredictability around variable costs tied to external API consumption



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

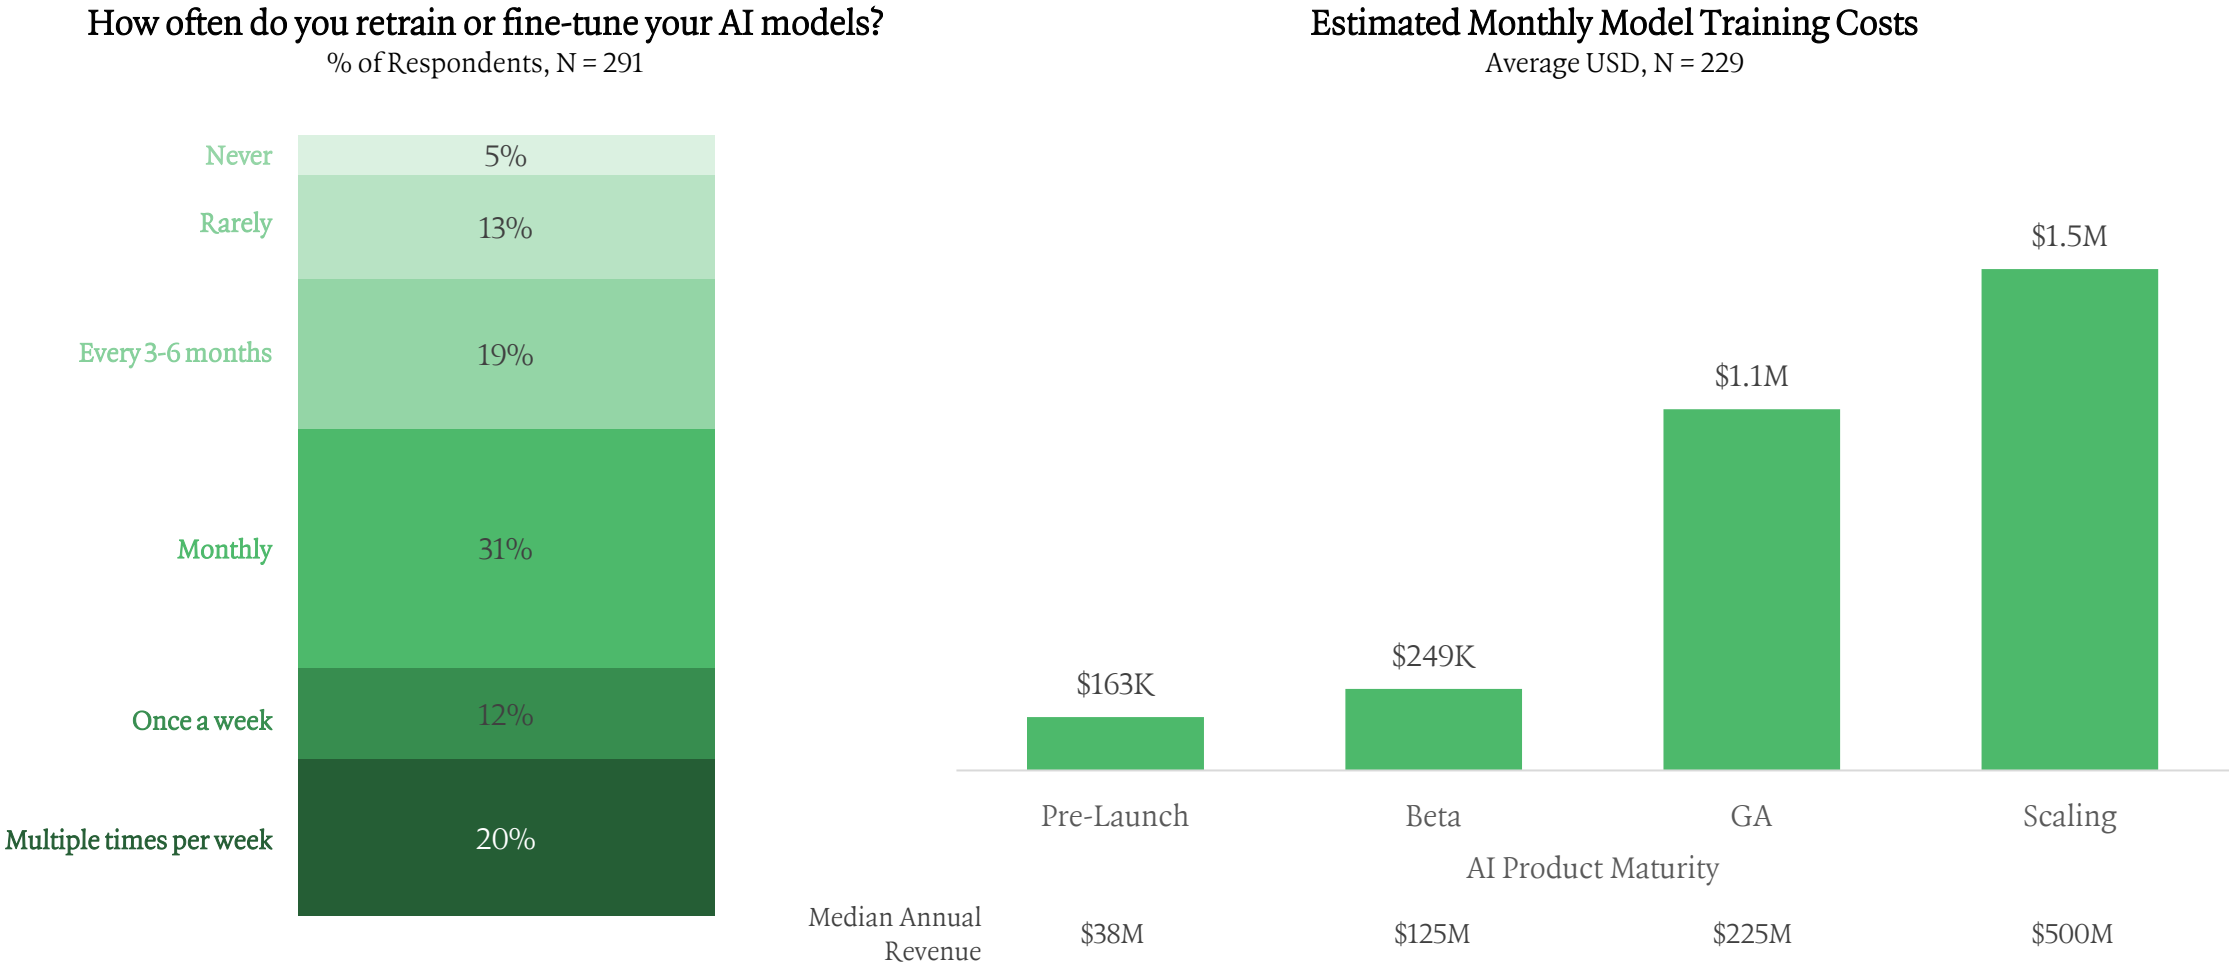
To cut AI infrastructure costs, organizations are exploring open-source models and ways to optimize inference efficiency



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

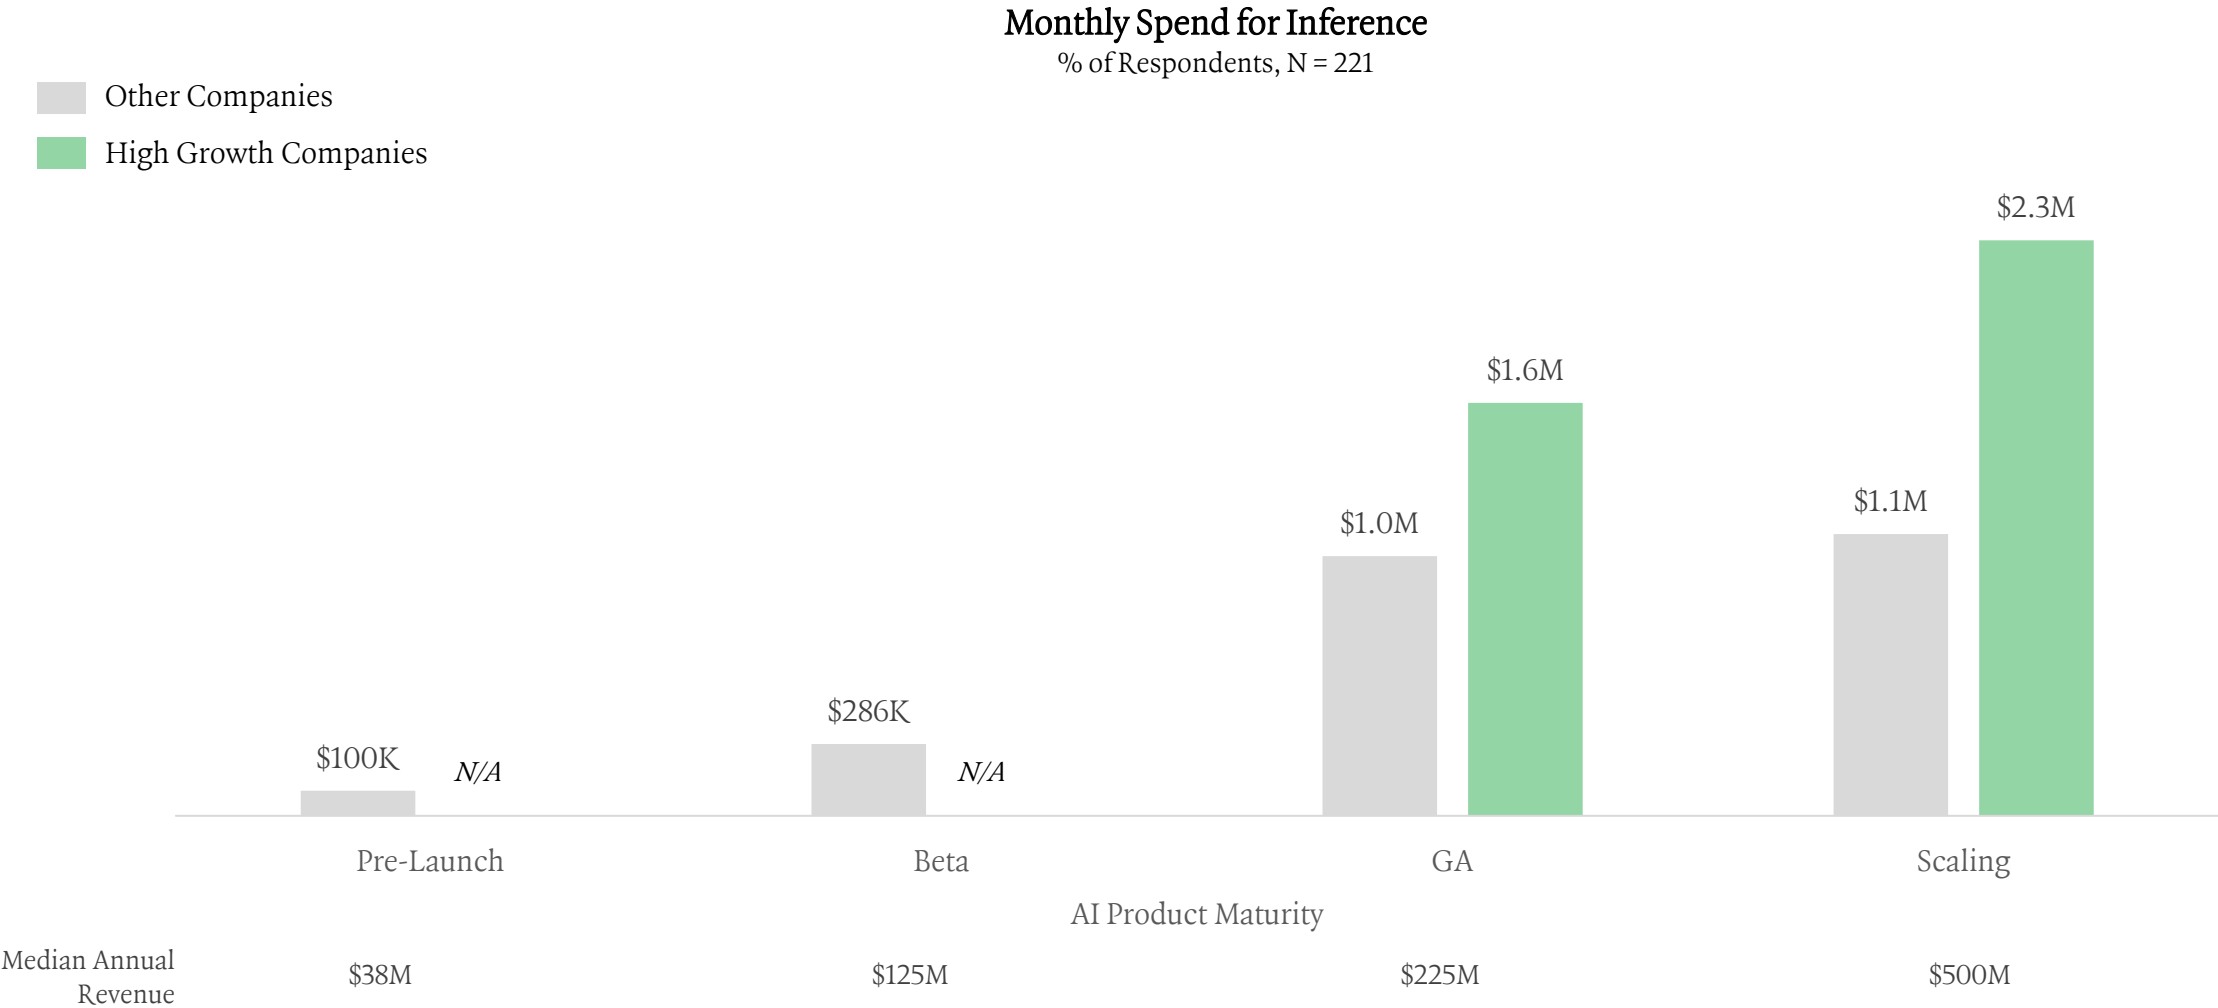
Model Training

Most respondents are training or finetuning models at least monthly, with estimated monthly model training costs ranging from \$160K-\$1.5M depending on the product maturity



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Inference costs surge post-launch with high-growth AI companies spending up to 2x more at GA and scale than their peers



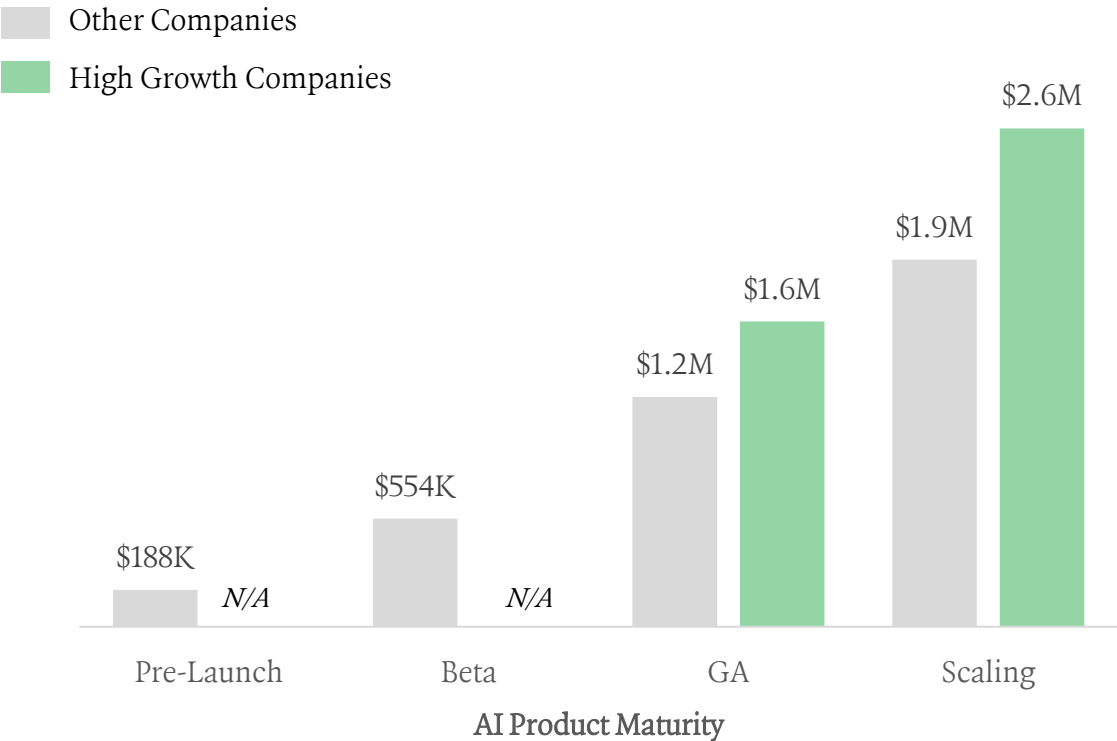
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Deployment Costs: Data Storage & Processing

Data storage & processing costs also climb steeply from GA stage onward, with high-growth AI builders spending more on data storage and processing than their peers

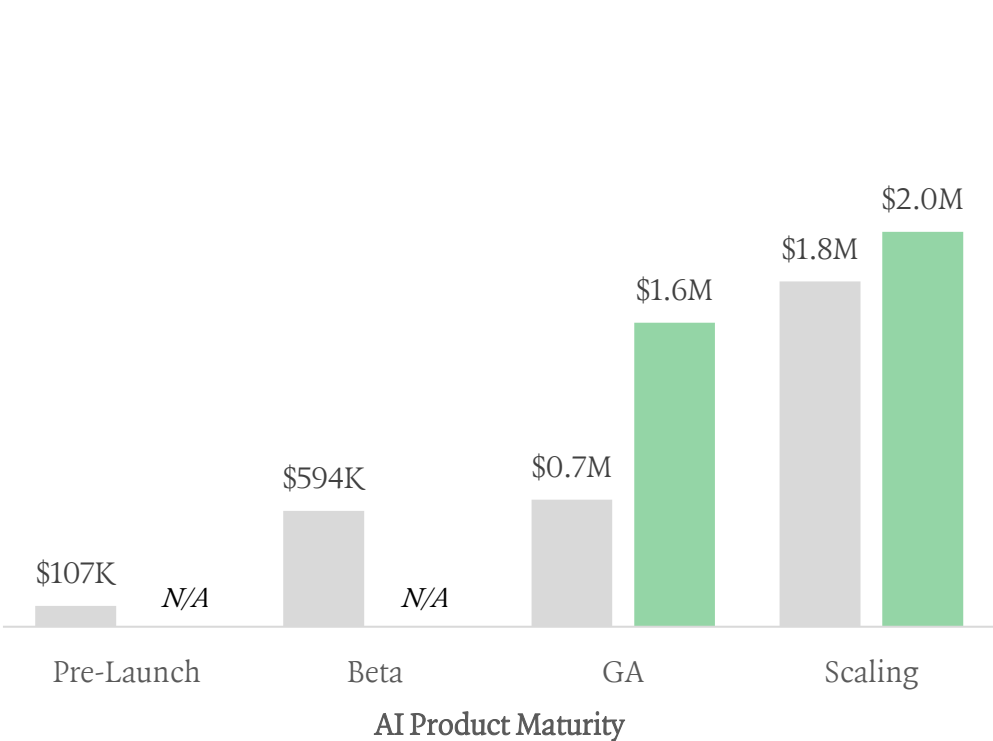
Monthly Spend for Data Storage

% of Respondents, N = 221



Monthly Spend for Data Processing

% of Respondents, N = 226



Median Annual Revenue \$38M \$125M \$225M \$500M

Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

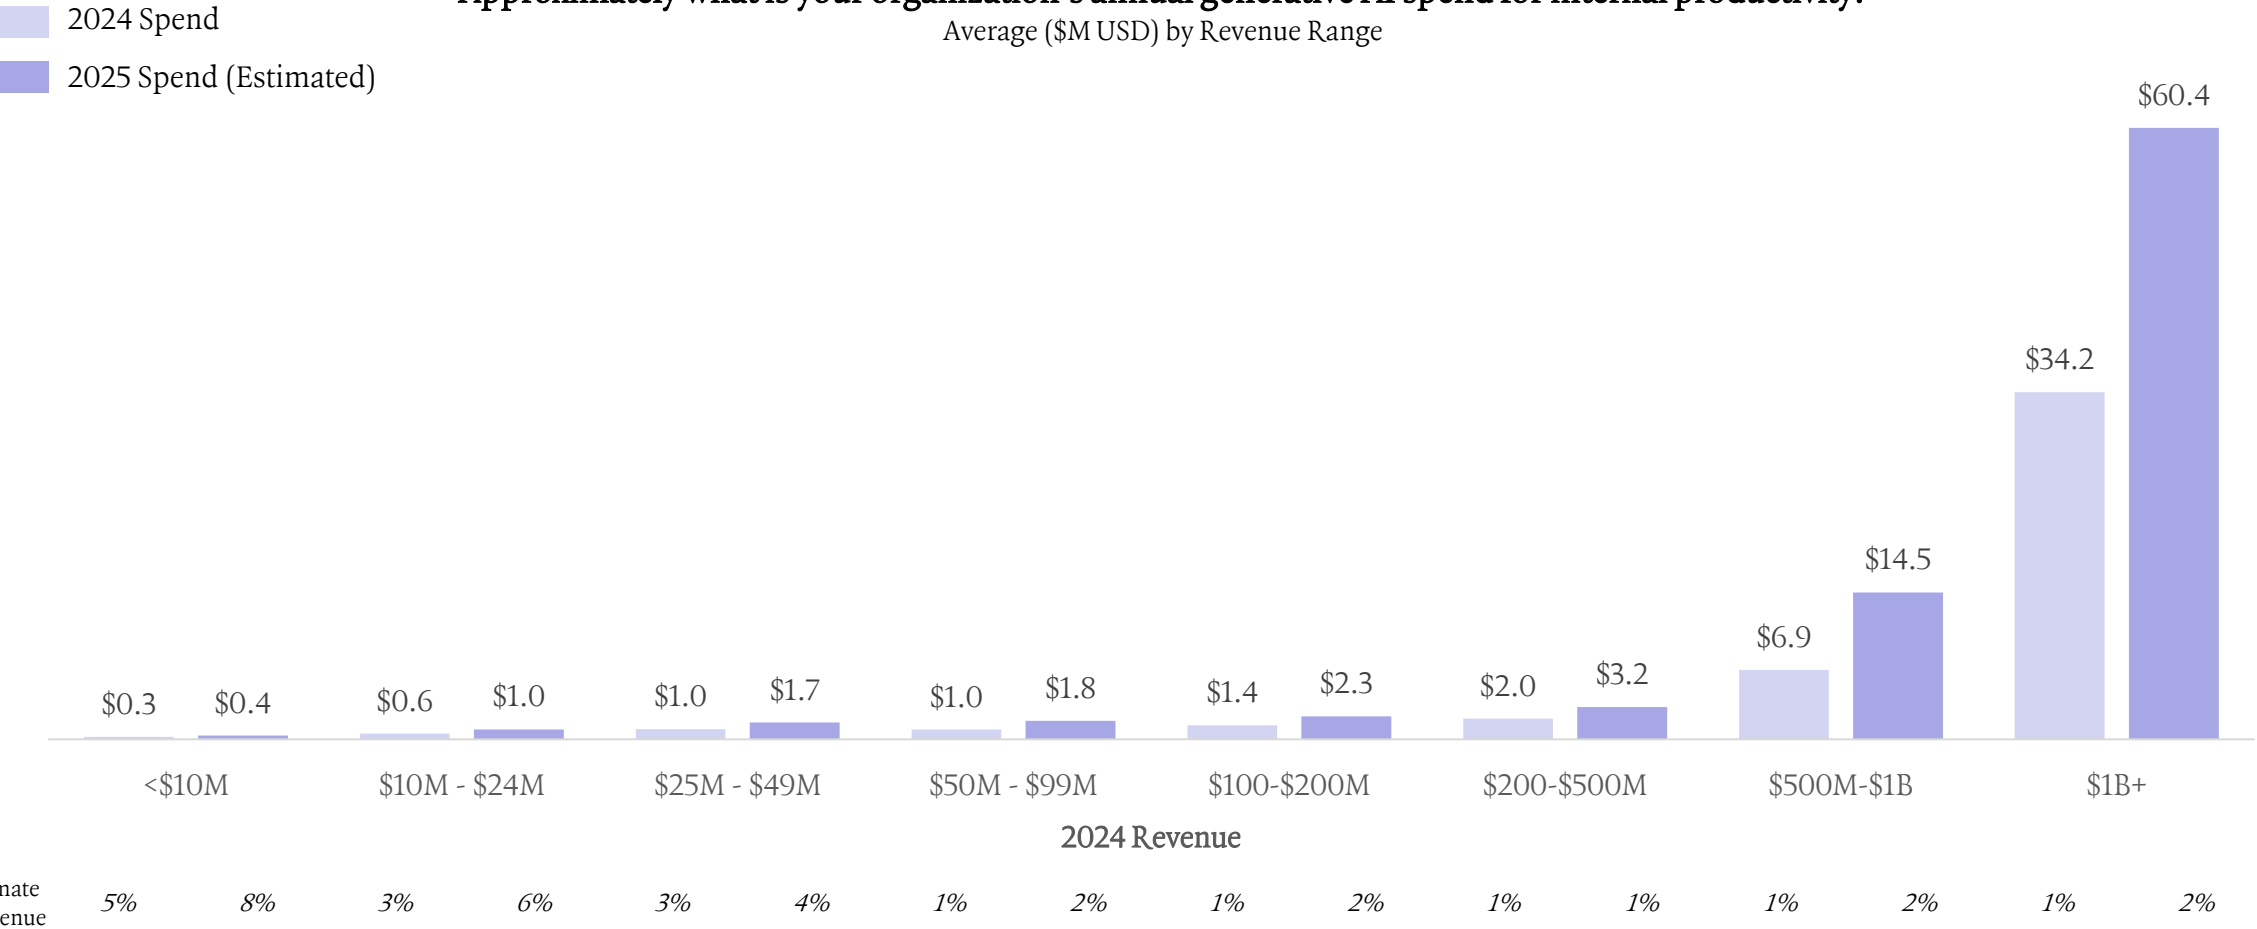
Internal Productivity



Annual Internal Productivity Budget

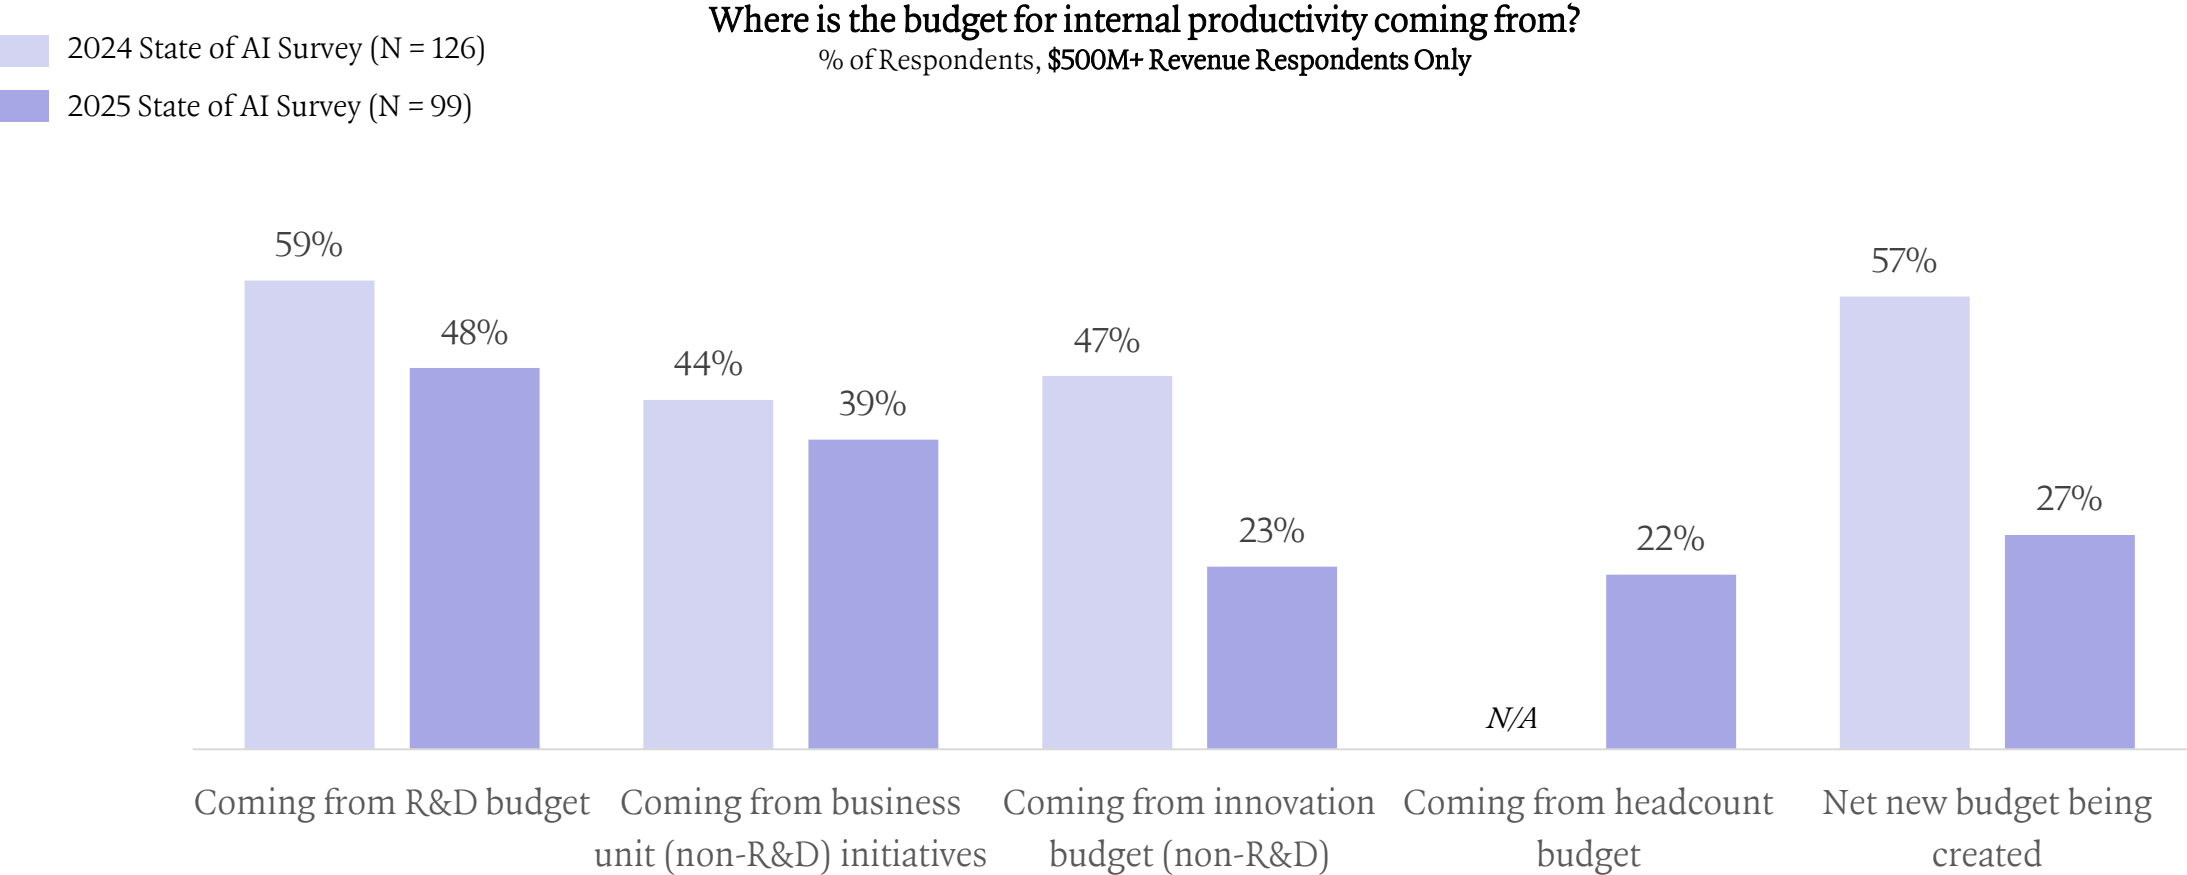
Internal AI productivity budgets are set to nearly double in 2025 across all revenue tiers, with companies spending anywhere from 1-8% of total revenue

Approximately what is your organization’s annual generative AI spend for internal productivity?
Average (\$M USD) by Revenue Range



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

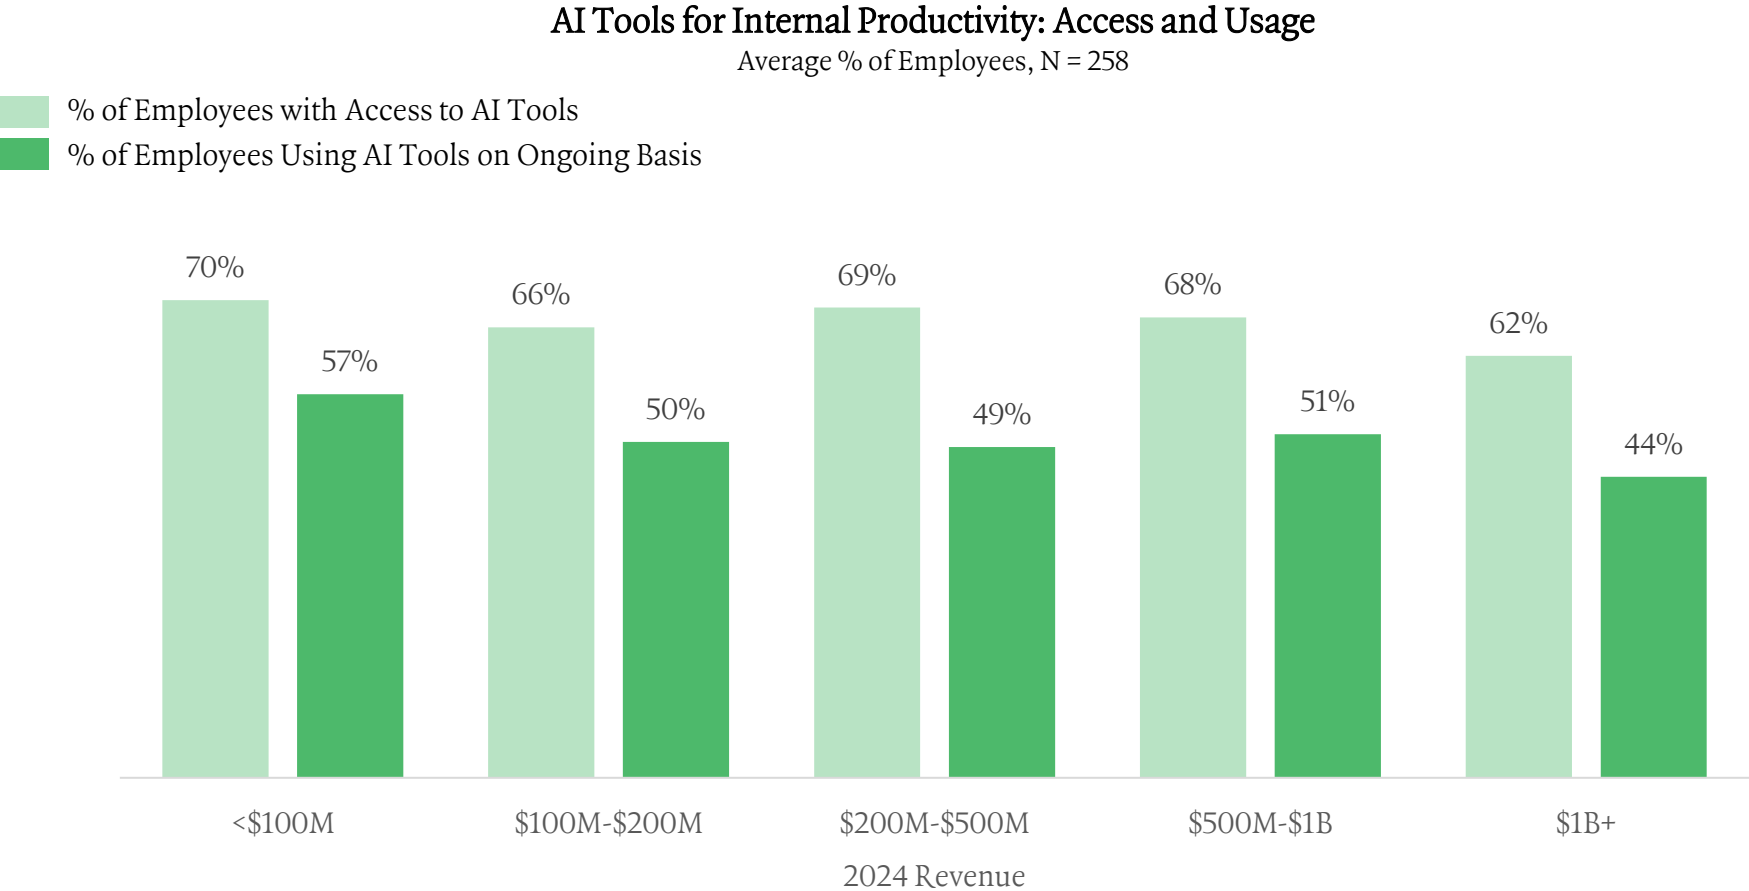
R&D budgets still remain the most common source of AI internal productivity budgets for enterprise companies; however, we are also starting to see headcount budgets being used for internal productivity spend



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

AI Access and Usage

While around 70% of employees have access to various AI tools for internal productivity, only ~50% of employees are using AI tools on an ongoing basis with adoption more difficult in mature Enterprises (\$1B+ revenue)



Don Vu
SVP, Chief Data & Analytics Officer,
New York Life

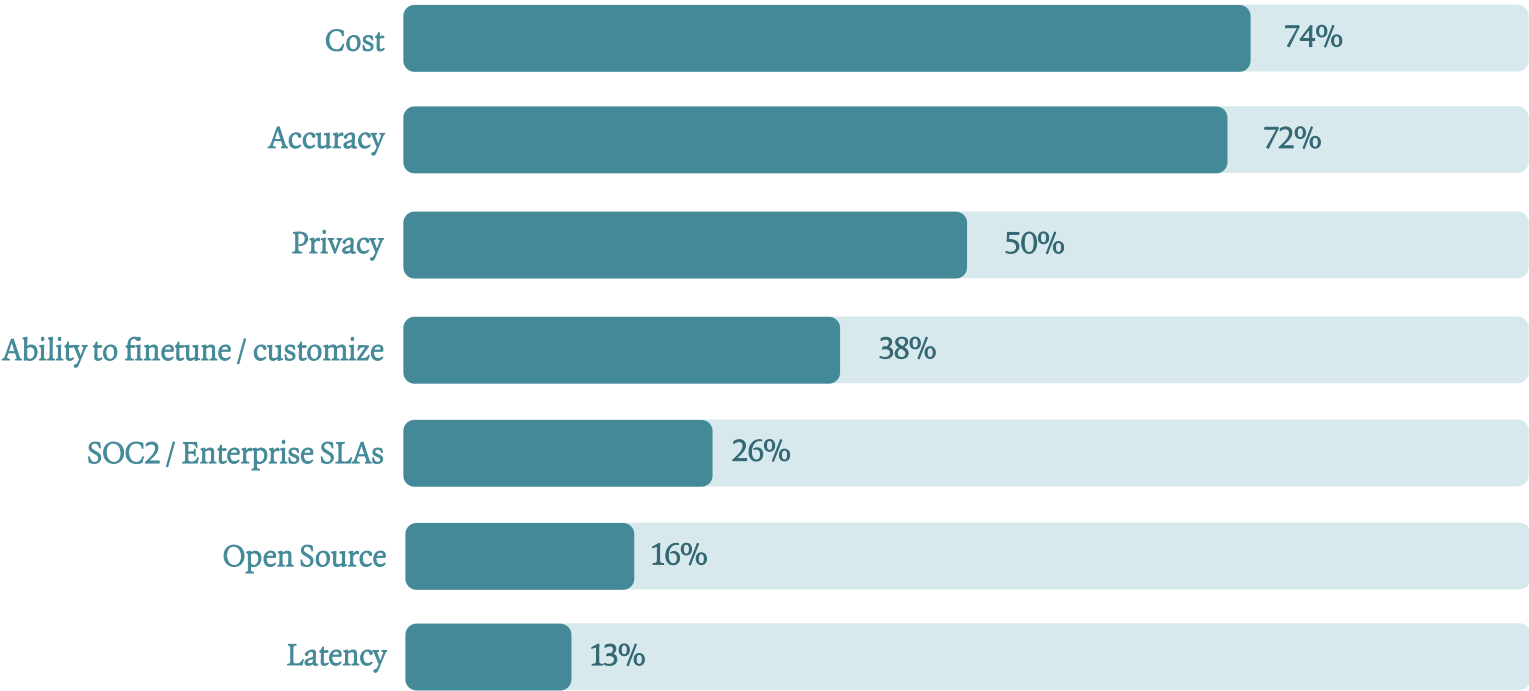
Just deploying tools is a recipe for disappointment, particularly for large enterprises. To truly empower employees, you need to **pair availability with scaffolding** that includes training, spotlighting champions, and most importantly relentless executive support.

”

Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

When choosing foundational models for internal use cases, cost is the most important consideration followed by accuracy and privacy

Top Considerations When Choosing a Foundational Model for Internal Use Cases
% of Respondents who ranked each aspect in Top 3, N = 265

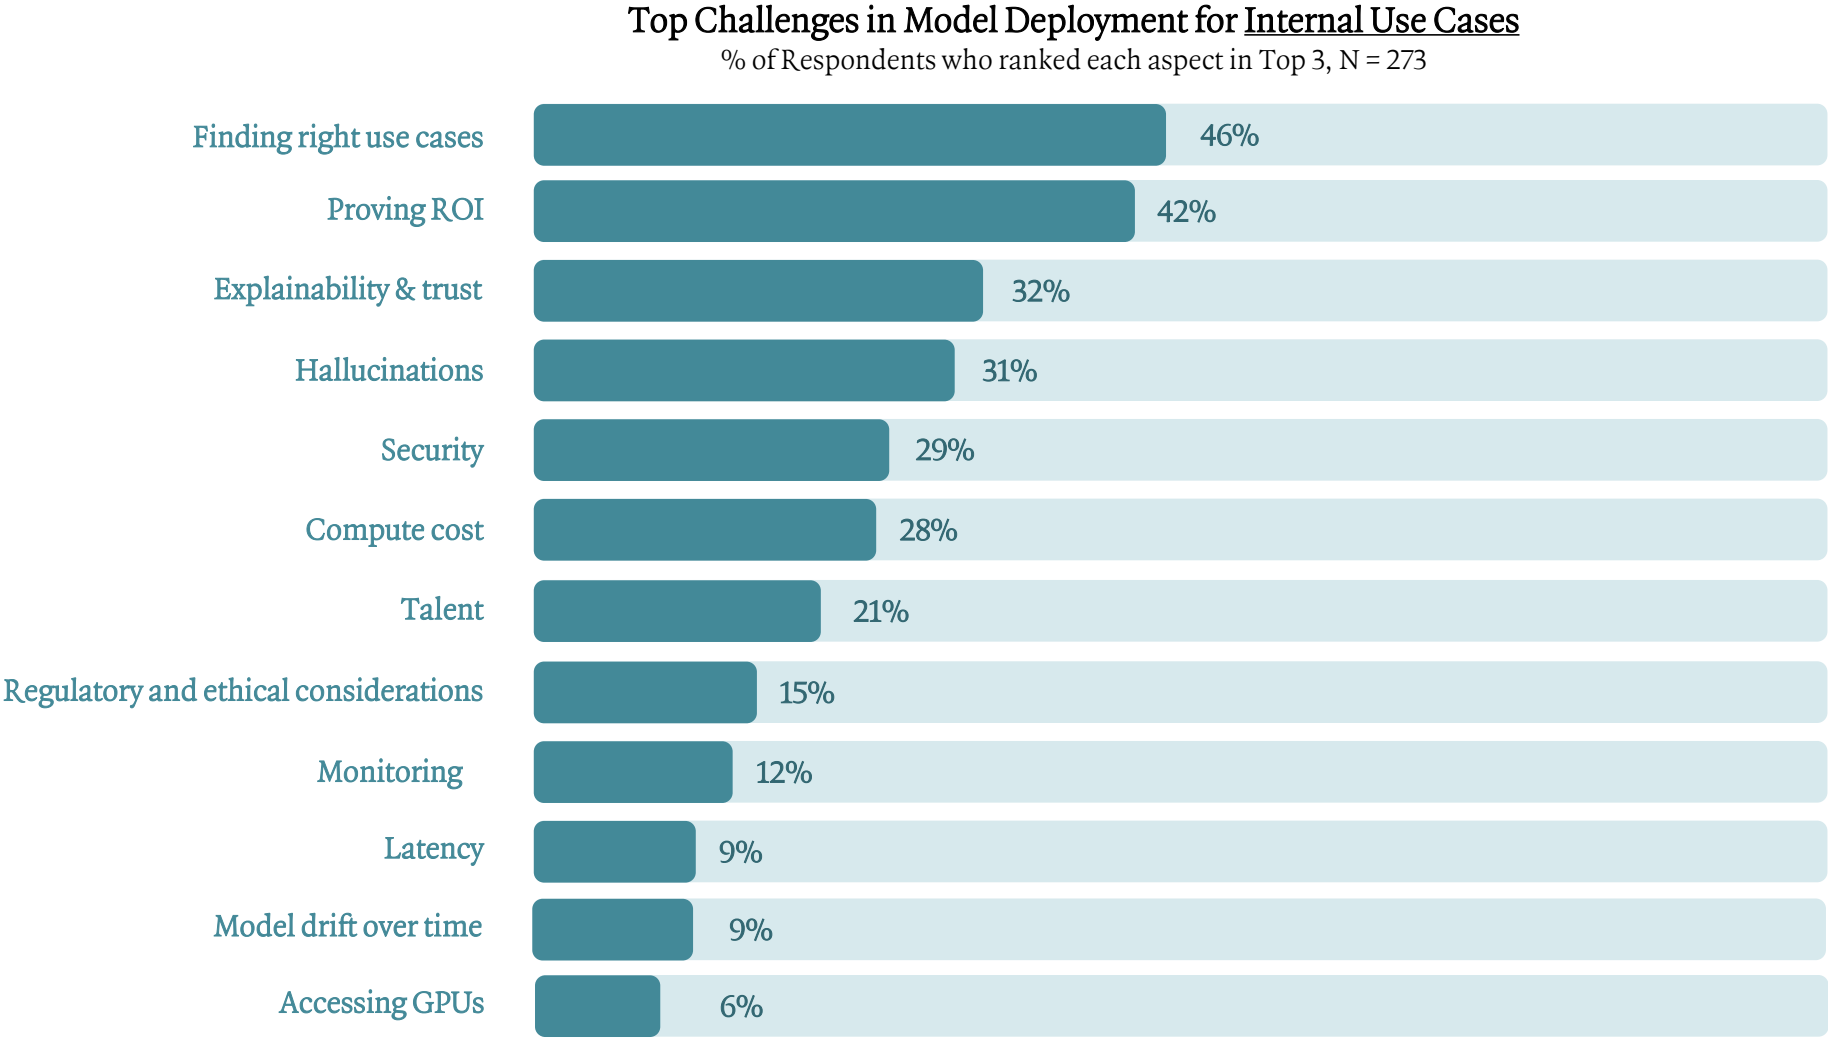


Whereas accuracy ranked as the most important factor when deploying external AI products, **cost is the most important consideration when choosing models for internal AI use cases.**

Privacy also becomes a more important consideration for internal use cases compared to external.

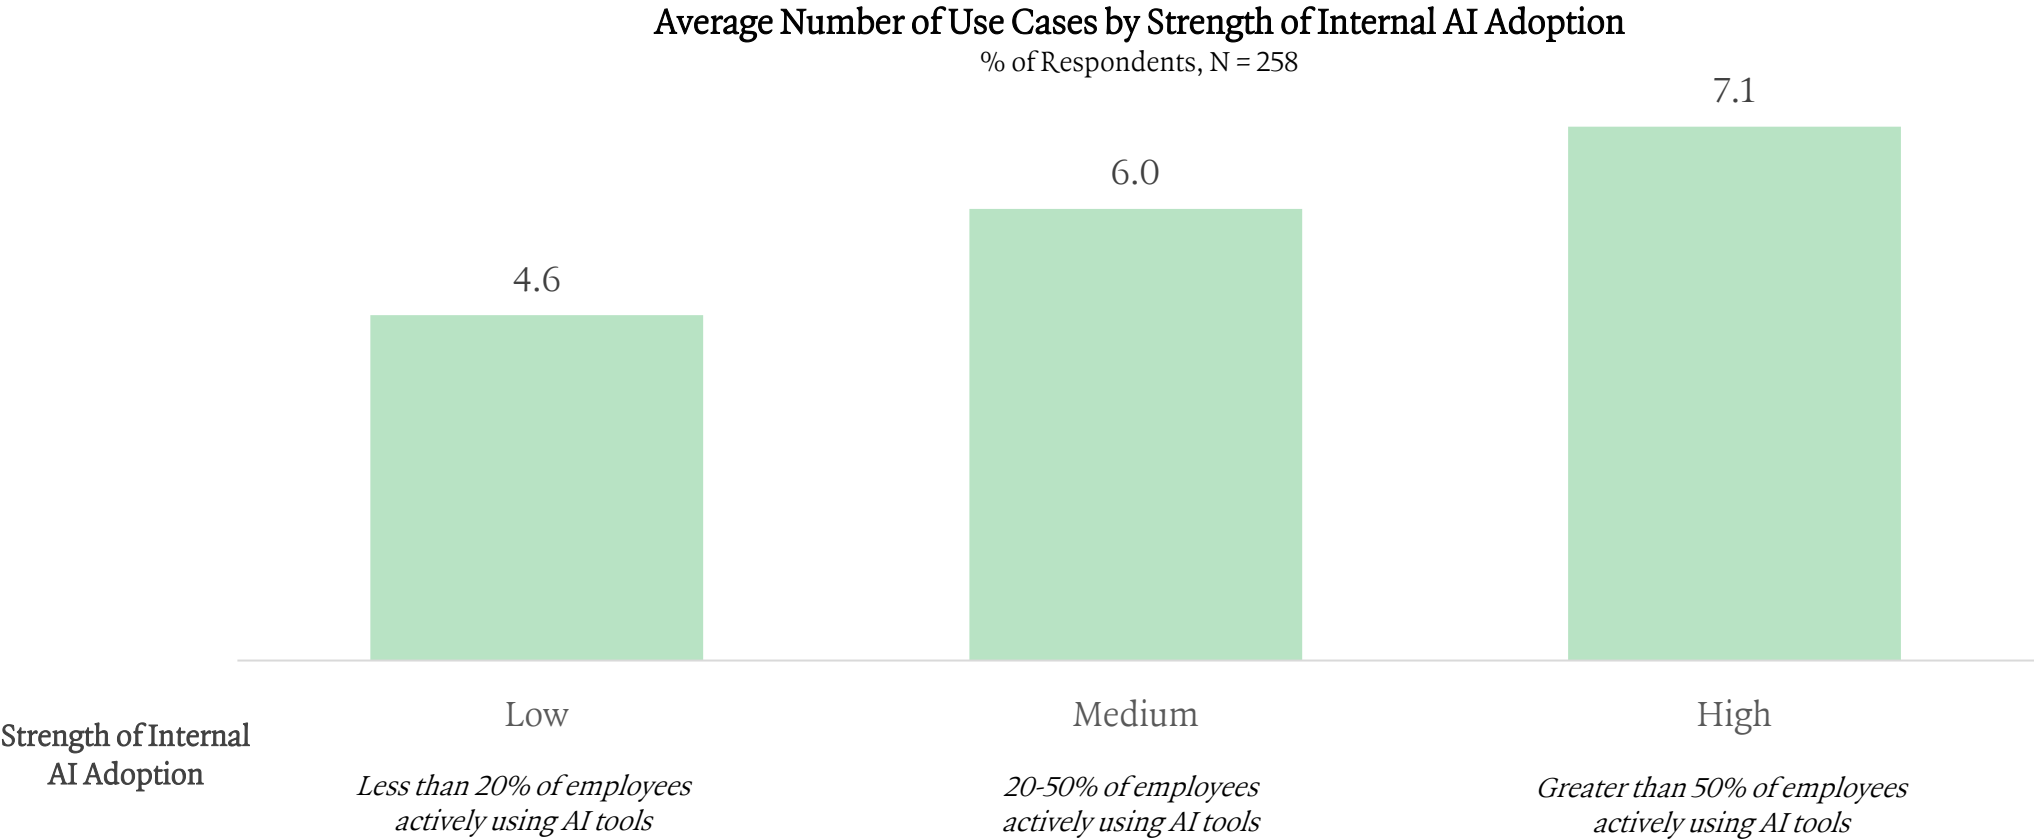
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/CDOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

The biggest challenges facing organizations deploying AI for internal use cases are often strategic (i.e. finding the right use cases and proving ROI) vs technical



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

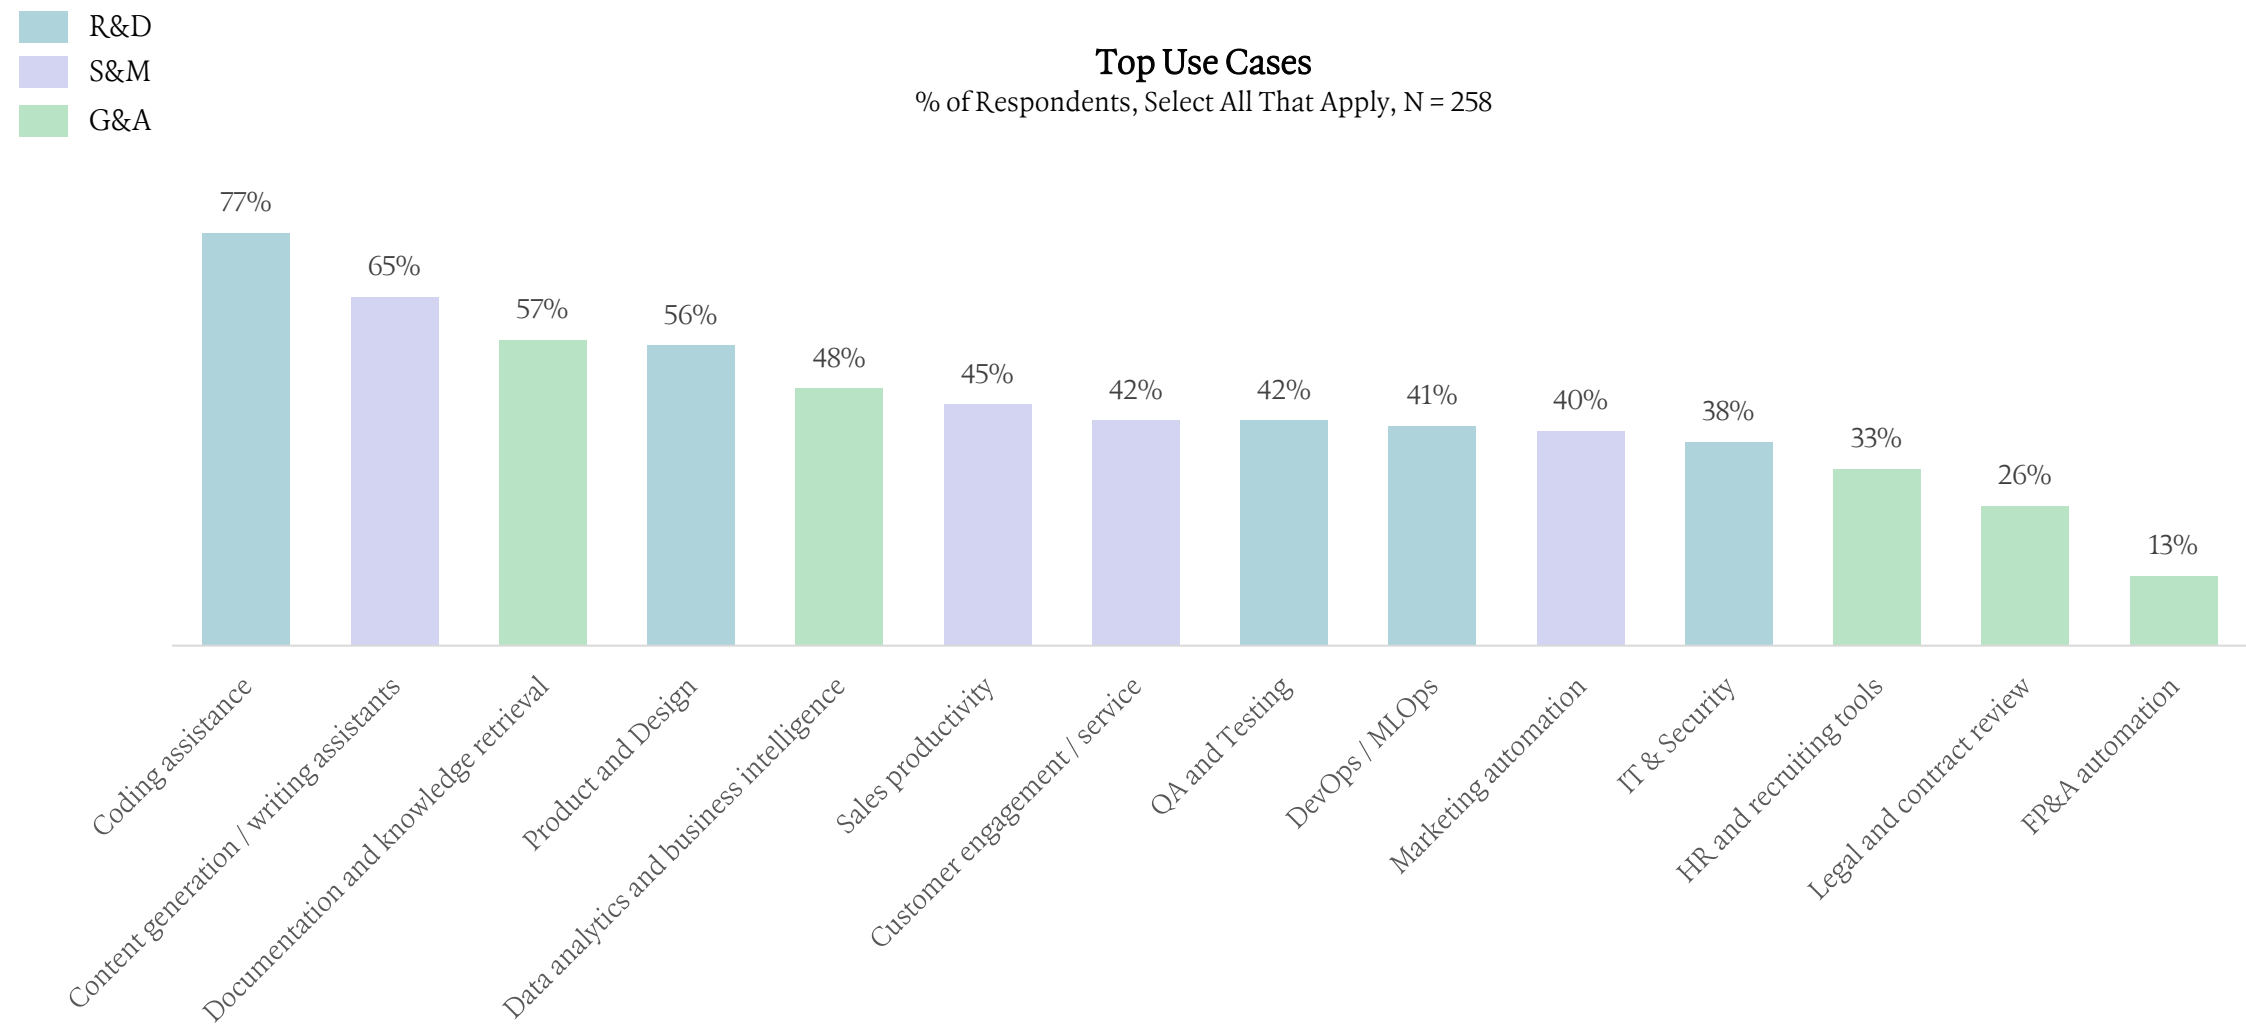
Companies are typically exploring multiple GenAI use cases across functions, with companies that have high employee adoption using GenAI across 7+ use cases



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Top Use Cases: By Popularity

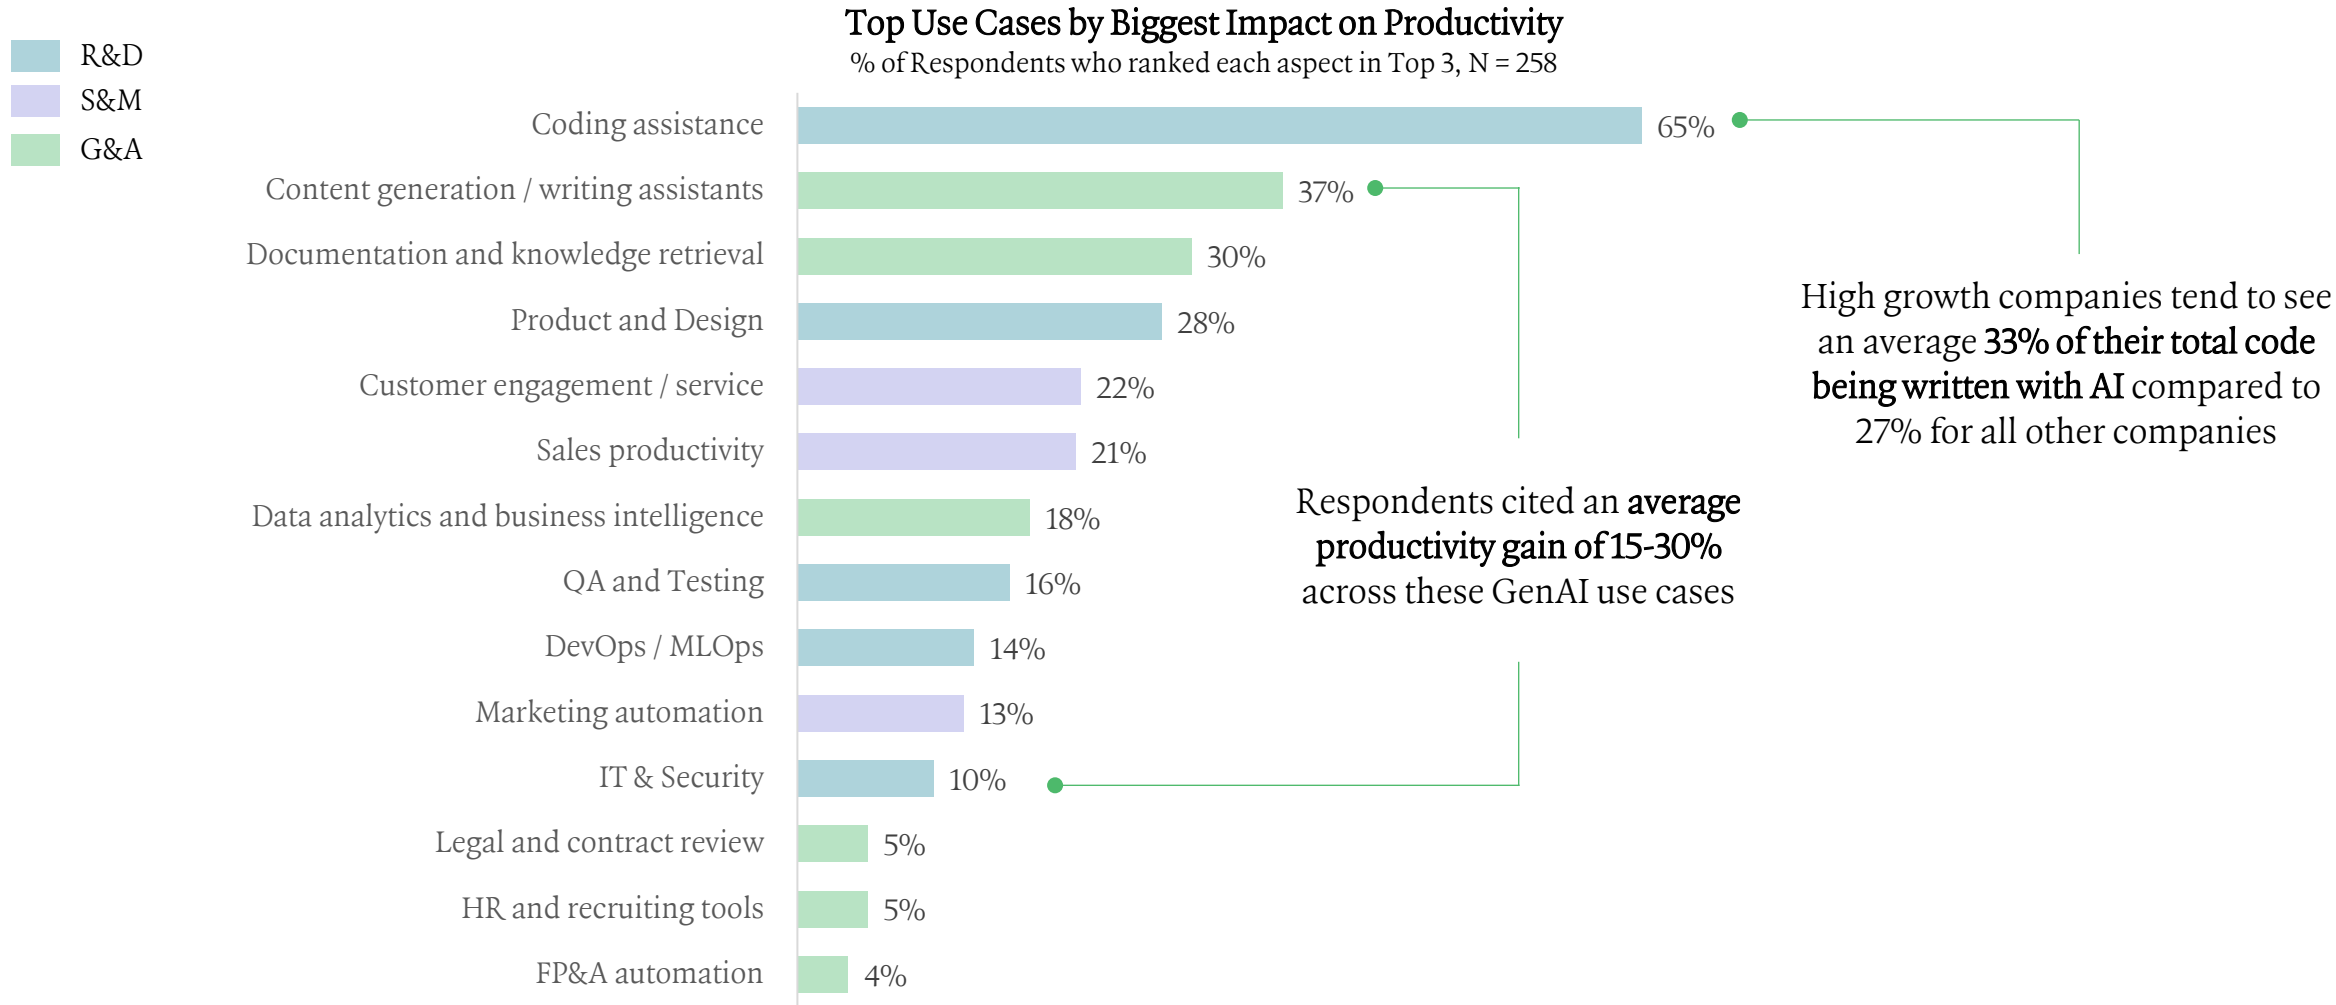
R&D and S&M use cases lead in popularity, while G&A use cases still lag in comparison



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Top Use Cases: By Impact

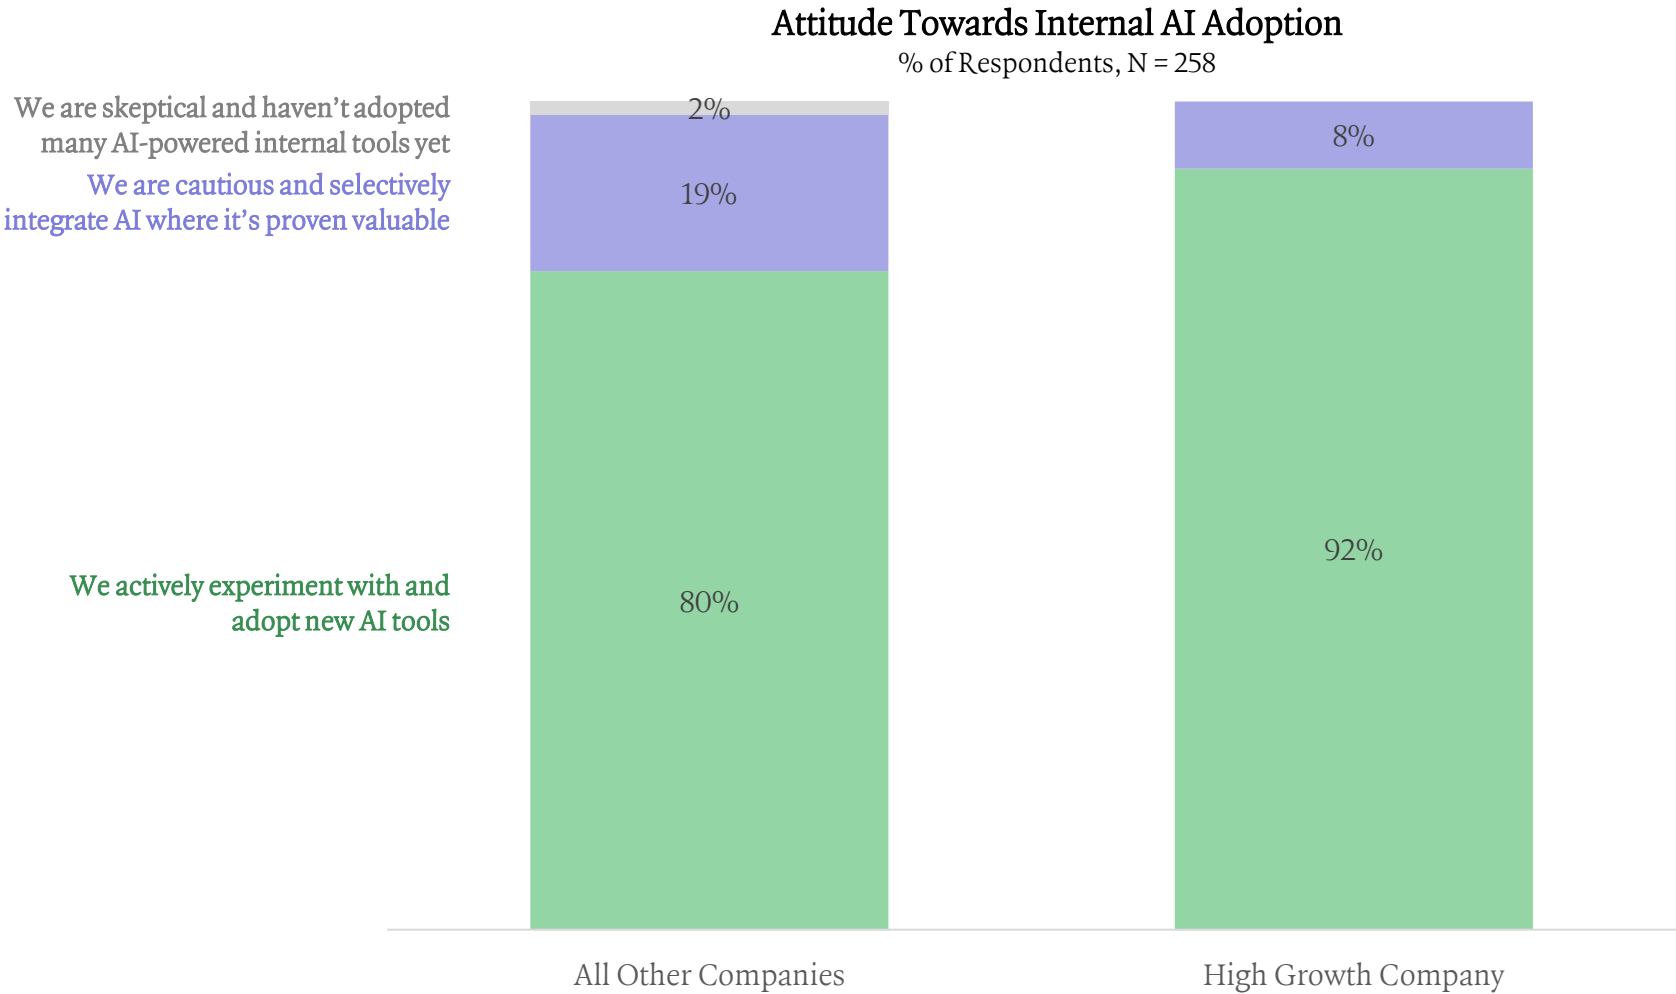
Top use cases by impact mirror usage trends with coding assistance by far outpacing other use cases in terms of tangible impact on productivity



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Attitude Towards Internal AI Adoption

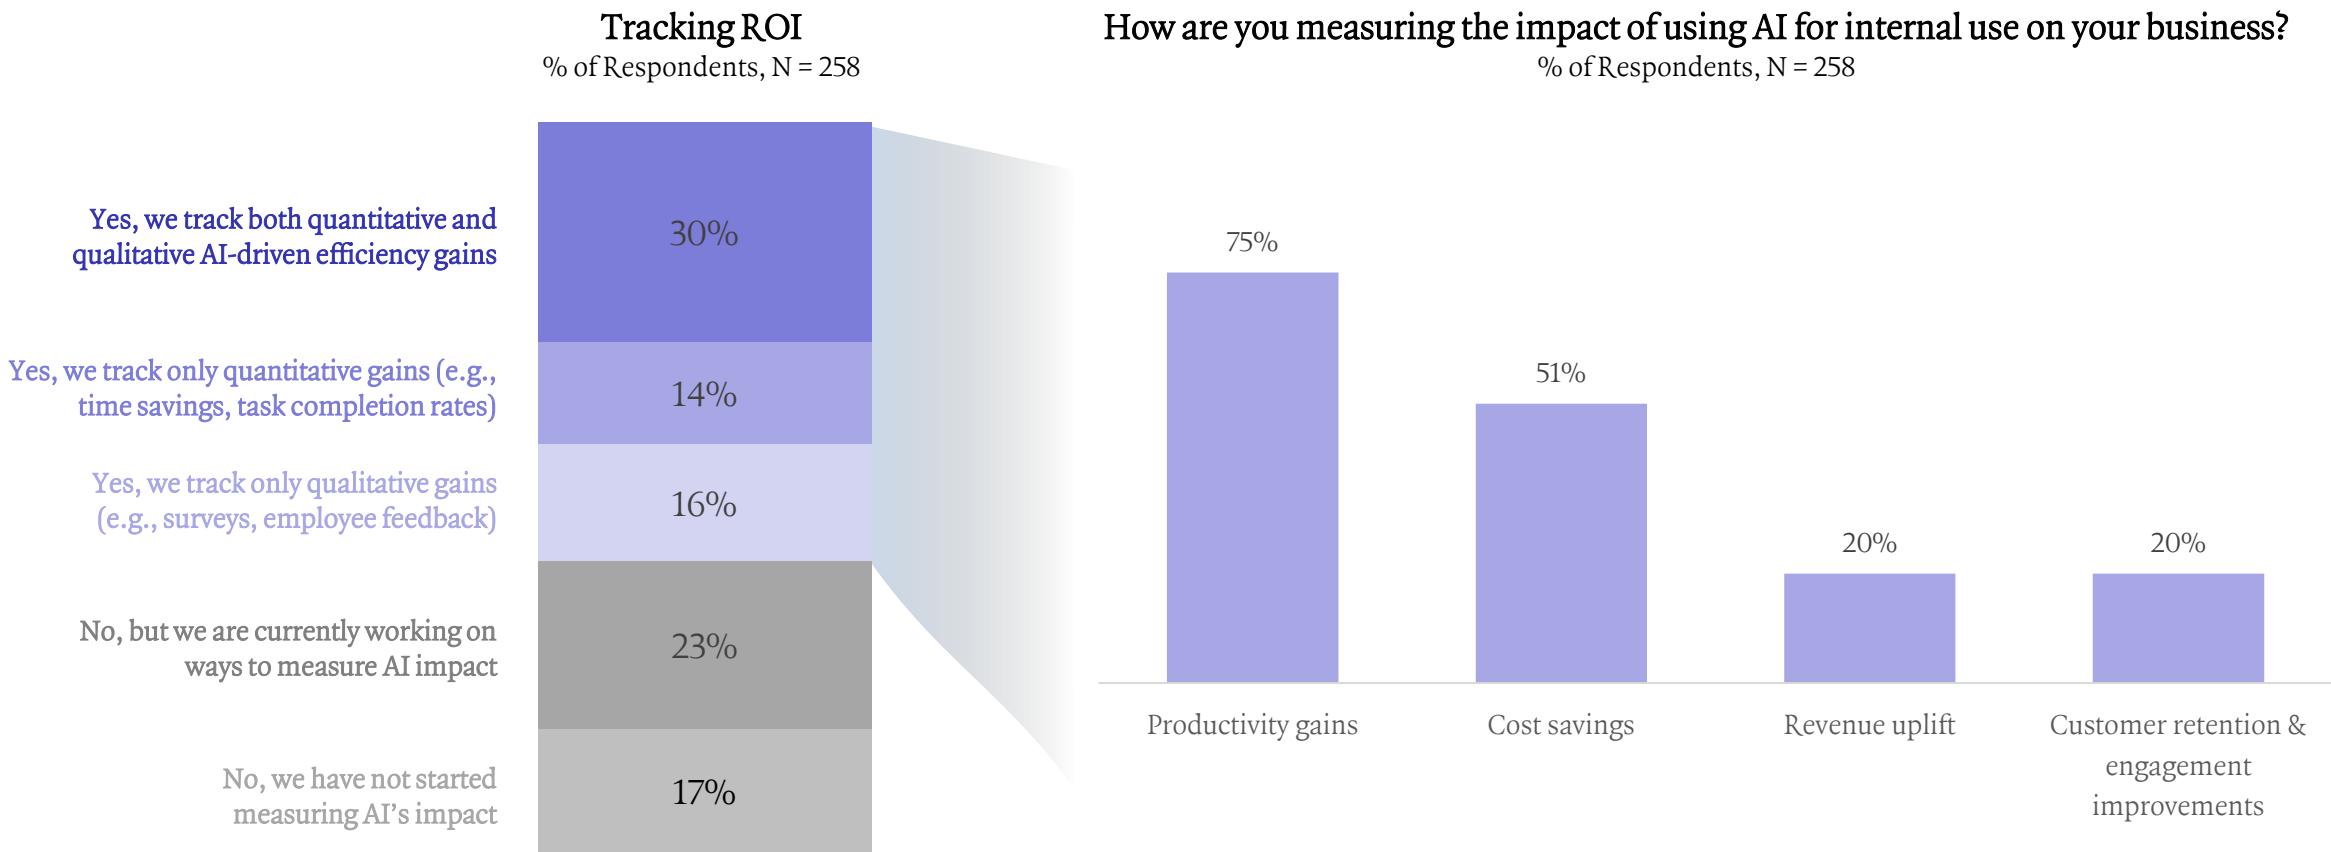
High growth companies tend to more actively experiment with and adopt new AI tools, suggesting that leading companies view AI as a strategic lever and are moving faster to integrate it across internal workflows



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Tracking ROI

Most companies are measuring productivity improvements and cost savings from internal AI use



Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

AI Builder Tech Stack



Most Used Tools: Model Training & Finetuning

Key Takeaways

Frameworks vs Managed Platforms

- **Core deep-learning frameworks remain popular** with PyTorch and TensorFlow accounting for over half of all usage across respondents
- But they're nearly **matched by fully managed or API-driven offerings** - prevalence of AWS SageMaker and OpenAI's fine-tuning service show that teams are **split between “build your own” and “let someone else run it” approaches**

Ecosystem Players Gaining Traction


- The Hugging Face ecosystem and Databricks' Mosaic AI Training are carving out meaningful niches, **providing higher-level abstractions over raw frameworks**
- Meanwhile, more specialized or emerging tools (AnyScale, Fast.ai, Modal, JAX, Lamini) landed in the single-digit percentages, suggesting **experimentation is underway but broad adoption remains nascent**

Enterprise-Grade Needs

- Later-stage companies typically have **larger data teams, more complex pipelines, and stricter requirements around security, governance, and compliance**
- Databricks' unified “lakehouse” architecture (which blends data engineering, analytics, and ML) and AnyScale's managed Ray clusters (which simplify distributed training and hyperparameter tuning) both speak directly to those enterprise needs with **more respondents in the \$500M+ revenue range using these solutions**

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 Used Tools: LLM & AI Application Development

Key Takeaways

Orchestration Frameworks Reign Supreme

- Top frameworks used include LangChain and Hugging Face's toolset which signals that teams clearly **value high-level libraries that simplify prompt chaining, batching, and interfacing with either public or self-hosted models**
- Around 70% of respondents also specified that they **use private or custom LLM APIs**

Safety and Higher-Level SDKs Gaining Traction


- Roughly 3 in 10 respondents use Guardrails to enforce safety checks, and almost a quarter leverage Vercel's AI SDK (23%) for rapid deployment which shows **growing awareness that production LLM apps need both guardrails and streamlined integration layers**

Long-Tail Experimentation

- Emerging players like CrewAI, Modal Labs, Instructor, DSPy, and DotTXT had weaker usage, indicating that **while experimentation is widespread, broad standardization has yet to settle beyond the big players**

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) and perspectives from the ICONIQ team and network of AI leaders consisting of our community of CIO/COOs overseeing AI initiatives in enterprises, CTOs, our Technical Advisory Board, and others in our network

Most Used Tools: Monitoring and Observability

Key Takeaways

Incumbent Infrastructure Still Rules

- Nearly half of teams lean on their **existing APM/logging stacks** (Datadog, Honeycomb, New Relic, etc.) **rather than adopting ML-specific tools** - underscoring that ease of integration and organizational standardization often outweigh the benefits of bespoke AI monitoring

Early Traction for ML-Native Platforms


- Both LangSmith and Weights & Biases have broken through to reach ~17% adoption, showing **real appetite for turnkey solutions that instrument prompt chains, track embeddings, and surface drift without bolt-ons to legacy systems**

Fragmented Long Tail & Knowledge Gaps

- Beyond the top two ML-native names, **usage quickly fragments across players** like Arize, Fiddler, Helicone, Arthur, etc, and 10% of respondents didn't know which tool they used; this points to both **a nascent ecosystem and confusion around what "observability" even means for generative AI**

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 Used Tools: Inference Optimization

Key Takeaways

NVIDIA's Grip on Production

- TensorRT and Triton Inference Server together command over 60% adoption, underscoring **how dominant NVIDIA's stack remains for squeezing latency and throughput out of GPU-based deployments**

Cross-Platform Alternatives Gaining Share

- The ONNX Runtime (18%) is the top non-NVIDIA solution, reflecting teams' **desire for hardware-agnostic acceleration across CPUs, GPUs, and accelerators**
- TorchServe (15%) likewise shows that pure-PyTorch serving still has a foothold, especially for CPU-only workloads or simpler containerized setups

Knowledge Gaps & Untapped Potential

- With 17% respondents they didn't know which optimization they use and 14% reporting "None," there's clear **confusion or inexperience around inference tuning**, suggesting an opportunity for education (and tooling) around quantization, pruning, and efficient runtimes - especially for teams running at scale

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 Used Tools: Model Hosting

Key Takeaways

Direct-from-Provider Is King

- The majority of teams hit model hosts directly via OpenAI, Anthropic, etc. underscoring that the path of least resistance remains **calling the vendor's own inference APIs rather than building or integrating through a middle layer**

Hyperscalers Close Behind


- AWS Bedrock and Google Vertex AI have carved out substantial share, reflecting **strong demand for unified, enterprise-grade ML platforms that bundle hosting with governance, security, and billing in a single pane**
- In particular, a **greater number of later-stage companies (\$500M+ revenue)** reported using hyperscaler solutions

Fragmented Alternatives & Emerging Players

- Beyond the big three, **usage quickly fragments across players like Fireworks, Modal, Together.ai, AnyScale, Baseten, Replicate, Deep Infra, etc.**
- This long tail suggests teams are **still exploring specialty hosts, often driven by unique pricing, performance SLAs, or feature sets** (e.g., custom runtimes, on-prem options)

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 Used Tools: Model Evaluation

Key Takeaways

No Clear Stand-alone Leader

- Nearly 1 in 4 teams use mostly built-in evaluation features from platforms like Vertex, Weights & Biases, or Galileo while 20% of respondents simply “didn’t know” which tool they use, signaling many organizations are still leaning on the evaluation capabilities baked into their existing ML stacks rather than adopting a dedicated framework

Emerging Specialized Frameworks


- LangSmith and Langfuse lead the pack of purpose-built evaluation tools, with HumanLoop and Braintrust also showing traction; these platforms are winning mindshare by offering richer prompt-level metrics, customizable test suites, and drift detection out of the box

Knowledge Gaps and DIY

- Almost a quarter of respondents did not know which evaluation tool they used or did not have an evaluation tool in place, signaling **both confusion around what “evaluation” entails for generative AI and the risk of unmonitored model regressions**
- Meanwhile, some respondents are also **rolling their own evaluation pipelines**, suggesting off-the-shelf tooling hasn’t yet covered all use cases

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 Used Tools: Data Processing & Feature Engineering

Key Takeaways

Classic Big Data Tools Still Dominate

- Apache Spark (44% of respondents) and Kafka (42% of respondents) lead the pack, underscoring that **at scale, teams default to battle-tested, distributed batch-and-stream frameworks for ETL and real-time data ingestion**

Python Power Base

- Despite heavy big-data footprints, 41% of respondents still lean on Panda - showing that **for smaller datasets, prototyping, or edge cases, the simplicity and flexibility of in-memory Python tooling remain indispensable**

Feature Stores on the Horizon

- Only 17% are using a dedicated feature store, indicating that **while the concept of “build once, serve everywhere” for features is gaining visibility, most organizations haven’t yet operationalized it at scale**
- **As maturity grows, we’ll likely see feature stores and lightweight orchestrators (Dask, Airflow, etc.) climb the ranks - but for now the Apache ecosystem rules**

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 Used Tools: Vector Databases

Key Takeaways

Search Engines Evolve Into Vector Stores

- Elastic and Pinecone lead adoption, reflecting how teams either retrofit existing full-text search platforms for embeddings or adopt purpose-built, managed vector engines

Redis & the “Long Tail”

- Redis shows the appeal of leveraging in-memory data stores you already run, while other solutions like Clickhouse, AlloyDB, Milvus, PGVector, etc, underscores that many organizations are experimenting with different backends to balance cost, latency, and feature needs

Rise of Open-Source Solutions

- Specialist open-source tools like Chroma, Weaviate, Faiss, Qdrant, and Supabase’s vector addon are chipping away at the early leaders, signaling a competitive battleground for ease-of-use, scaling, and cloud-native integrations

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 Used Tools: Synthetic Data & Data Augmentation

Key Takeaways

In-House Reigns Supreme

- Over **half of teams (52%) build their own tooling**, suggesting that off-the-shelf providers still struggle to cover every use case or integrate with existing pipelines

Scale AI is the clear vendor leader

- At 21% adoption, **Scale AI is the go-to third-party synthetic-data platform** - but even it only reaches one in five organizations

Early Traction for Programmatic Frameworks

- Snorkel AI and Mostly AI show that **programmatic labeling and generation tools are gaining mindshare, but remain far behind custom solutions**

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 Used Tools: Coding Assistance

Key Takeaways

Dominance of First Movers


- **GitHub Copilot is used by nearly three-quarters of development teams**, thanks to its tight VS Code integration, multi-language support, and backing by GitHub’s massive user base
- Copilot’s network effects and product-market fit make it hard to dislodge, but the **strong second-place showing for Cursor (used by 50% of respondents)** signals appetite for diverse IDE integrations

Long Tail of Offerings Lag

- After the top two, adoption drops off sharply with a **fractured long tail of solutions**, suggesting that **while most teams have trialed at least one assistant, very few have standardized on alternatives**
- Low-code or no-code solutions like Retool, Lovable, Bolt, and Replit also had honorable mentions indicating that there is **increasing appetite in the market for idea-to-application solutions**

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 Used Tools: DevOps and MLOps

Key Takeaways

MLflow Leads—but No Monopoly


- MLflow was used by 36% of respondents and the **clear frontrunner for experiment tracking, model registry, and basic pipeline orchestration** – this is only just over one-third of teams, indicating plenty of room for alternatives
- Weights & Biases also holds strong share with 20% of respondents using, reflecting **its appeal as a managed SaaS for tracking, visualization, and collaboration**
- **Beyond the top two, usage quickly fragments** – 16% “don’t know” which tools power their MLOps and other tool mentions include Resolve.ai, Cleric, PlayerZero, Braintrust, etc. This points to both confusion around responsibilities (DevOps vs. MLOps) and a market still sorting itself out

Gap between Tracking and Full-Scale Ops

- The dominance of tracking-first platforms like MLflow and W&B suggests that **many teams haven’t yet adopted end-to-end MLOps suites - continuous delivery, drift monitoring, or automated rollback remain work in progress for most**

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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 Used Tools: Product and Design

Key Takeaways

Figma's Near-Universal Reach

- With 87% adoption, **Figma is effectively the de-facto standard for UI/UX and product design** - teams overwhelmingly stick with its real-time collaboration, component libraries, and plugin ecosystem rather than seeking out AI-specific design tools

Miro for Higher-Level Collaboration

- With 37% adoption, **Miro remains the go-to for wireframing, user-journey mapping, and cross-functional brainstorming**; its whiteboard-style interface complements Figma's pixel-perfect canvases, especially in early ideation phases

Rise of AI-Enabled Product Wireframes

- Design teams aren't yet feeling the urgent need for AI-native product/design platforms, however **many are using low/no-code solutions to Bolt, Lovable, and Vercel V0 for rapid prototyping**

Most Widely Used Tools

From survey respondents; By alphabetical order



Notes: Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.
Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Internal Productivity Use Cases (Part 1 of 2)

Use Case

Key Trends

Sales Productivity

- Many teams are getting their AI-powered sales features straight out of Salesforce - indicating that an easy path is to **lean on your existing CRM's built-in recommendations, forecasting, and opportunity-scoring rather than bolt on a separate service**
- Other respondents are also using **sales-engagement platforms** like Apollo, Salesloft, Gong, etc, while others are also leaning into **AI driven prospecting tools** like Clay and People.ai
- As embedded capabilities mature, we will likely see **consolidation around a handful of platforms or clearer differentiation from the point-solution upstarts**

Marketing Automation & Content Generation

- Marketers overwhelmingly turn to **Canva's generative features for on-brand visuals and quick content iterations**, making it by far the most common "AI" touchpoint in the marketing stack
- Many respondents are also using solutions like n8n or homegrown solutions, indicating that **marketing use cases sometimes require a high degree of in-house customization**
- Many respondents are also using **specialized AI writing tools** like Writer and Jasper, with adoption higher for later stage companies (\$100M+ revenue)

Customer Engagement

- Teams overwhelmingly rely on Zendesk or Salesforce's embedded AI features for customer interactions, signaling that **ease of plugging into existing ticketing and CRM workflows still beats adopting a standalone conversational AI platform**
- A sizable minority lean on **specialist tools** like Pylon, Forethought, Grano.la, or Intercom **when they need deeper bot customizations, self-service wizards, or tight in-app support widgets** - suggesting that best-of-breed still has a role when out-of-the-box AI falls short

Documentation and Knowledge Retrieval

- Most teams either build on existing wikis and note-taking tools or standardize on Notion; this shows that organizations often **default to whatever's already in place for knowledge capture before experimenting with AI-powered overlays**
- However, a sizable proportion of respondents are also leaning into **purpose-built AI tools** like Glean and Writer for **indexing and semantic search**

For more information on specific tools in each category, please reach out to [ICONIQ Insights](#)

Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

Internal Productivity Use Cases (Part 2 of 2)

	Use Case	Key Trends
For more information on specific tools in each category, please reach out to ICONIQ Insights	IT & Security	<ul style="list-style-type: none">• ServiceNow (used by 33% of respondents) and Snyk (used by 30% of respondents) lead the pack, showing that large organizations are still defaulting to their existing ITSM and security-scanning platforms rather than standing up brand-new AI tools• Zapier and Workato were also commonly mentioned, underlining how much teams value low-code orchestration for stitching together alerts, ticket creation, and remediation scripts across disparate tools
	Legal	<ul style="list-style-type: none">• Legal departments are dipping toes into AI primarily through ChatGPT and ad hoc scripts, but purpose-built legal assistant platforms are starting to gain traction• As regulation and security concerns mount, we'll likely see a bifurcation: mainstream LLMs for informal research and compliance-focused suites for mission-critical contract workflows
	HR & Recruiting	<ul style="list-style-type: none">• Nearly half of teams rely on LinkedIn's built-in AI features - profile suggestions, candidate matching, and outreach sequencing - underscoring that recruiters lean on platforms they already use daily rather than integrating standalone solutions• However, niche platforms like HireVue for AI-driven video interviews and Mercor for candidate engagement are starting to see modest uptake
	FP&A Automation	<ul style="list-style-type: none">• Many teams are using Ramp for FP&A automation, likely leveraging its spend management and data sync features in an all-in-one platform• Specialized suites like Pigment, Basis, and Tabs are also starting to pick up traction, showing growing interest in driver-based planning and multi-scenario modeling platforms• Around one-third of respondents are also using homegrown solutions, reflecting investment in custom scripts, Excel macros, and bespoke pipelines to glue together ERP, billing, and BI systems

Source: Perspectives from the ICONIQ GenAI Survey (April 2025) 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

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 Venture and Growth to disclose publicly). Trademarks are the property of their respective owners. None of the companies illustrated have endorsed or recommended the services of ICONIQ.

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 © 2025 ICONIQ Capital, LLC. All rights reserved.

Technology matters. Strategy matters. People matter most.

Meet the ICONIQ Venture and Growth team



ICONIQ

San Francisco | Palo Alto | New York | London

Join our community   