



In partnership with  Microsoft

AI PARTNERSHIPS ON THE PRODUCT ROADMAP



TABLE OF CONTENTS

Introduction	01
95% of Companies Have AI Product Initiatives	02
AI Innovation Leads, Revenue Still Matters	03
Generative AI is the Most Popular AI Investment	03
Hybrid Wins: In-House Build With Embedded AI Tools	04
Security: The Primary Bottleneck in AI Development	05
Hyperscaler Marketplaces Give You an Edge	06
Key Takeaways	07
Conclusion	08



INTRODUCTION

2022 was the year the ChatGPT chatbot interface launched to rocketship adoption, but 2025 was the year investments in Artificial Intelligence exploded. 40% of Enterprise companies are investing in AI capabilities to increase revenue growth and drive operational efficiency across their organizations. 2% of US GDP in 2025 was driven by AI infrastructure and investments.¹

With the global AI economy forecast to reach \$4.8TN by 2033, we wanted to understand where technology companies were investing to meet this massive customer demand.

We created this report with input from Microsoft's partners to help leaders navigate the decisions that matter for the year ahead:

- Where should companies invest for long-term growth?
- Which AI capabilities and product features should be a priority?
- How do companies package and sell those features so they are commercially viable?


What insights are included in this report?

- An in-depth analysis of AI trends across the ISV ecosystem
- Top hyperscaler benefits you can leverage to drive better outcomes for your AI product

To create this report, we surveyed industry leaders across the ISV ecosystem, 51% of whom are current Microsoft Partners, and conducted interviews with Microsoft Partners to capture context that does not show up in the numbers.

¹ AI market projected to hit \$4.8 trillion by 2033, emerging as dominant frontier technology, UNCTAD, April 7, 2025.

Below is a breakdown of our survey respondents' profiles per team size:



Use this report as a strategic reference to identify key industry trends, build a robust AI product roadmap, and ensure responsible deployment of AI.

95% OF COMPANIES HAVE AI PRODUCT INITIATIVES


In our survey, 95% of respondents report that AI product work is on the roadmap in the next 6-12 months. We see this as a response to the broader industry trend of heavy investment in AI.

Product teams are prioritizing AI features to keep up with customer expectations and competitor releases. As more companies build AI, the real challenge is creating value that customers can clearly see and are willing to pay for. We expect that pressure will shape most roadmaps, and we address what “value” needs to look like in the next section.

AI INNOVATION LEADS, REVENUE STILL MATTERS

Product innovation is currently the top motivator for companies building AI, with 62% of respondents citing it as their primary objective. This is unsurprising, as customers increasingly expect AI capabilities to be a natural part of the product experience. For many teams, the most effective strategy is to integrate AI into an existing product or portfolio, where it can improve functionality, quality, or usability without forcing customers to adopt a separate tool.

What is the key business objective driving your AI initiatives?




GENERATIVE AI IS THE MOST POPULAR AI INVESTMENT

Generative AI is the dominant category of functionality, with 58% of respondents naming it their primary AI capability focus. Customer adoption is highest here, and the use cases are being proven quickly in real workflows, especially prompt-based experiences that generate drafts, summaries, and answers inside existing tools.

What's interesting is what's not leading the list. Despite the significant attention it receives in industry conversations, agentic AI does not show up as a primary investment focus for most respondents, relative to generative AI. With frontier models investing heavily in agentic functionality to keep the AI working vs requiring constant user prompts and oversight, we expect Agentic AI investments from ISVs to dramatically increase in 2026.

Which emerging AI capabilities (e.g., generative AI, advanced ML) is your organization primarily planning to invest in over the next 6–12 months?



HYBRID WINS: IN-HOUSE BUILD WITH EMBEDDED AI TOOLS


Building an AI product rarely means building every layer yourself. In our survey, 57% of partners report using a hybrid approach that combines in-house development with embedded third-party tools.

Model training costs are eye-watering, and even if they were not, the low availability of chips and infrastructure to develop models means a hybrid approach is required for all but the very largest and capital-rich companies.

We also see growing use of connector-style tooling, where teams take an existing large language model and connect it to their product systems and knowledge sources to deliver a working AI feature faster.

Overall, hybrid development strikes a practical balance. Teams keep ownership of the core value proposition while deciding exactly where AI should enhance it. They also avoid over-investing in scarce specialist talent and infrastructure by embedding proven tools and outsourcing targeted expertise, thereby reducing build time and limiting cost exposure.

How is your organization approaching the development of AI solutions?



Note: “Others” (4%) combines responses that individually represented 1% each, including: providing AI solutions for enterprises; memory in-house with Azure OpenAI deployment; in-house with AI partners; and embedding/white-labeling options.

SECURITY: THE PRIMARY BOTTLENECK IN AI DEVELOPMENT

Across our partner interviews, we’ve found security to be a barrier to AI adoption. This follows a pattern similar to previous major digital transformations, such as social media, which raised privacy and data management concerns.

AI features frequently need to fetch information to answer questions or complete tasks. This data retrieval process risks exposure of customers’ personally identifiable information, such as names, email addresses, and home addresses. If a user can prompt an assistant to receive information they should not see, the product has created a permissions bug with real consequences.

“Many Copilot deployments remain in pilot. Security is the blocker. AI inherits the user’s access model, and most customers lack clarity on where sensitive data lives and who can access it.”

-Microsoft Partner

Making a product safe and controllable is the only way to gain the approval of security teams and build long-term customer trust. This reality highlights the critical role of service partners. A responsible and secure rollout often depends on partners who can implement governance and security controls correctly.



The partner plays a critical role in helping customers adopt AI securely by identifying sensitive data exposure and issues most customers aren't even aware of.


-Microsoft Partner

The primary challenge is often a surprise in cost. While customers may expect a straightforward rollout, projects frequently stall when the full scope of security requirements surfaces and the final service cost estimates are delivered.

HYPERSCALER MARKETPLACES GIVE YOU AN EDGE

When we asked Microsoft partners which benefits matter most, 34% cited access to Microsoft resources as their most impactful advantage. We read this as a signal that many teams rely on hyperscalers to get from product work to a shippable AI offering. That reliance is a practical necessity, as AI features run on cloud infrastructure, depend on platform services, and are frequently discovered or purchased through digital marketplaces.

Which program benefits have been most impactful for your partnership?



In our partner input, access to technical expertise consistently emerges as the value driver for partnering with a hyperscaler to get AI functionality launched.

Top 3 Benefits of Microsoft AI Cloud Partner Program



We also see partners tie hyperscaler support to execution outcomes such as faster time-to-market and stronger co-selling motion. That is consistent with the pattern in this report: hyperscaler resources and field collaboration are a material advantage when you are trying to ship AI responsibly and sell it at scale.

KEY TAKEAWAYS

1. Make revenue part of the product decision.

Our survey shows that product innovation is the leading objective for AI development at 62%, while revenue remains a priority at 51%. That combination means AI work is expected to ship and pay for itself. In exec discussions, the defensible position is a clear line from the AI feature to a commercial outcome.

2. Use hyperscaler resources to reduce execution risk.

Based on our findings, most partners will continue to rely on hyperscalers to move from AI product work to shipping and co-selling. Building your relationship with them can help your product reach revenue faster.

With hyperscalers, users have access to products and consultations that help them understand what success looks like in the Marketplace and how to meet launch requirements.

3. Plan for detailed security reviews and pushback on service costs.

Involve the customer's security team at the start of the project, while you are still shaping the feature and data access rules. When possible, sell and deploy through the hyperscaler marketplace, since customers often trust security solutions there more because they have already been vetted by the hyperscaler's partner and publishing requirements.

On services costs, assume the price tag will surprise customers and plan your messaging around that. Qualify early so you know what support the customer can handle themselves, then set expectations on the service scope and cost before the pilot starts. Be explicit about why the services work is required, including governance setup, permission design, and security validation, so it does not feel like an add-on that appeared late.

CONCLUSION

While the AI opportunity is large, the daily work for partnership leaders comes down to making the right bets, shipping responsibly, and proving value in ways customers will pay for.

In this next phase, customers will judge AI products less on novelty and more on whether they can trust the outcomes and explain results within their own organizations. Internally, executive teams will judge AI work on whether it strengthens the core business model or adds cost and risk with little to no lift.

This shift is why the next stage of the AI ecosystem is about product and partner discipline. For those leading these relationships, a strong roadmap must be explicit about which workflows matter and how to use partner resources to deploy AI responsibly.

Use this report as your reference point for conversations around investment decisions, feature choices, and the commercial logic that makes an AI product profitable and sustainable.

