snyk

REPORT

2026 State of Agentic AI Adoption

Anonymized Insights from 500+ Evo by Snyk AI Discovery Assessments

snyk.io

Introduction

Enterprise AI has crossed a structural threshold. What began as experimentation with prompts, chatbots, and copilots is rapidly evolving into agentic systems: AI software capable of reasoning, calling tools, accessing enterprise data, and taking autonomous action inside production environments.

This report provides an empirical view of that transition, based on anonymized AI Bill of Materials (AI-BOM) telemetry data across 500+ scans of customers' AI environments in Q4 of 2025 with Evo by Snyk. The analysis focuses on how AI is embedded, composed, and operationalized.

The findings reveal a clear inflection point. Agentic architectures are already in production, with roughly one in four organizations deploying autonomous agent frameworks or Model Context Protocol (MCP) servers. These systems represent a shift from passive AI use to platforms that can orchestrate workflows, integrate with internal systems, and execute actions with minimal human intervention.

At the same time, enterprise visibility and governance have not evolved at the same pace. Model-centric views dramatically underrepresent the true AI footprint. System-level analysis shows that each model is typically supported by two to three additional components, tools, datasets, packages, and orchestration layers. This fundamentally expands the operational and risk surface far beyond what most organizations track today.

This gap is no longer theoretical. As AI systems gain access to tools and internal services, risk compounds expanding from data exposure to unintended or unauthorized actions. Independent threat intelligence confirms that adversaries are actively probing enterprise AI deployments, exploiting the same blind spots that limit internal oversight.

The implication is clear: governing AI at scale now requires system-level visibility into agentic software, not just models or infrastructure. Organizations that treat AI as isolated components risk falling behind both operationally and defensively, while those that align adoption with disciplined governance are positioned to unlock the full strategic value of autonomous AI.

This gap is no longer theoretical. As AI systems gain access to tools and internal services, risk compounds expanding from data exposure to unintended or unauthorized actions.

At a glance: Top metrics from the report

1

Agentic adoption has arrived.

- 28.4% of organizations use agentic architectures (either Agents or MCP Servers)
- 20.4% of organizations use agentic frameworks (Agents).
- 18.2% deploy MCP servers.
- 10.1% use both, signaling platform-level agentic architectures.

Takeaway: Agentic AI is no longer experimental, it is actively entering production.

2

AI's footprint is ~3x larger than model counts suggest.

- 0.24 models per repository (model-only view).
- 0.68 total AI components per repository (system-level view).

Takeaway: Models are just the visible tip of a much larger AI supply chain.

Note: Density metrics are calculated across the total volume of scanned repositories, including legacy applications and 'maintenance-mode' codebases. A density of 0.24 in a brownfield environment signals significant penetration of new development initiatives.

3

Depth, not volume, defines AI maturity.

- Healthcare, financial services, and specialized industries average ~50 AI components per account.
- Technology firms average ~30 AI components.

Takeaway: Regulated and high-stakes industries are embedding AI most deeply.

4

External dependencies dominate the AI stack.

- 82.4% of AI tools come from third-party packages.
- Respondents report using roughly 5 external tools for every 1 custom tool.

Takeaway: AI innovation is accelerating — but so is supply chain exposure.

5

Data lineage is the weakest link

- Only 0.46 datasets per model are explicitly tracked.

Takeaway: Most deployed models lack clear data provenance, limiting governance and compliance readiness.

snyk

Chapter 1: From experimentation to agentic systems

What began as experimentation with prompts and chat interfaces is now maturing into agentic AI systems embedded directly into enterprise operations.

This shift is already visible in the data. Across the analyzed organizations, 21.8% have deployed agentic frameworks, signaling autonomous behavior in real systems, while 19.7 % have deployed Model Context Protocol (MCP) servers, indicating investment in the infrastructure required to connect AI to enterprise tools, services, and data. These agents are not isolated experiments. They're being intentionally integrated into production systems. AI has moved from consultation to execution.

Binary indicators of AI presence, however, understate the true depth of adoption. When measured solely by models, AI usage appears limited. When the full AI supply chain is considered — including models, tools, packages, datasets, and orchestration layers — the footprint expands dramatically. On average, each model is supported by two to three additional AI components, revealing that enterprises are deploying interconnected AI systems, not standalone models.

This chapter examines how AI adoption progresses from experimentation to platform-level integration. By analyzing adoption intensity, system composition, and industry patterns, it distinguishes exploratory usage from production-grade, agentic architectures and identifies where AI is becoming foundational to enterprise operations.

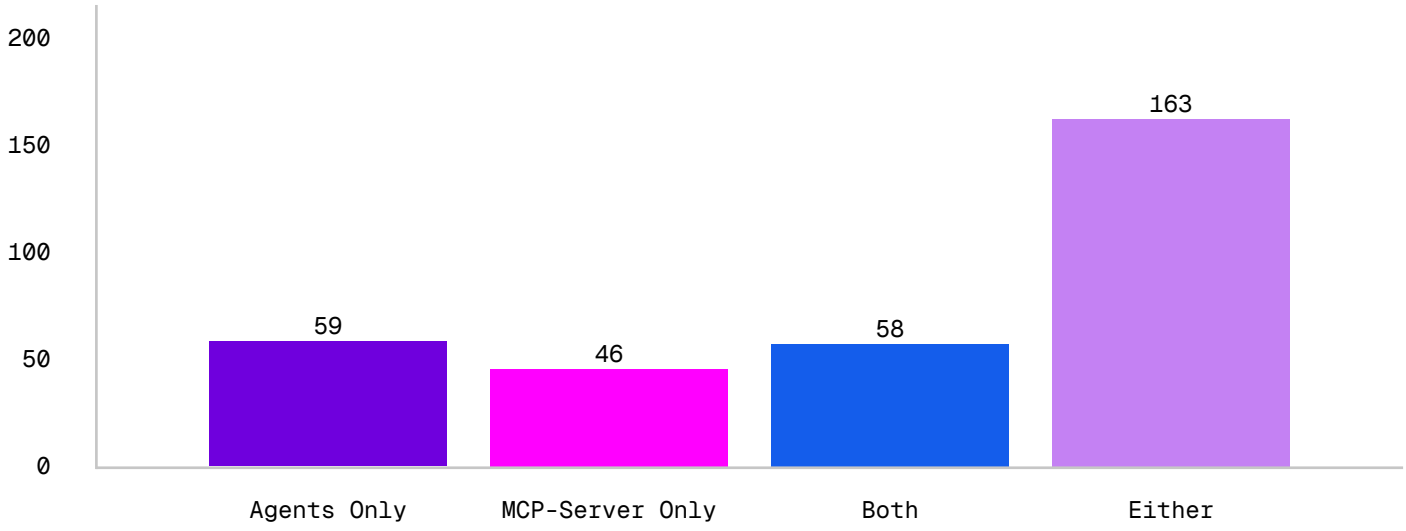

To assess the early adoption of agentic and platform-level AI architectures, we analyzed customer-level usage of agentic frameworks (Agents) and Model Context Protocol (MCP) servers. This metric reflects adoption rate, measuring the percentage of organizations that have deployed at least one instance of each technology, rather than the total number of components in use.

Across the analyzed customer base, this helped track the industry's move from simple chatbots to autonomous Agents.

20.4% of active organizations have adopted agentic frameworks, such as autonomous agent frameworks (e.g., LangChain, AutoGen), indicating active experimentation or deployment of autonomous AI behaviors. This is the strongest signal of AI maturity. Organizations are rapidly moving beyond "chat" to "action."

18.2% of active organizations have deployed MCP servers, reflecting early investment in infrastructure that connects AI systems to internal tools, data, and services. This represents the cohort that has crossed the critical threshold into agentic architectures. MCP is effectively becoming the standard "connective tissue" for the agentic ecosystem.

Furthermore, the significant overlap **(10.1% of respondents use both)** indicates that organizations experimenting with autonomous AI are simultaneously investing in the infrastructure required to operationalize those systems. Rather than treating agents as isolated experiments, early adopters appear to be building platform-level architectures that enable AI systems to interact with enterprise resources in a structured, repeatable way.

IMPLICATIONS

Agentic AI adoption is no longer niche. More than 1 in 4 organizations has already moved beyond prompt-based AI toward autonomous systems.

Platform infrastructure is emerging alongside autonomy. MCP servers act as a connective layer between AI models and enterprise systems, expanding the scope and potential impact of AI deployments.

Risk scales with adoption. As agents gain access to tools and internal services, operational and security risks shift from data exposure to unintended or unauthorized actions.

This pattern marks a transition from AI usage focused on content generation and decision support toward systems capable of autonomous action, orchestration, and execution.

From presence to adoption intensity

The presence of agents or MCP servers signals early maturity, but it does not capture the extent to which AI is embedded across the enterprise. True AI adoption is reflected in density, distribution, and composition how many components exist, how widely they are deployed across repositories, and how tightly they are integrated into everyday systems. As organizations graduate beyond binary indicators of adoption to quantify the operational footprint of AI, this paradigm reveals a moment when AI shifts from isolated capability to foundational infrastructure.

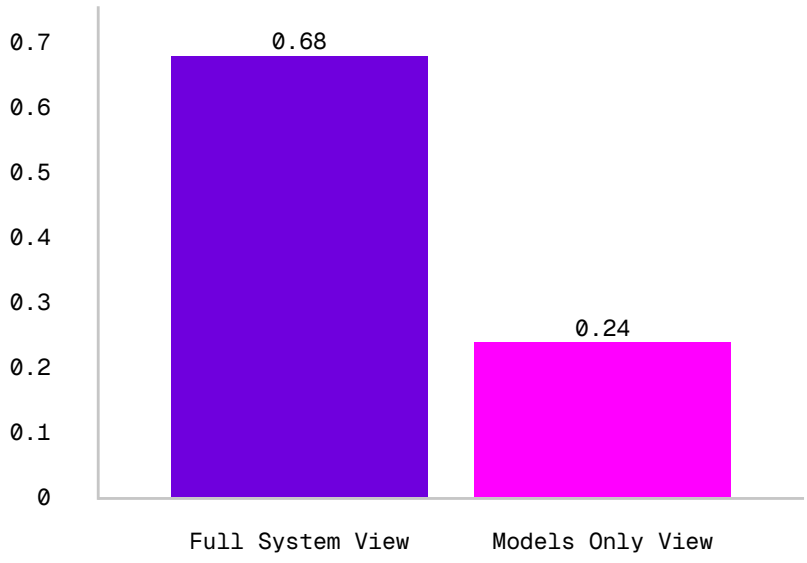

To understand the true scale of AI adoption within enterprise codebases, we measured AI density per repository using two approaches: a model-only view and a system-level view that includes the full AI supply chain. This metric reflects the average number of AI-related components present in a repository and provides insight into how AI is operationalized in practice.

Findings

When AI usage is measured solely by the presence of models, the observed density is relatively low, averaging 0.24 AI components per repository. This narrow view suggests that AI adoption remains limited and that the associated operational footprint is modest.

However, when the analysis is expanded to include the broader AI supply chain, including packages, tools, datasets, and other supporting components, AI density increases substantially to 0.68 components per repository, representing nearly a threefold (2.8x) increase over the model-only measurement.

The disparity between these two measurements highlights a consistent pattern: AI is rarely deployed as a standalone model. Instead, models are typically embedded within a network of supporting components that enable data ingestion, orchestration, execution, and integration with other systems. On average, each model is accompanied by two to three additional AI-related components, significantly increasing the complexity of the repository.

This finding indicates that AI adoption at the code level is more extensive and interconnected than model-centric metrics alone would suggest.

IMPLICATIONS

Model-focused metrics underrepresent AI adoption. Organizations relying on model counts as a proxy for AI usage may significantly underestimate the size of their AI footprint.

Operational complexity grows with AI systems, not models. Supporting components introduce additional dependencies, integration points, and potential failure modes.

Security and governance requirements extend beyond models. Datasets, tools, and third-party packages contribute meaningfully to the overall risk profile and must be included in oversight strategies.

AI adoption by industry: Volume vs. intensity tells different stories

To understand how deeply AI is embedded across industries, we analyzed adoption from two complementary perspectives: total AI component volume and average AI components per account (intensity). Viewed together, these charts reveal not just who is using AI, but who is operationalizing it at scale.

Chart 1: Total AI components by industry — where AI is most widespread

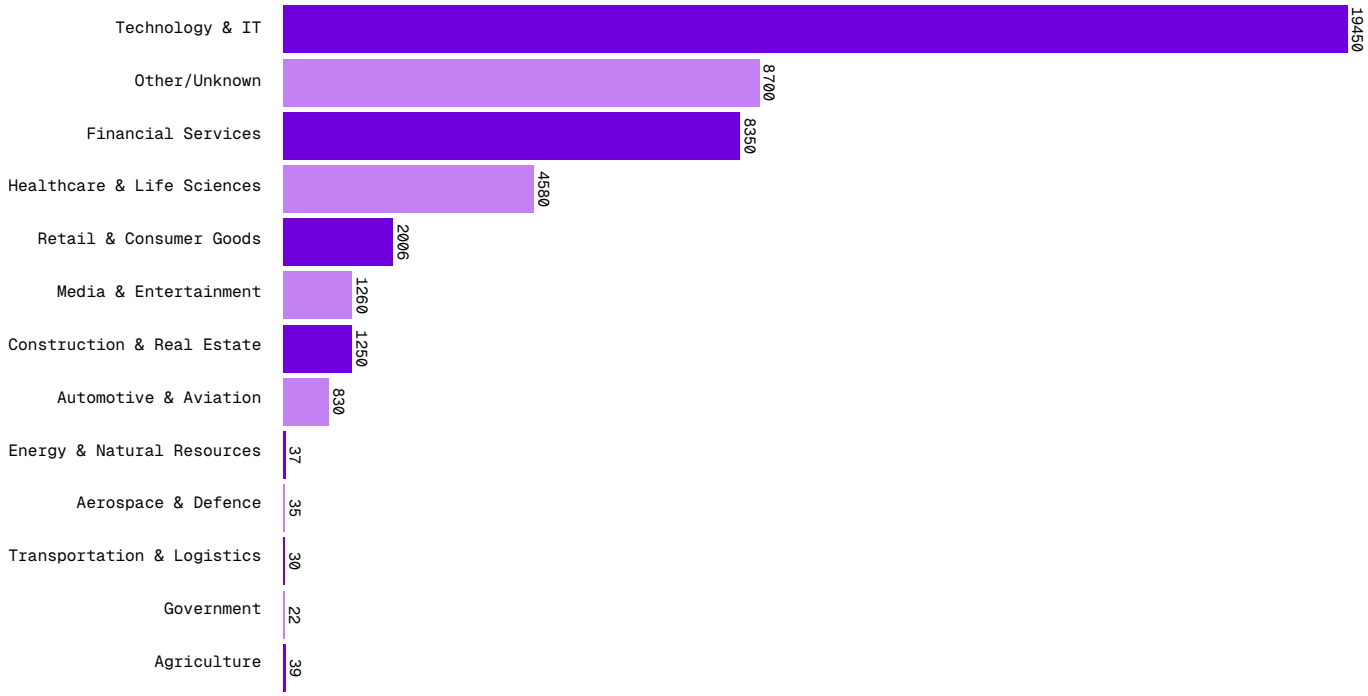

The total volume view shows where the total sum of all AI components, such as models, tools, and datasets, is most prevalent across the customer base. It measures breadth and shows where AI is common, but not how deeply it is embedded within any single organization. Here, Technology & IT dominate, accounting for the largest share of total AI components by a wide margin. This reflects the sector's sheer scale: hundreds of technology companies integrating AI into products, platforms, and developer workflows.

Financial Services and the "Other" category follow, with Healthcare & Life Sciences trailing further behind in total count. At first glance, this suggests that AI adoption is primarily a technology sector phenomenon, with regulated industries playing a secondary role.

This reflects the sector's sheer scale: hundreds of technology companies integrating AI into products, platforms, and developer workflows.

Chart 2: Average AI components per account-where AI is deepest

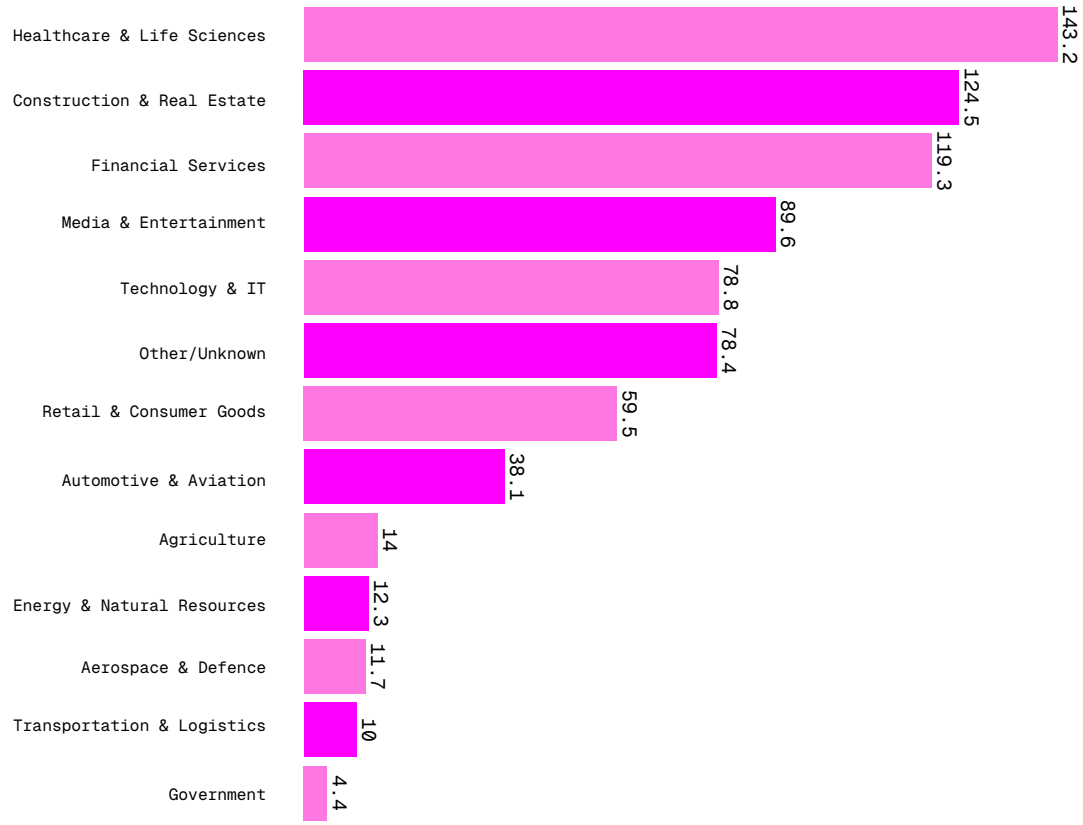

The intensity, defined as the average number of all AI components per active account, tells a very different story. AI intensity reveals depth. Regulated and specialized industries are building larger, more complex AI systems per organization than the average technology company. When AI components are normalized per account, Technology & IT moves down, replaced by industries with far denser AI estates.

This shift is especially pronounced in Healthcare & Life Sciences, which moves from a mid-tier position in total volume to near the top in intensity. This indicates that healthcare adoption is not superficial or experimental. AI is being woven deeply into research, diagnostics, and operational workflows.

Regulated and specialized industries are building larger, more complex AI systems per organization than the average technology company.

Chart 3: Agentic and model-focused components: From usage to autonomy

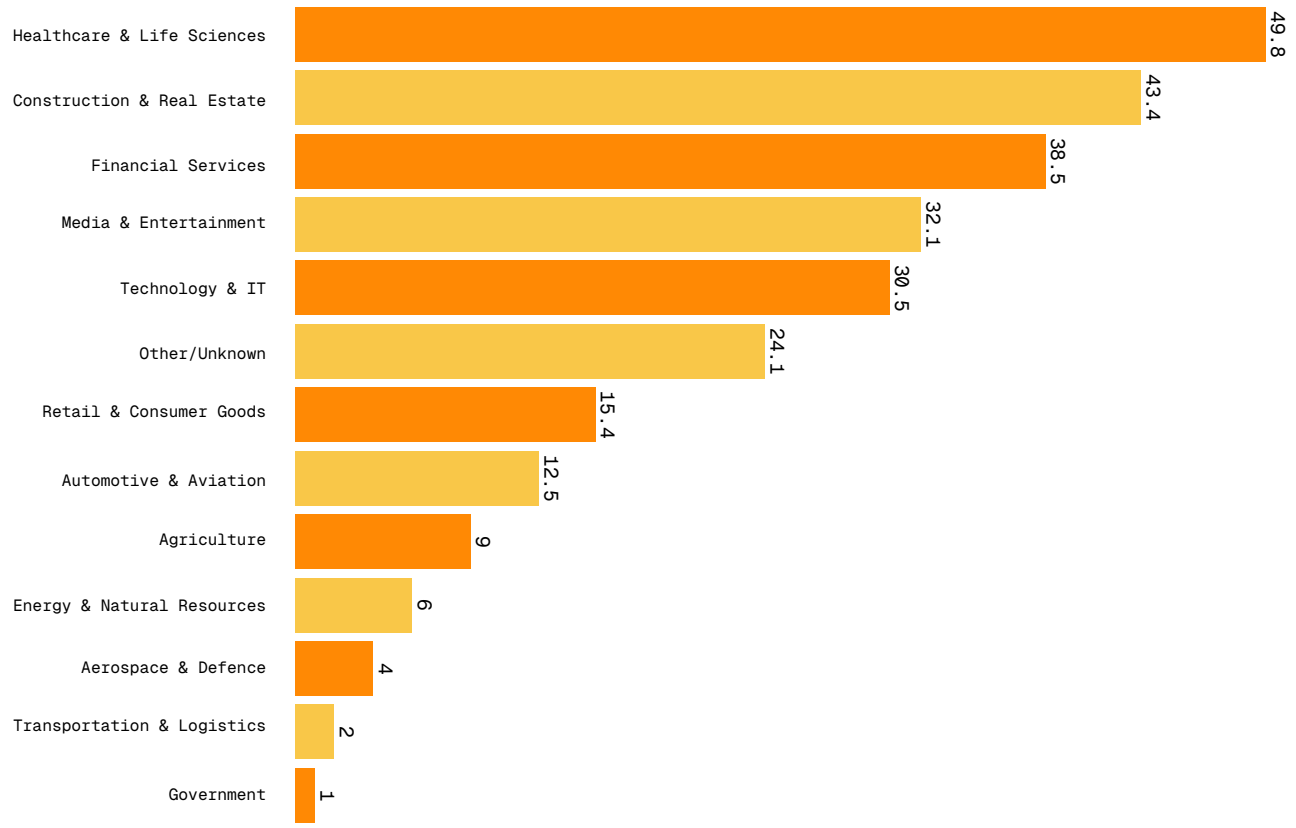

When the analysis narrows further to core system components, models, MCP servers, and agents (focusing on the brain of the AI system), the signal becomes even clearer. High-intensity industries remain leaders, demonstrating that their AI footprint is not driven by scattered experimentation or tooling sprawl, but by intentional, production-grade architectures.

These components represent a shift away from simple API calls toward orchestrated, autonomous systems where models use tools, access data, and execute actions. The persistence of Healthcare, Financial Services, and Construction & Real Estate at the top confirms that agentic AI is emerging first where trust, compliance, and operational impact matter most.

High-intensity industries remain leaders, demonstrating that their AI footprint is not driven by scattered experimentation or tooling sprawl, but by intentional, production-grade architectures.

The top 5 most adopted AI services

The AI supply chain data reveals a clear power structure rather than a fragmented market. Adoption is now decisively concentrated around two dominant providers: OpenAI and Google which together account for approximately 60% of all AI services in use. OpenAI is currently the clear leader at 43.6%. This helps confirm that most enterprises are standardizing on just one or two primary foundation model platforms for core workloads.

Beneath this top tier, a stable, more established layer has emerged, led by Hugging Face, Meta, Microsoft, and Anthropic signaling that enterprises are deliberately diversifying their model portfolios to balance capability, safety, and vendor risk. Most notably, the refined data reveals a new challenger class: Mistral, Alibaba, and Cohere, demonstrating that regional and specialized providers are already penetrating real production environments, particularly for embeddings and cost-efficient inference. Even after normalization, a meaningful long tail remains, underscoring continued experimentation outside standardized platforms and based on developer preferences tied to different projects.

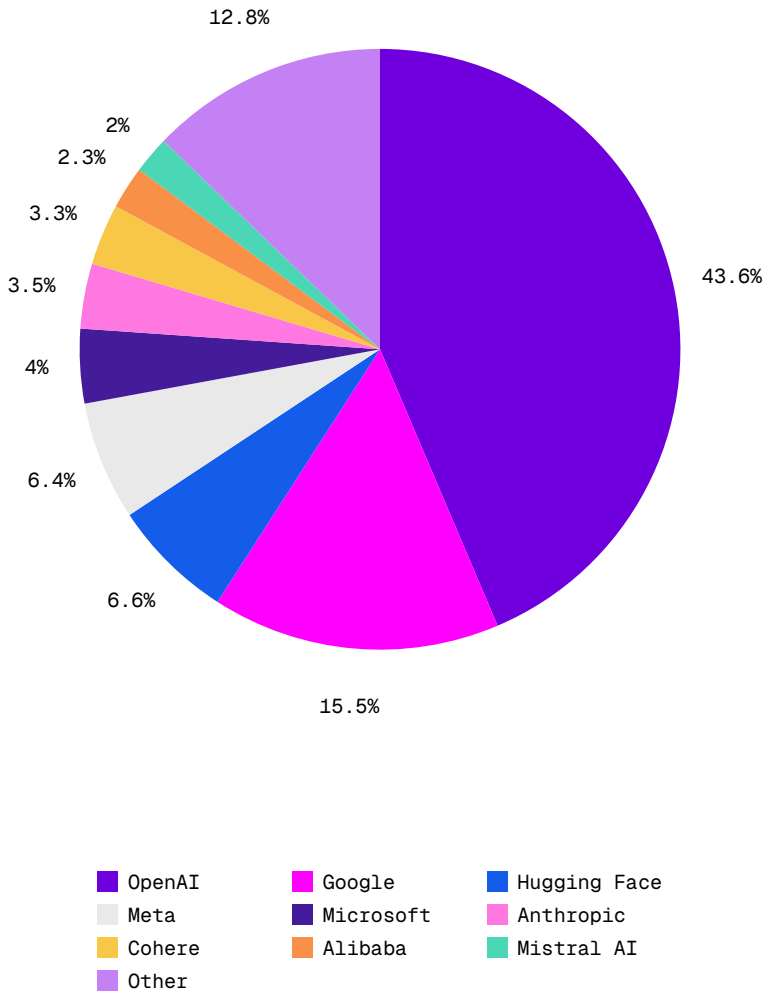

The result is a hybrid AI reality. A highly concentrated core supply chain surrounded by a rapidly evolving edge, increasing both innovation velocity and the need for centralized governance, visibility, and risk control.

Complexity of AI adoption (model density per customer)

We analyzed adoption complexity by measuring the number of unique models deployed per active customer. This reveals a stark "barbell" distribution in the market.

- **The power users (41.7%):** A large group of 41.7% of active customers use 5 or more distinct models. This suggests a high degree of heterogeneity, where teams orchestrate multiple vendors (e.g., GPT-4 for reasoning, Llama 3 for local tasks) rather than standardizing on one.
- **The emerging builders (37%):** On the other end of the spectrum, 37% of active customers have 0 detectable models. However, notably, in nearly half of this group (**18.8%**) AI tools, agents, or frameworks are being detected, indicating they are actively building the infrastructure for AI even if they haven't yet deployed a traceable foundation model.
- **The middle ground:** Only ~21.2% of customers sit in the middle with 1-4 models, suggesting that once companies move past the initial build phase, they scale rapidly to multi-model architectures.

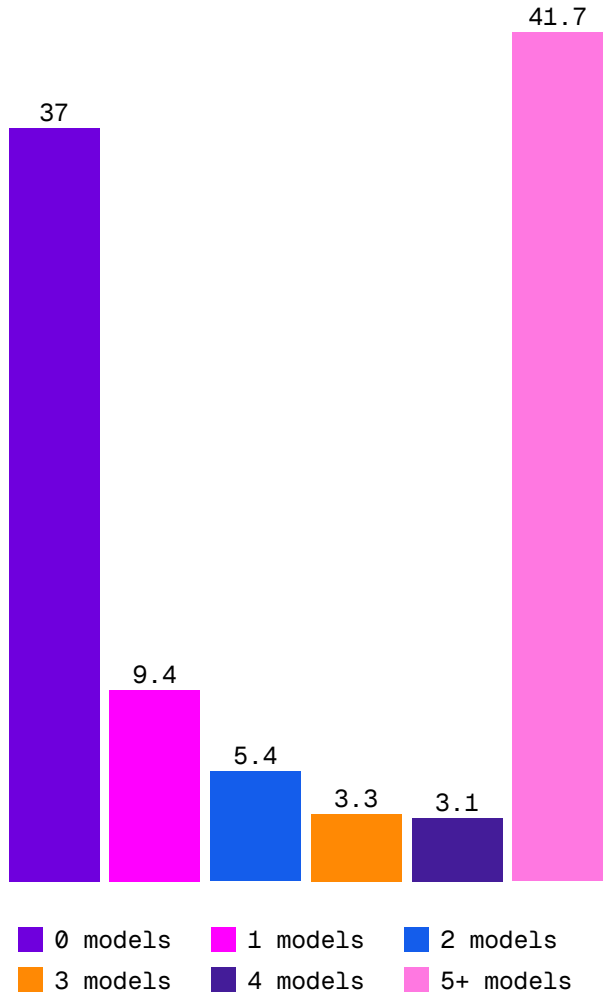

Customers using multiple models (particularly five or more) are likely to employ multi-vendor solutions or a diverse mix of specialized models (e.g., GPT-4 for reasoning, Haiku for speed). Single-model deployments generally reflect simpler use cases such as basic chatbots.

Chapter 2: Emerging structural risks in agentic AI systems

This chapter focuses on the operational and security implications of enterprise AI adoption. As AI systems become more autonomous and interconnected, the fundamental nature of risk is shifting from individual components to system composition and behavior.

Our analysis of the 500+ scans of customers' AI environments with Evo by Snyk highlights several structural characteristics that define this new landscape:

It is important to recognize that these patterns are not failures; they are the natural outcomes of modern, high-velocity software development. However, they create environments where autonomous systems can act in unexpected ways without centralized visibility.

The new risk class: Agentic exposure

Traditional AI risk discussions have largely focused on data leakage or model misuse. Agentic systems introduce a different, more volatile class of risk: Agentic exposure.

When AI systems can autonomously call tools, access APIs, pull models, or execute workflows, risk shifts from what AI knows to what AI can do. Agentic exposure risk emerges when these autonomous systems operate without a clear understanding of their composition, permissions, and execution paths.

External threat context: Why this matters now

This risk is not theoretical. Independent threat intelligence confirms that adversaries are already exploiting this reality. Recent campaigns demonstrate systematic enumeration of enterprise LLM infrastructure, using benign prompts and behavioral fingerprinting to identify deployed models and exposed integrations. These techniques specifically target hidden proxies, internal services, and indirect model access, precisely the areas least visible through traditional monitoring. Threat actors are mapping the agentic surface now.

Why cloud-only views are insufficient

To mitigate this, organizations often turn to cloud and CSP-centric security tools. While essential, these approaches cannot explain or govern agentic AI on their own. Cloud security answers where the infrastructure runs and who has access. It does not answer which agents exist, which tools they can invoke, or how models, tools, and data compose into autonomous behavior.

Supply chain composition and complexity

Agentic AI risk arises from software composition, not from infrastructure configuration. This section examines the specific data points from supply chain dependencies to governance gaps that reveal where these risks are hiding in your environment today.

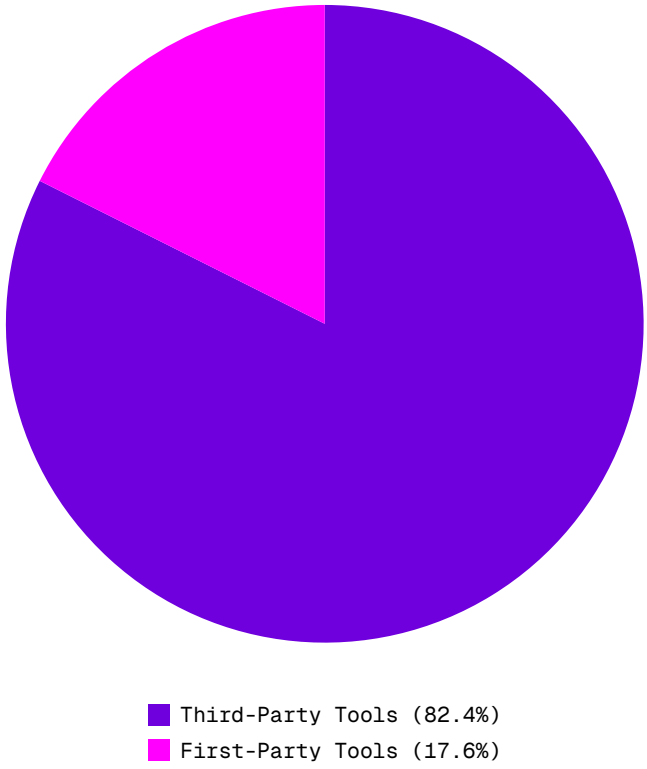

We looked to compare tool components from third-party libraries versus those defined locally in the code. The data shows that 82.4% of AI tools across the customer base are sourced from third-party packages, while only 17.6% are locally developed. This translates to roughly five third-party tools for every one custom-built tool, highlighting strong reliance on external ecosystems to accelerate AI development. While this approach reduces deployment time, it also broadens the dependency surface, increasing exposure to supply chain vulnerabilities, unvetted updates, or malicious package insertion.

The open source and proprietary model mix

Proprietary models such as GPT-4, Gemini Pro, and Claude models slightly lead in production, representing 57% of models, while Open source accounts for 43%. Enterprises rely on proprietary models for core reasoning and generative tasks, but a large minority of AI adoption is driven by the long tail of open source models used for embeddings, text analysis, and computer vision, with top examples including ResNet50, BERT, RoBERTa, and sentence-transformers. This highlights a hybrid AI ecosystem where open source drives volume and flexibility, and proprietary models drive strategic value.

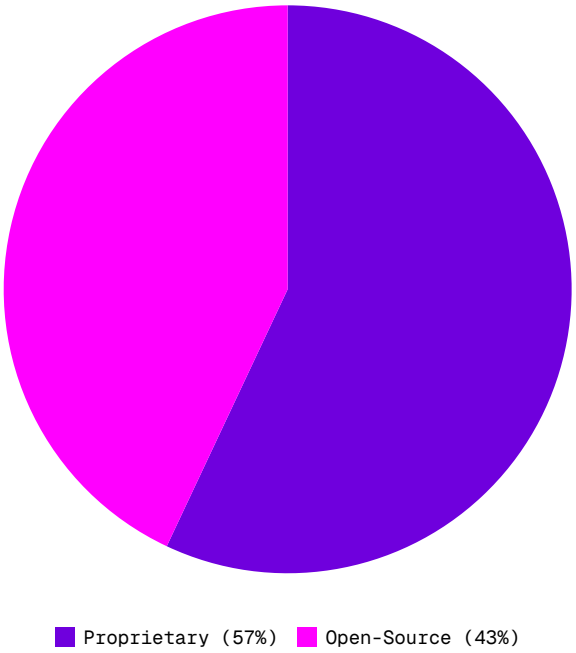

However, this long tail of open source models also introduces potential risk. Unlike proprietary models, open source models vary in quality, may lack built-in safeguards, and often require local infrastructure for deployment. The large number and distributed nature of these models make inventory, versioning, and monitoring challenging, increasing the potential for security, operational, or regulatory issues. Enterprises that adopt many open source models need structured governance and oversight to manage these risks while still benefiting from their flexibility and task-specific utility.

The governance gap

Dataset-to-model ratio: On average, only 46% of accounts with Models have tracked datasets to support them. This suggests that more than half (54%) of accounts with active models lack explicitly associated training or fine-tuning data in the AI Bill of Materials (AI-BOM).

This gap severely limits data provenance, bias assessment, regulatory compliance, and incident response. Even teams that are successfully deploying models are largely failing to document the data that powers them.

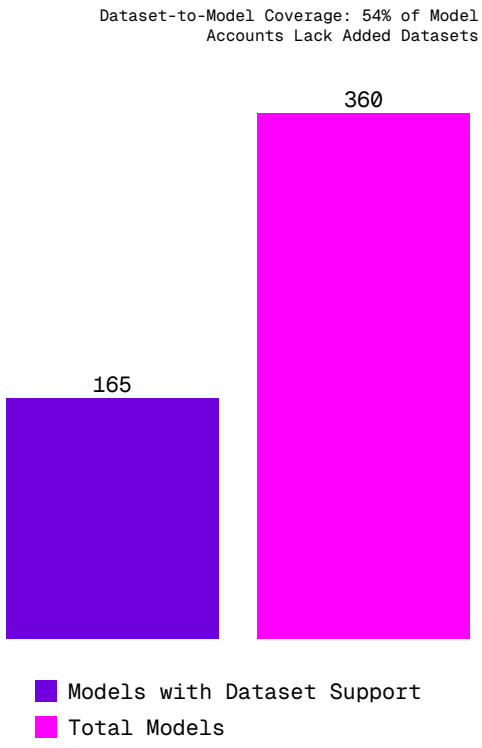

Prompt template exposure: Hardcoded prompt templates were identified in 4.4% of active accounts. Although low in volume, these represent high-value risk points.

Untracked or exposed system prompts can lead to IP leakage and indicate immature “PromptOps” practices where prompts are treated as "magic strings" in code rather than managed assets.

IMPLICATIONS

External dependency risk is the default: Heavy reliance on third-party tools expands the attack surface and creates potential supply chain vulnerabilities.

Data lineage visibility is insufficient: Missing dataset-to-model relationships impede governance, auditability, and compliance readiness.

Risk compounds across layers: The combination of untracked datasets, widespread third-party dependencies, and occasional hardcoded prompts amplifies overall system complexity and operational risk.

Developer-to-AI asset ratio: Supervision overhead

Across over 250k developers managing 46,620 AI assets, the ratio is approximately 0.18 AI assets per developer, roughly one AI asset for every three developers. While this may appear low, each AI asset represents a specialized, high-risk component. In practice, a typical team of 10 developers collectively manages approximately 2 AI models or agents, creating non-trivial supervision overhead.

Key takeaways:

1. AI assets are not yet "every developer's responsibility," but they require dedicated attention from small teams, introducing operational risk.
2. Even with limited adoption, the complexity and risk per asset are high, making oversight, monitoring, and governance essential.
3. Organizations need tools and processes to manage AI dependencies safely, ensuring that human oversight scales with AI adoption.

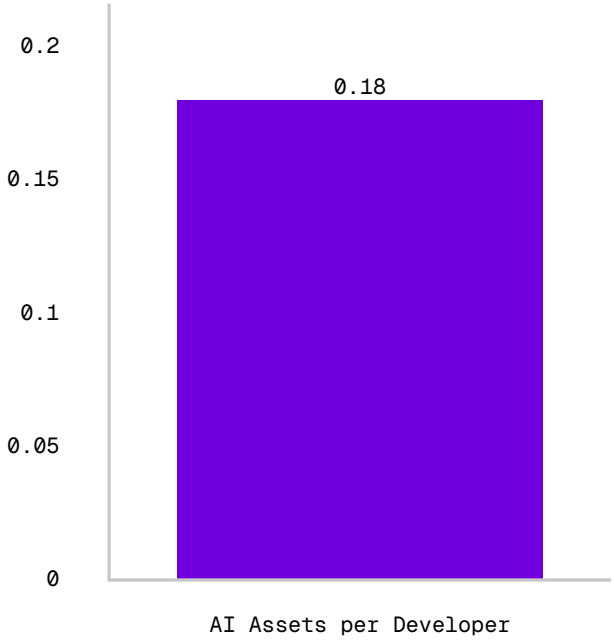

The AI-centric road ahead

Enterprise AI is no longer an experimental capability. It is becoming a foundational infrastructure. The analysis of 500+ scans of customers' AI environments with Evo by Snyk reveals a market that has moved faster than top-level metrics suggest. While a cohort of "Emerging Builders" (37%) is still laying the architectural groundwork, a massive segment of "Power Users" (41.7%) has already scaled to complex, multi-model production environments.

Most critically, a meaningful minority of organizations, roughly 28.4% have already crossed the threshold into agentic architectures. These organizations are no longer just calling APIs; they are embedding autonomous systems, Model Context Protocol (MCP) servers, and agentic frameworks deeply into production software.

The "agentic gap" in visibility and governance has not caught up to this reality. System-level metrics reveal that for every model deployed, organizations introduce nearly three additional software components, including tools, datasets, and packages. This expands the operational risk surface threefold, yet traditional governance often remains fixated solely on the model.

This disconnect creates structural vulnerabilities:

- **Supply chain exposure:** With 82.4% of AI tools coming from third-party packages, organizations are importing vast amounts of unmanaged external code.
- **Data blind spots:** With less than half of models (0.46 ratio) having tracked datasets, data provenance remains opaque.
- **Operational risk:** As agents gain the ability to execute actions, the lack of visibility into these dependencies allows operational errors or malicious actors to propagate undetected.

The path forward is clear: **AI governance must evolve from model-level abstraction to a system-level understanding of agentic software.**

Organizations that aim to harness AI at scale must move beyond siloed monitoring. They require comprehensive supply chain governance, rigorous dataset lineage, and oversight frameworks that can see the entire system, not just the model inference. Enterprises that deeply and autonomously integrate AI without the corresponding oversight expose themselves to disproportionate risk.

snyk

snyk.io

Data sourced from 500+ customers who successfully scan their AI environment in Q4 2025 with Evo by Snyk.