MANIFEST

Beyond the Black Box: How AI is Forcing a Rethink of Software Supply Chain

What 300+ Security Leaders Reveal About AI and Software Supply Chain Risk

Table of Contents

Executive Summary

Organizations have made progress securing the software supply chain, but AI is forcing a rethink of what "supply chain security" means. AI systems reintroduce familiar problems in new forms: opaque dependencies, uncertain provenance, and unclear accountability—now spanning models, datasets, agents, and third-party AI services.

Visibility is slipping: 63% of organizations report shadow AI usage, and AppSec/Product Security consistently see more gaps than leadership. Meanwhile, legacy approaches are struggling to keep up, teams are overwhelmed by noisy signals that drive work without consistently improving real-world risk decisions. The next frontier is moving beyond the black box with verifiable transparency and operational control across the software and AI supply chain.

A Novel Form of Supply Chain Risk

As AI enters every product and workflow, organizations face a new supply chain risk: systems they can't reliably see, explain, or verify. Security teams need practical mechanisms to inventory AI components, validate what's introduced through vendors and developers, and manage risk continuously as systems evolve.

We recently surveyed 300 security leaders and practitioners across major industries and regions (U.S. and EMEA, including France and Germany), and this report includes country-specific and industry breakout insights for notable findings.

> "We're seeing history repeat itself. AI systems are recreating the same supply chain blind spots the industry has spent a decade trying to solve in open source, including unverified dependencies, limited provenance, and unclear accountability. The fact that 63% of organizations report shadow AI usage underscores how much exposure remains unmanaged. Governance frameworks are evolving, but organizations need operational mechanisms to inventory and manage AI supply chains in practice."
>
> — Daniel Bardenstein, CEO & Co-founder, Manifest

Key Findings

60%

of respondents create SBOMs, but more than 50% of organizations aren't consuming or managing their SBOMs in practice.

72%

of respondents say that their organizations already receive verifiable transparency data during procurement most or all of the time.

63%

of survey participants report their organizations have shadow AI, indicating widespread unmanaged usage.

56%

of survey participants agree that SCA tools are noisy and 51.7% believe that SCA findings lead to meaningful remediation.

40%

of AppSec teams believe there is a lack of mature AI security practices, while nearly 80% of security executives report thorough practices, reinforcing the gap between perception and reality.

A NOTE BEFORE YOU BEGIN

Application Security Teams are the Canary in the Coal Mine

Manifest researchers noted throughout the study that responses from Application Security (AppSec) teams differed from their leadership counterparts. AppSec teams encounter the operational reality first—where tooling, ownership, and workflows break down across the software and AI teams build and buy. Because of this hands-on exposure, their perspectives provide a more accurate indicator of readiness. Heeding these signals offers a realistic snapshot of where software supply chain security and AI security stand today. Their feedback is especially valuable because true technology transparency requires the ability to inspect, inventory, and analyze every component of digital systems, from software dependencies to supplier code and machine learning models.

AppSec and Product Security report more noise, more gaps, and more friction. “Noise” most often appears as high volumes of alerts without clear exploitability context, inconsistent component identifiers, and duplicate findings across tools—conditions that create friction without materially improving risk outcomes.

When ownership is ambiguous, organizations rarely maintain a shared record of what software is deployed to production, which risks have been formally accepted, or how decisions were made, complicating both audit readiness and incident response.

These findings point to a deeper issue than tooling alone. Even as organizations add new security capabilities, fragmented ownership and disconnected workflows prevent teams from translating signals into consistent risk reduction. Addressing this requires tighter collaboration across AppSec, Product Security, Legal, and Compliance, supported by shared tooling and a common system of record. In practice, that means a centralized intake of SBOMs—including supplier-provided artifacts—normalization across formats, consistent policy enforcement, and continuous monitoring as software evolves and new vulnerabilities emerge. Without shared ownership and shared records, improvements in tooling will continue to generate more activity—but not necessarily better outcomes.

47%

of respondents reported siloed teams and unclear ownership as the top blockers in improving security and reducing the risk of their software supply chain.

Even as organizations add new security capabilities, fragmented ownership and disconnected workflows prevent teams from translating signals into consistent risk reduction.

THE SBOM PARADOX

Generated But Not Operationalized

Most organizations have checked the SBOM box. They can generate them, often through open-source generators or their existing Software Composition Analysis (SCA) tools. But generation is only the first step—one that more than half of organizations haven't actually begun.

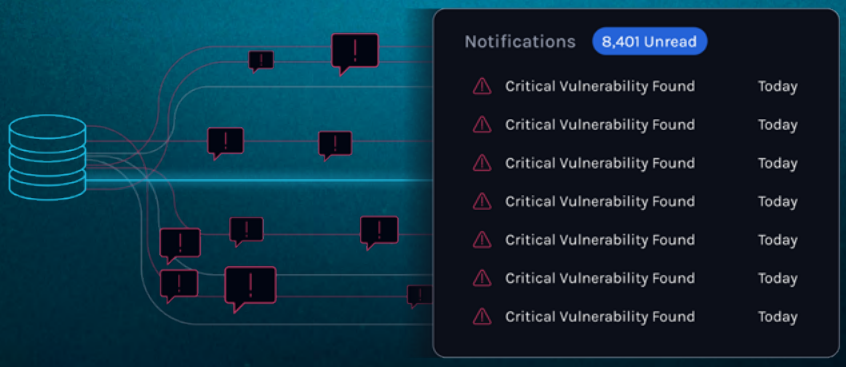

Without centralized intake, normalization, policy enforcement, or continuous monitoring, SBOMs remain noisy, point-in-time artifacts that provide visibility without assurance. What's missing is continuous supply chain assurance.

SCA alone doesn't provide the capabilities needed for ongoing risk management.

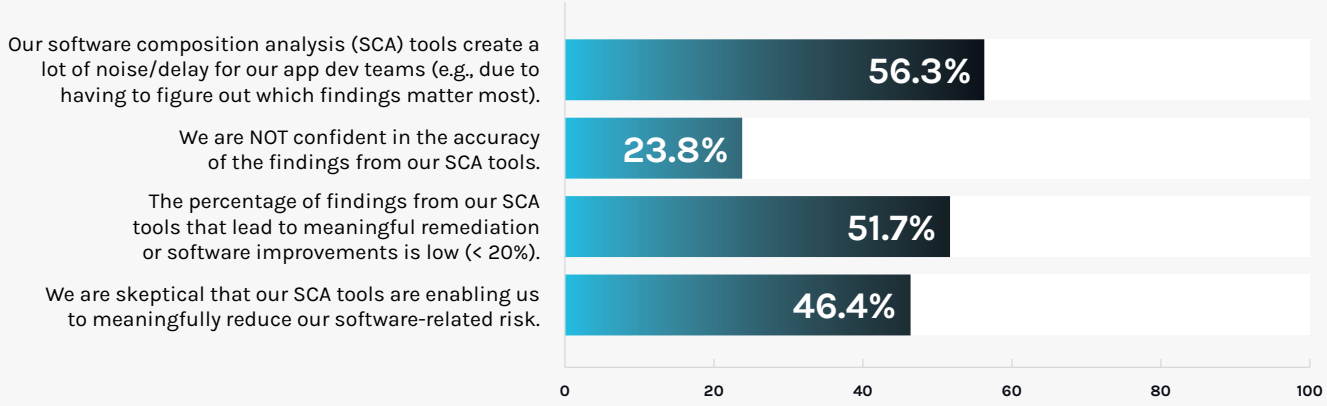

We expected to find that legacy SCA tools surface thousands of low-priority findings without improving visibility of reducing real-world risk, and that's exactly what we found. More than half of respondents agreed that SCA tools are noisy (56.3%). Yet a similar share reported that SCA findings still prompt remediation work (51.7%).

At the same time, nearly half of survey participants (46.4%) remain skeptical that SCA is enabling meaningful risk reduction. **In effect, SCA is driving action, but many teams are unconvinced it improves understanding of which risks actually matter.**

Regional differences reinforce this tension. Respondents in EMEA are significantly less skeptical of SCA's ability to reduce software risk (28.2%) than their U.S. counterparts (53.2%), suggesting potential cultural differences, as Americans tend to be more polarized.

Industry Trends

Healthcare stands out as the most skeptical industry.

69% of healthcare respondents believed SCA tools create a lot of noise/delay.

58% of healthcare respondents were skeptical of SCA tools reducing software-related risk.

This skepticism likely reflects the operational realities of healthcare environments, where disruption or misprioritized risk can have direct consequences for patient safety. Recent high-profile ransomware attacks have reinforced this concern, demonstrating how cybersecurity events can translate into real-world impacts on patient care and, in some cases, mortality. Against this backdrop, it's plausible that healthcare organizations approach automated findings from SCA with a more conservative or cautious mindset.

At the same time, this critical stance may signal an opportunity rather than resistance. Healthcare teams may be hesitant to fully trust point-in-time SCA results. SBOMs—paired with continuous monitoring—can offer a more transparent and auditable approach to tracking risk across clinical systems and internet-connected medical devices already in market. SBOM-driven visibility, especially when tied to ongoing monitoring of deployed devices, could better align with healthcare's need for high assurance, traceability, and patient-safety-centric risk decisions.

The Growing Need for Technology Transparency

Manifest researchers observed a clear shift toward technology transparency, driven by the need for continuously updated visibility into the software—and increasingly AI—that organizations build and buy.

SCA remains the most widely adopted capability, used by nearly 4 in 5 organizations (79.6%).

SBOM usage, while lower overall at 41.8%, is notable given its relative newness—and increases steadily with organizational size.

Adoption rises from 31.7% among the smallest organizations to 59.4% for the largest, suggesting that external pressures like regulatory requirements and industry standards may drive adoption for larger firms.

These organizations must manage higher volumes of third-party components and supplier artifacts, making SBOM collection, validation, and tracking an operational necessity, not just a compliance exercise.

SBOM Usage: The practice of generating SBOMs or managing and consuming them, including storage, distribution, analysis, and operational use.

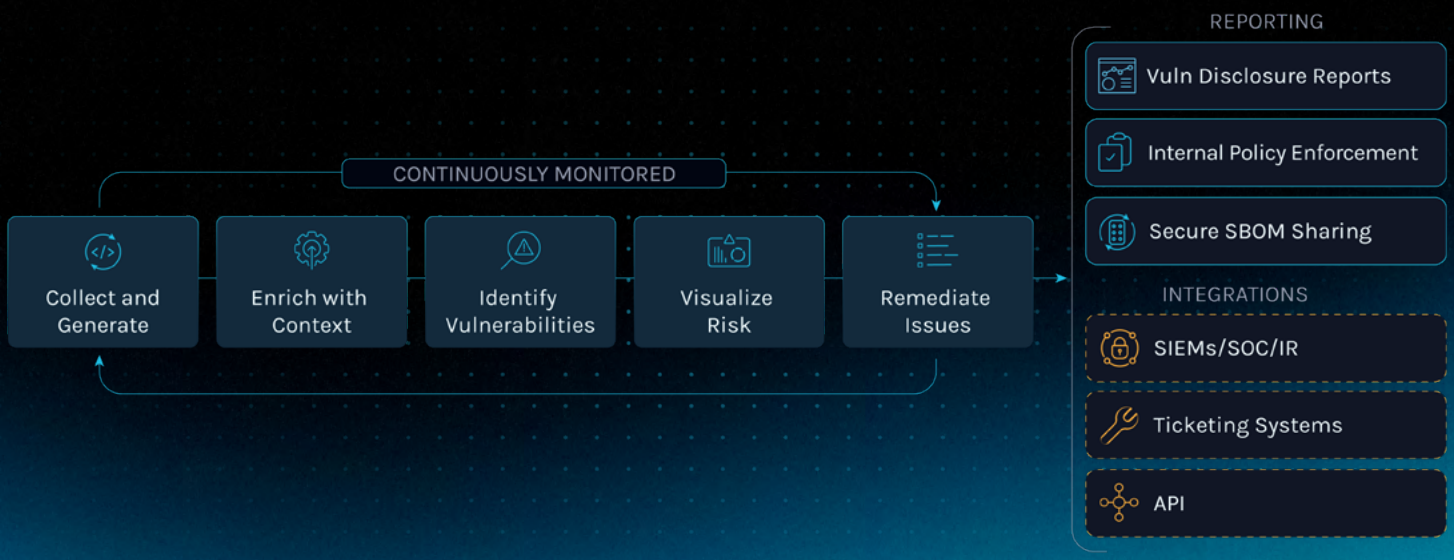

Other patterns are consistent with expectations. AppSec teams tend to prioritize Dynamic Application Security Testing (DAST) and Application Security Posture Management (ASPM) in line with executive security priorities. U.S. organizations lead European counterparts (FR/DE) in adopting SCA, SBOM, and ASPM. Industry differences are pronounced: manufacturing reports the highest SBOM adoption (50.0%), while technology companies dominate SCA usage (85.2%) but lag on SBOM adoption (37.5%). These divergences may suggest sector-specific drivers—DevOps velocity in technology versus regulatory and lifecycle pressures in manufacturing.

Across segments, the differentiator is not whether SBOMs exist, but whether organizations have the operating model and infrastructure to govern and use them continuously. While most organizations generate SBOMs, fewer than half are actively using tools to consume, manage, and operationalize them. Many SCA vendors provide basic SBOM generation, but stop short of delivering the ongoing risk management capabilities teams need: centralized intake, normalization across formats, policy enforcement, and continuous monitoring as vulnerabilities and components evolve.

This exposes a critical maturity gap. SCA alone delivers point-in-time visibility; it does not provide continuous software supply chain assurance. Because risk changes after release, SBOMs are most effective when treated as continuously refreshed inventories—kept current as software evolves rather than generated only at major release milestones. Smaller organizations often struggle to reach this level of maturity due to resource and structural constraints, reinforcing the need for solutions that reduce operational complexity.

The findings underscore the essential role of SBOMs relative to SCA. SBOMs enable organizations to move from knowing what components exist to actively governing risk across the software supply chain. SCA tools cannot meaningfully address an organization's supplier or vendor risk, whereas an operationalized SBOM provided by the vendor can. This same lifecycle approach increasingly extends to AI systems, where organizations require comparable transparency into models, datasets, and provenance alongside traditional software components. As regulatory expectations and best practices evolve, closing the gap between SCA and SBOM adoption—particularly within technology organizations—will be critical to maturing software supply chain security programs.

THE MISSING MIDDLE

Verifiable Transparency Data: Bridging from "Generated" to "Operationalized"

Expectations for higher-quality transparency artifacts from software providers are rising, even as many organizations have yet to operationalize them. AppSec professionals are driving this change.

Receiving transparency data does not equate to being able to act on it. As the previous section demonstrated, many organizations remain trapped at the "generation" stage of SBOM maturity. They have not yet built the processes required for centralized intake, normalization, or continuous monitoring. As a result, the transparency artifacts they receive today often function as static documents rather than operational assets.

49.4%

Nearly half of respondents obtain verifiable transparency data—such as SBOMs, provenance records, or signed binaries—from vendors most of the time during procurement.

Industries That Receive the Lowest Rates of Verifiable Transparency Data From Vendors, Despite Needing it the Most

This gap is most visible in highly regulated sectors. Financial services (14.3%) and healthcare (19.5%) report the lowest rates of receiving verifiable transparency data from vendors—despite having the greatest need for it.

Receiving transparency data does not equate to being able to act on it. As the previous section demonstrated, many organizations remain trapped at the "generation" stage of SBOM maturity. They have not yet built the processes required for centralized intake, normalization, or continuous monitoring. As a result, the transparency artifacts they receive today often function as static documents rather than operational assets. It's important to note that transparency is relative: not all SBOMs are created equal, and variations in generator quality can significantly affect risk visibility.

The research also highlights an important shift: when software vendors do provide clear, accurate, and verifiable information about components and dependencies, the downstream effects are substantial. Organizations report faster resolution of security issues and quicker deployment of new technology into production. In these cases, transparency shifts teams from a reactive "wait and see" posture to a proactive "see and act" model This eliminates the guesswork that typically occurs when new vulnerabilities emerge or underlying components change without notice.

Transparency alone is not enough—but when it is operationalized, the impact is measurable.

Transparency is an Investment, Not a Burden.

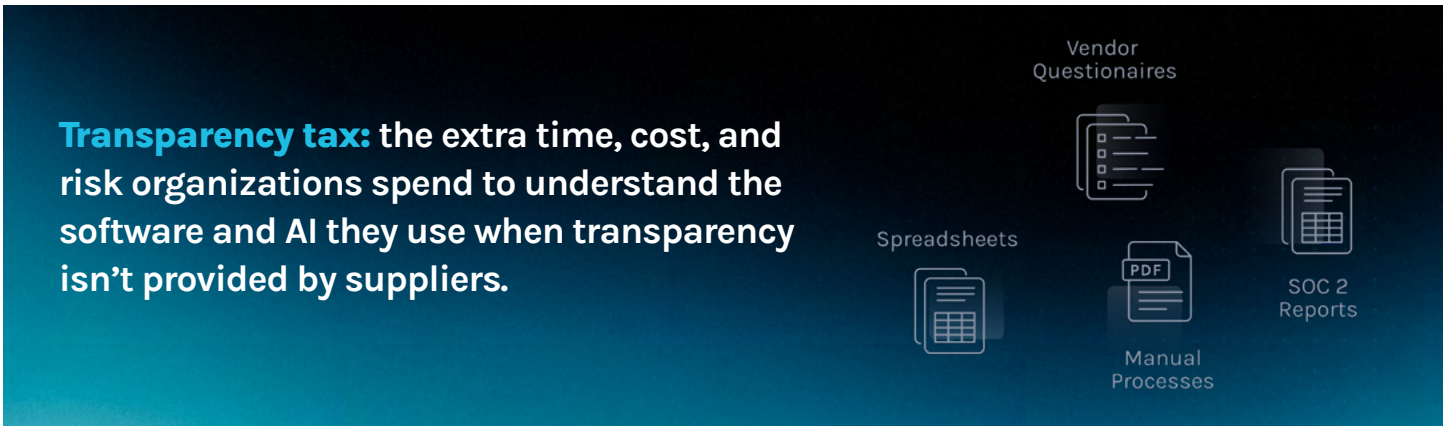

The "transparency tax" is real, but organizations with strong transparency do not pay it.

About two-thirds of organizations now receive verifiable transparency data most or all of the time during procurement—but consistency varies sharply by role and region.

	Receiving Transparency Data	Sharing Transparency Data
AppSec	**46.7%**	**53.3%**
EMEA	**35.2%**	**51.4%**
U.S.	**24.2%**	**29.2%**

AppSec teams report the highest exposure to transparency gaps, with 46.7% consistently receiving data, while 53.3% report having to share or interpret it downstream for others. This highlights a common challenge: transparency often exists, but does not flow cleanly across teams or systems.

Geography and regulation also shape outcomes. EMEA respondents both give and receive less transparency data than U.S. counterparts, suggesting that forthcoming regulations like the EU AI Act and Cyber Resilience Act (CRA) have not yet translated into day-to-day procurement practices. In contrast, regulated industries such as healthcare (19.5% receive) and financial services (14.3% receive) benefit from stricter rules, resulting in more transparency and effectively avoiding much of the "tax."

Transparency is an Investment, Not a Burden. (Cont.)

Transparency drives real customer outcomes.

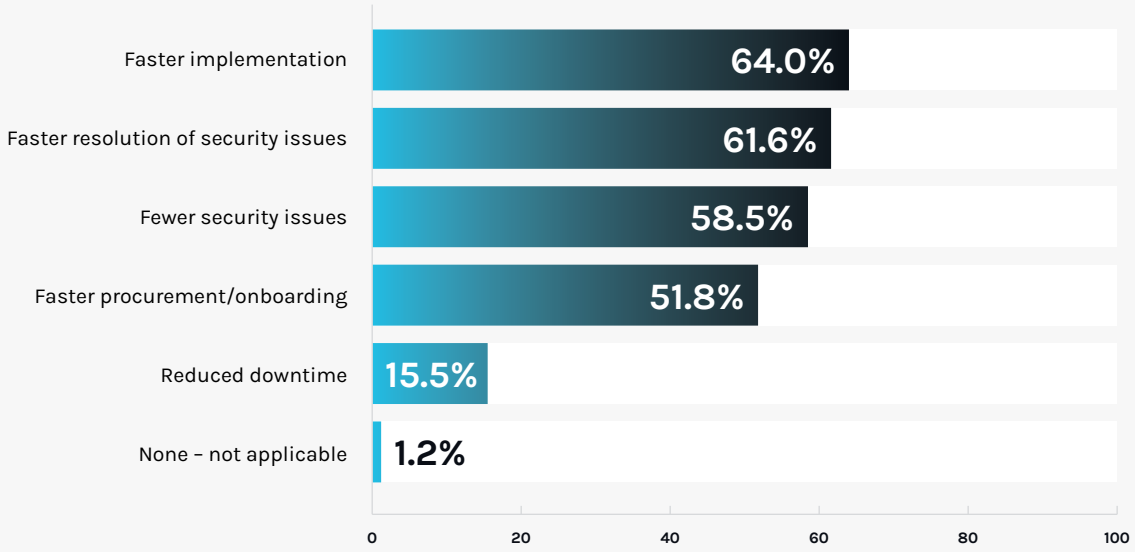

When vendors provide clear, accurate, and verifiable information about software components and dependencies, the benefits are tangible:

- **64%** report faster implementation of new technology.
- **61.6%** report faster resolution of security issues.
- **15.5%** report reduced downtime which is still meaningful given internal constraints like incident workflows and change management.

These gains are most pronounced among AppSec teams (62.2%), large organizations (71.9%), and EMEA respondents (67.6%). This signals that teams with stronger transparency foundations are already operating more efficiently.

Organizations without this visibility remain reactive—slower to respond, slower to deploy, and burdened by manual investigation. Vendors that provide verifiable transparency turn it into a competitive advantage, helping customers reduce risk, accelerate adoption, and operate with confidence.

Organizations that embed transparency move faster, reduce risk, and gain a clearer view of their software and AI supply chains

AI INSIGHTS

AI Adoption Is Outrunning Governance, Licensing, and Legal Risk Management

While our survey found that virtually no organizations report avoiding AI entirely, high adoption doesn't equate to mature governance.

56.7% of respondents are training open-weight models on internal data, and a significant minority (29%) are tuning their own models, which is likely to grow.

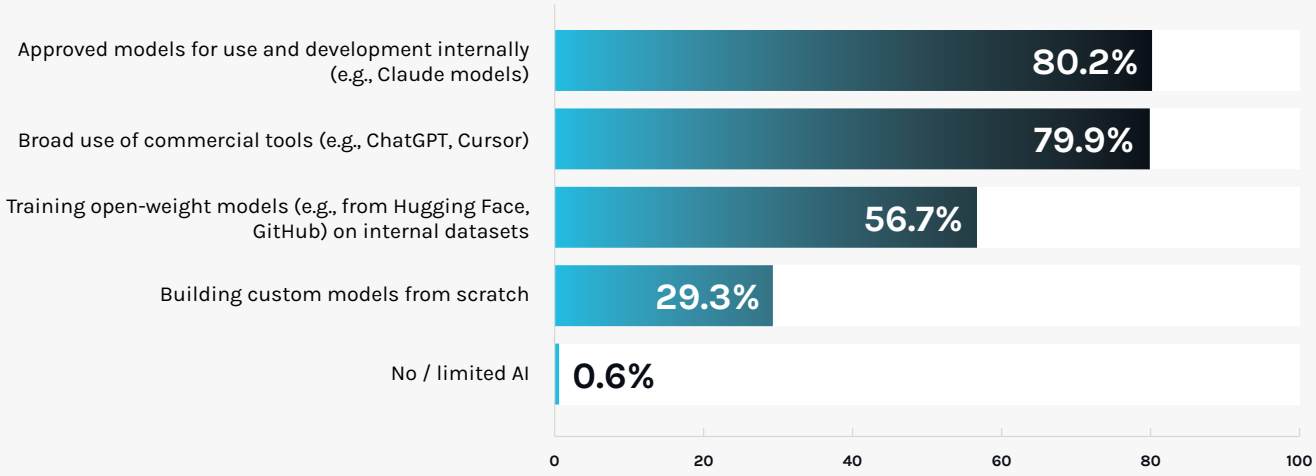

When looking more specifically at which AI models are currently being used, the data suggests that EMEA has significantly lower adoption of open-weight models and custom model development compared to other regions. Instead, organizations in EMEA appear to be leaning more heavily on commercial AI models.

This pattern likely reflects a combination of regulatory expectations, data protection norms, and risk tolerance. EMEA teams appear to favor the assurances, support, and perceived compliance that come with commercial offerings. They are less inclined to take on the operational responsibilities of hosting, tuning, and governing open-weight models themselves—a stance reinforced by the EU's proactive AI governance, including the EU AI Act. What's behind this pattern? Potentially a mix of regulatory expectations, data protection norms, and risk appetite.

Relying on commercial AI has tangible consequences. On the one hand, it provides alignment with vendor-backed controls and service-level agreements. On the other, it may limit flexibility and reduce visibility into the underlying model behavior and licensing, leaving potential gaps in control and insight.

Why it matters:

- Most teams are putting guardrails in place by leveraging familiar software governance frameworks.
- AppSec is an early signal of AI-specific risk awareness, but their separate track highlights potential fragmentation across the organization.
- Without alignment, “normal software” and AI could be governed by different processes, creating gaps in risk coverage.
- You don’t need to reinvent the entire governance function for AI, you need to extend and adapt what you already have. Update governance playbooks to address AI-specific risks like data provenance, model behavior, and continuous retraining, ensuring a unified approach as AI adoption scales.
- Government adoption is advancing faster than anticipated, demonstrating early-stage regulatory readiness and signaling that public-sector organizations are already integrating AI governance practices at scale.

The market is moving in the right direction, treating AI as part of the software and data supply chain, rather than building parallel, duplicative structures that are harder to maintain over time.

Governance Maturity Gap

56.4%

of all teams govern AI like other software, reusing existing review, approval, and risk processes.

42.4%

of all teams treat AI separately, operating outside standard governance structures.

42.2% of AppSec teams govern AI the same, while

57.8%

govern it separately.

Watch the Pace-Setters in AI Adoption

Financial services and technology companies are ahead of the AI adoption curve, operating with agility and strong incentives to move quickly on new capabilities. For technology firms, AI is central to their products and platforms, while in financial services, AI can directly influence revenue, risk management, and competitive advantage. This structural motivation drives both sectors to experiment, iterate, and invest in the engineering and governance needed to operationalize AI at scale.

Survey data confirms this leadership:

89.8% of financial services teams and over **80.5%** of technology teams report formally approved AI models for internal use.

83.7% of financial services organizations and **80.5%** of technology companies use platforms such as ChatGPT or Cursor.

46.9% of financial services teams build models from scratch, followed by technology at **35.9%.**

Financial services organizations are not just adopting AI, they are actively shaping it to meet their risk posture, compliance requirements, and business strategies. Nearly half of respondents in financial services build custom models from scratch, compared with the overall baseline of 29.3%. Years of investment in quantitative analytics, risk modeling, and fraud detection have primed these institutions to own models tailored to their proprietary data and competitive differentiation. Rather than relying primarily on off-the-shelf AI, they are more willing to develop and fine-tune models for highly specific use cases within strict regulatory constraints, potentially setting a benchmark for other industries.

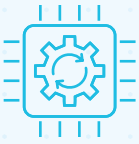

Financial services and technology companies demonstrate that AI adoption is not just about using new tools. It's about embedding AI into core operations and governance structures, turning experimentation into operational capability at scale.

Shadow AI Exposes Gaps in AI Visibility and Governance

Shadow AI is already a known problem, not a hypothetical one. But we can clearly see that traditional controls and inventories aren't keeping up. The majority of respondents report that they've discovered AI (tools, models, or integrations) that were adopted without clear oversight from security, procurement, or risk teams.

This suggests that AI is following a similar trajectory to past waves of "shadow IT," where business units and developers move faster than centralized governance.

Roughly two-thirds of organizations have shadow AI, indicating widespread unmanaged usage.

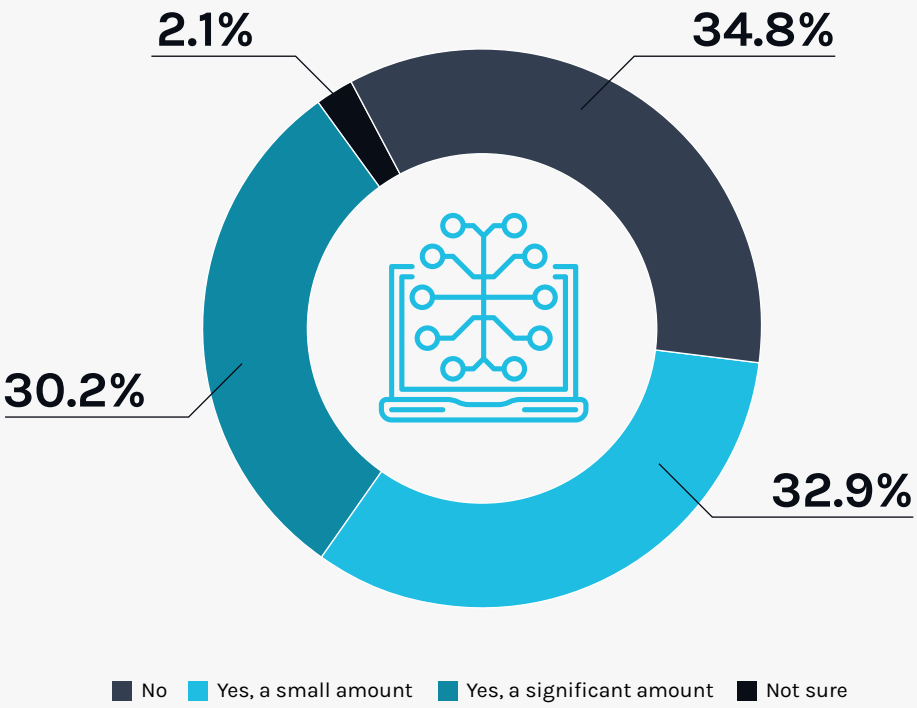

Regional AI Adoption Patterns Shape Shadow AI Risk

In EMEA, the share reporting "no shadow AI" is notably higher at

45.7%

EMEA

- Higher reliance on off-the-shelf AI.
- AI flows through formal procurement, vendor management, and compliance processes.

Other regions

- Greater use of open-weight and custom models.
- More developer-led experimentation.

AI Licensing and Legal Awareness Are Active Pain Points

Survey results reveal widespread challenges around understanding and managing AI licensing, intellectual property, and usage obligations. Most respondents recognize significant room for improvement, with 93% agreeing and more than half (54.6%) agreeing strongly that their organization has significant room for improvement when it comes to understanding and managing the legal usage, licensing, and intellectual property obligations associated with AI components. Financial services respondents were nearly unanimous (98%).

Gaps Have Consequences

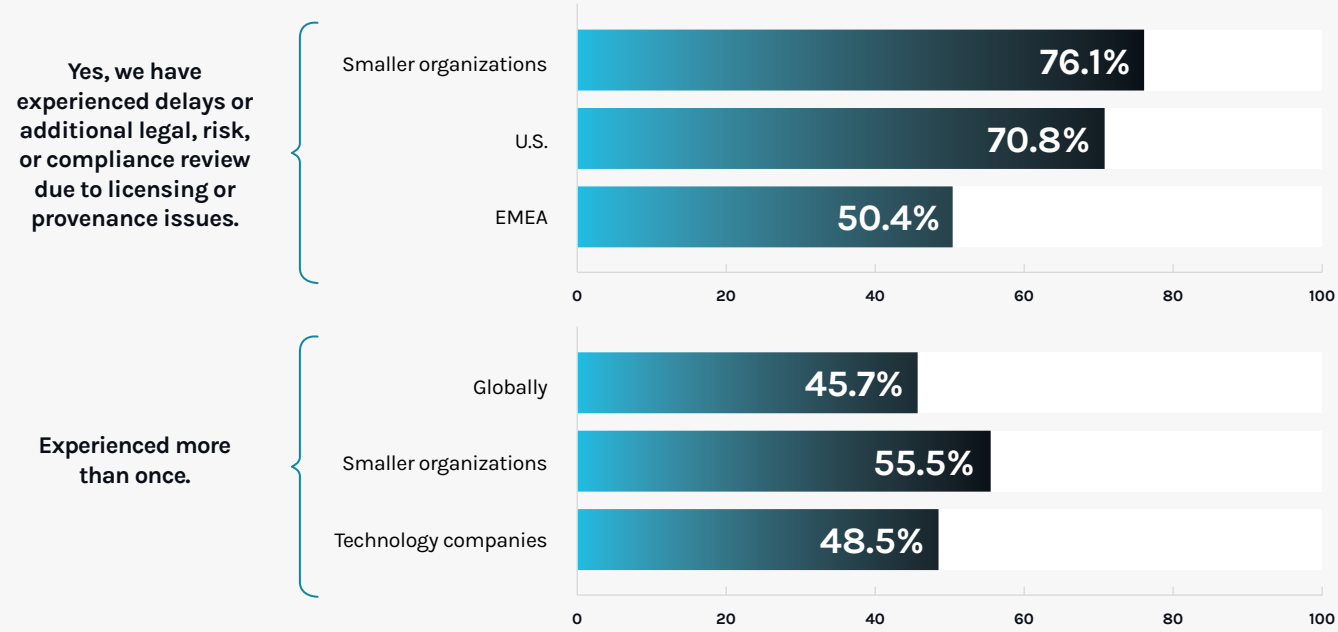

Usage Patterns

56.7% of respondents report training open-weight models on internal data. U.S. teams and smaller organizations are heavier users of open-weight or custom models compared with EMEA, which leans more on off-the-shelf AI. Taken together with the licensing gaps, this raises a concern: many organizations may be experimenting with or operationalizing open-weight models without fully understanding license terms, usage restrictions, or attribution requirements.

Blind Spot

The majority of organizations could be inadvertently introducing legal, compliance, and business risk. There is a clear opportunity for better guardrails, tooling, and processes that connect AI model selection and training decisions with robust license analysis and governance, helping organizations adopt AI confidently and responsibly.

How Strong is the AI Security Program? The Answer Depends on Who You Ask...

Executives consistently rate their security programs highly, but the view from AppSec and Product Security tells a more fragmented story. While overall responses suggest **72.9% of organizations consider their AI security programs complete or thorough,** that number is heavily influenced by executives.

AppSec's View of AI Security Programs

53.3%
complete/thorough

40.0%
haphazard

4.4%
not at all complete/thorough

AppSec's view of AI Governance

42.2%
Governed the same as software

57.8%
Governed separately from software

This discrepancy highlights a tension between perception and reality. Some teams may operate effectively while others rely on ad hoc or manual processes.

Even organizations with mature AppSec practices face gaps in AI visibility.
Traditional SCA and APSM tools were not designed to capture all AI components, models, and services in use, leaving teams confident, yet disorganized. This is particularly evident in large organizations, where more products, teams, vendors, and shadow AI usage make comprehensive tracking difficult—despite leaders feeling "good enough." Similarly, financial services report both high adoption of off-the-shelf AI solutions and the highest proportion of haphazard processes (38.8%), illustrating how rapid AI adoption can outpace governance and inventory efforts.

Government AI: Fast Adoption, Uneven Visibility

Government organizations show a nearly even split between thorough and haphazard AI inventory processes, even as adoption accelerates.

- Increasing use of off-the-shelf AI.
- Growing reliance on open-weight models trained on internal data.
- Inconsistent visibility into AI components across environments.

Fast adoption without consistent visibility creates a potential risk zone. Foundational practices, like unified AI component inventory and SBOM-like transparency, are critical to maintaining control.

Understanding the Greatest Obstacles to Software Supply Chain Security

There's an obvious opportunity to align AppSec, Product Security, Legal, and Compliance on shared tooling and records. By doing so, it provides a single source of truth, making it easier to spot gaps, respond quickly, and advance AI security programs.

Top Constraints

47.3% organizational constraints

36.3% insufficient skills

35.7% budget limits

34.8% management understanding

32.6% staffing shortages

Perceived vs. Actual Security

Organizations report high confidence in their security programs, especially AppSec teams, who rate internal software and AI components higher than the overall averages. Externally sourced software is slightly lower, with mid-sized orgs struggling most, and EMEA respondents generally more confident than the U.S.

Rate your organization's capabilities and progress for ensuring security and reducing risk

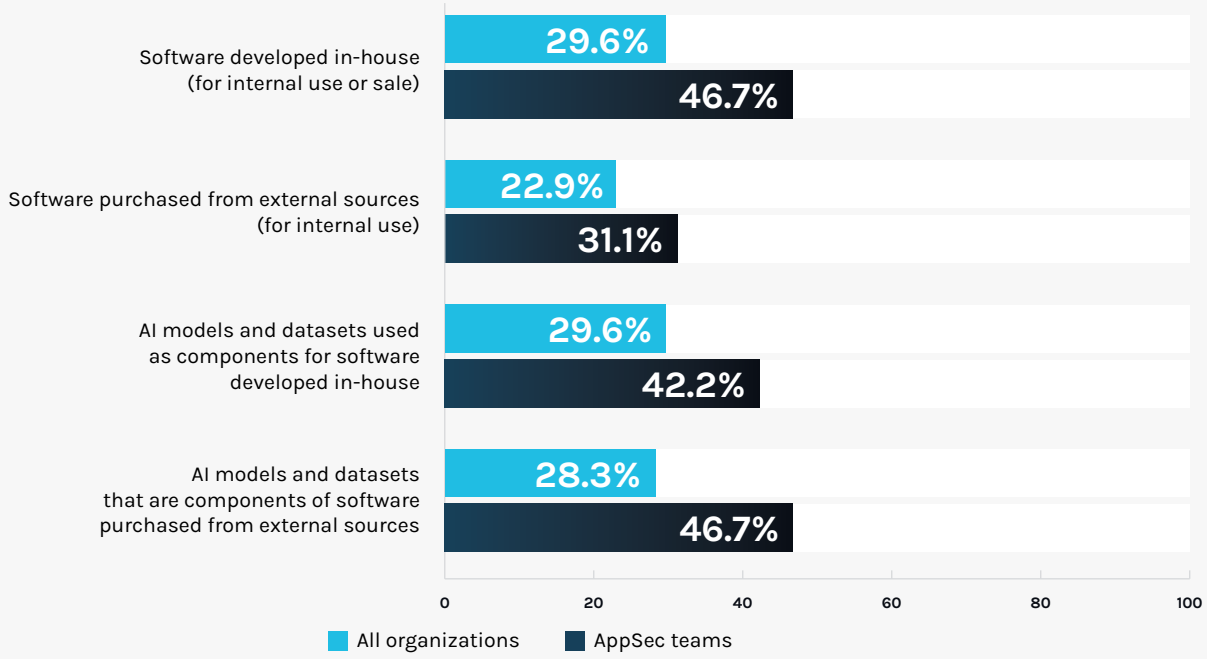

This optimism may mask gaps. Executives could be focusing on visible AI features—chatbots, copilots, recommendation engines—rather than the underlying models embedded in software. The result is a potential disconnect between perceived and actual visibility into AI, highlighting the need for SBOM-like approaches and unified AI inventories to ensure governance keeps pace with adoption.

Bringing Visibility to AI

As AI adoption accelerates, so do the risks; spanning outdated models, blackbox training data, non-compliant licenses, and geopolitical concerns. AI risk encompasses threats from integrating AI into software, including high-risk licenses, undocumented datasets, restricted or unmaintained models, shadow AI, and third-party dependencies. Ultimately, these mirror the longstanding challenges of the software supply chain.

Securing the Supply Chain. What's Needed:

Configure Policies: Ensure AI model development aligns with your organization's governance policies.

Explore AI Risk: Evaluate open-weight models from Hugging Face for risks related to undocumented training data, alignment with internal policies, and potential licensing or usage restrictions.

Track and Inventory Models: Keep track of all AI models in use, open-weight or custom tuned.

Enforce Policies: Detect models in source code and python notebooks, and trigger alerts and enforcement actions for developers and security teams for models that violate internal policies.

Scan Source Code: Detect AI models embedded in source code by scanning the CI/CD pipeline and GitHub repositories.

About Manifest

Manifest empowers private and public sector organizations to operate critical systems and applications with confidence. We detect and manage hidden software supply chain and AI risks at scale. The Manifest Platform provides end-to-end visibility and control across Product Security, AI Risk, and Supplier Risk, helping teams build secure, trusted software without losing velocity. Trusted across defense, healthcare, automotive, and other regulated industries to enhance product & AI security, third-party risk, and compliance.

Learn more →

MANIFEST

Survey methodology

Our survey was conducted online in November 2025 and included 328 respondents across the United States, EMEA, France, and Germany. Participants represented a broad cross-section of security roles. The majority—68%—were CIOs, CISOs, or other senior IT security executives. Another 13.7% were application security managers or practitioners, 6.4% worked in product security, and 8.5% held AI security or governance leadership roles. A smaller share (2.7%) worked in other IT security functions, while 0.6% represented third-party risk management.

Respondents came from organizations of all sizes, ranging from 500 employees to large enterprises with more than 25,000 employees. Nearly one-third (29.3%) worked at organizations with 1,000–4,999 employees, and another 26.5% came from companies with 5,000–9,999 employees. About 19.2% worked in mid-sized companies with 500–999 employees, 15.2% were from organizations with 10,000–25,000 employees, and 9.8% represented the largest global enterprises.

Participants spanned a wide set of industries, including defense and aerospace, financial services and FinTech, government and public sector, healthcare and medical device manufacturing, automotive and other manufacturing sectors, retail and ecommerce, and technology/software, along with additional respondents from other industries.