Delinea

Delinea AI Report

AI in Identity Security Demands a New Playbook

Artificial intelligence (AI) is reshaping how business gets done. That includes IT and security, where new agentic AI use cases are blossoming. AI in cybersecurity and IT operations offers tremendous benefits, but it also introduces commensurate risk to the enterprise if organizations are not able to set and enforce AI governance policies aligned with their risk appetite.

As organizations increasingly deploy AI, they must begin treating those entities with the same level of scrutiny and governance as privileged employees, contractors, or partners. This shift demands significant changes in how identity and access management (IAM) is architected.

To understand where the industry is today, with agentic AI vs generative AI deployments and risk management, Delinea recently commissioned a global survey of 1,758 IT decision-makers. This includes stakeholders from the U.S., UK, UAE, Australia, Singapore, and Germany.

The study shows that AI in IT operations is growing:

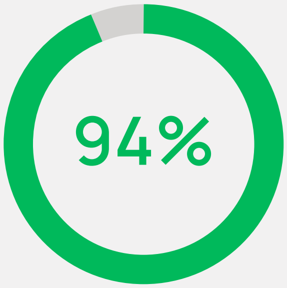

of global companies use or pilot some form of AI in IT operations

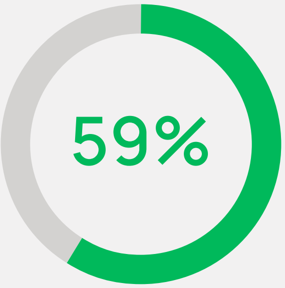

of organizations actively use both generative AI and agentic AI in IT operations

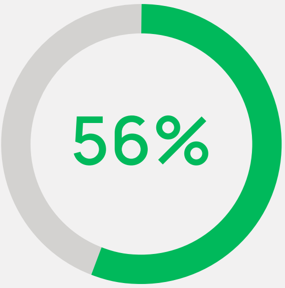

of organizations run into shadow AI issues at least monthly

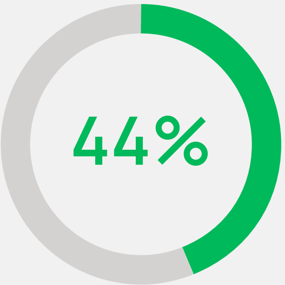

of organizations' security architecture is fully equipped to support secure AI today

Let's explore how organizations are using AI in security operations, how they're grappling with the risks of AI, and how AI can help mitigate identity security risks in the AI era.

Large firms need AI to keep up with speed of IT Ops

Most IT organizations are already betting big on AI to improve IT operations. A full **84%** are already actively using some form of AI in IT operations—whether generative AI, agentic AI, or some combination of the two. A further **10%** of firms are piloting or evaluating AI tools for use in their IT departments.

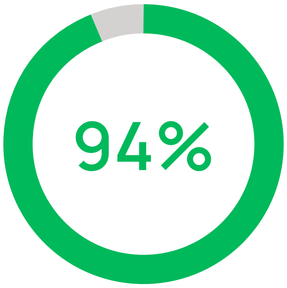

of global companies use or are using or piloting some form of AI in IT operations.

Both agentic AI and generative AI offer significant potential to streamline operations and improve efficiency in organizations of all sizes. These AI tools can help IT operators speed up processes, quickly generate code and summative reports based on prolific log data, create better documentation, and establish more intelligently automated workflows.

Larger companies are far more likely to already actively use AI in IT operations. As IT operations scale up in volume of work—more tickets, more systems managed, more development needed—they're turning to AI agents to help them keep up with the barrage. The need to keep up with content and automated action is greater for these firms, as are the resources to adopt and experiment with emerging tech. Our study showed that, while **89%** of companies with more than 500 employees are actively using AI, only **52%** of companies with under ten employees are doing so.

> **Agentic AI serves an important role in the automation and productivity programs across our IT organization. It's also very important to our broader enterprise use cases given the volume and diversity of interactions with customers and consumers.**
>
> – CISO, Healthcare Organization

However, the upsides of these new technologies also bring significant emerging risks that demand a rethink of traditional policy controls. Many companies have not yet been able to fully rearchitect their approach to modernize for the agentic AI era. Even as they're deploying these technologies headlong, their policies and controls are still stuck in the past.

Most regions closely tracked with the global response rates for active AI usage. Singapore leads in deployments, with 90% of respondents already using AI in IT operations.

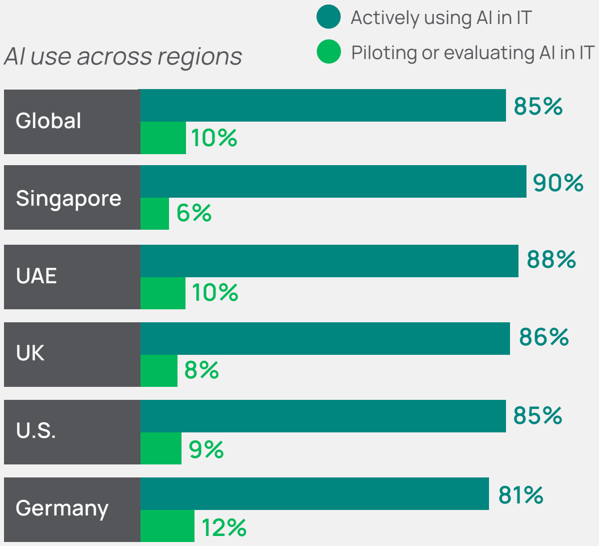

Organizations are grappling with AI risks

The vast majority of organizations are already using AI for IT operations and AI usage is proliferating across the business. But the hard reality is that their security architecture is not yet ready to support secure AI operations.

Less than half of organizations **(44%)** report that their security architecture is fully equipped to support secure AI. Meanwhile, **47%** of organizations are hopeful that their security roadmap will bolster the security of AI deployments within two years.

Two years, however, is a long time in the face of a rapidly evolving AI landscape that's introducing a spate of new cybersecurity and business risks by the day. Organizations across all sectors are already facing a surge in AI-enabled cyber threats, according to the Delinea 2025 State of Ransomware Report. Adversaries increasingly weaponize generative AI for phishing, deepfakes, and sophisticated social engineering attacks.

Delinea reports that organizations are responding in kind, with **90%** of corporate IT security teams leveraging AI to keep up with bad actors.

AI increases the attack surface of enterprise systems, posing an even greater risk to businesses. It also introduces many security assurance risks as new failure points are introduced into IT architecture.

When it comes to identity risk introduced by AI, respondents' most common concern was malicious actors' use of AI. Some **47%** of global firms reported that AI-generated phishing/deepfakes were their top concern, and **44%** reported that AI-driven credential theft is one of their biggest concerns.

Additional concerns included agentic AI systems with unchecked access, unsanctioned use of AI tools, and poor visibility into AI access workflows.

Top AI security concerns

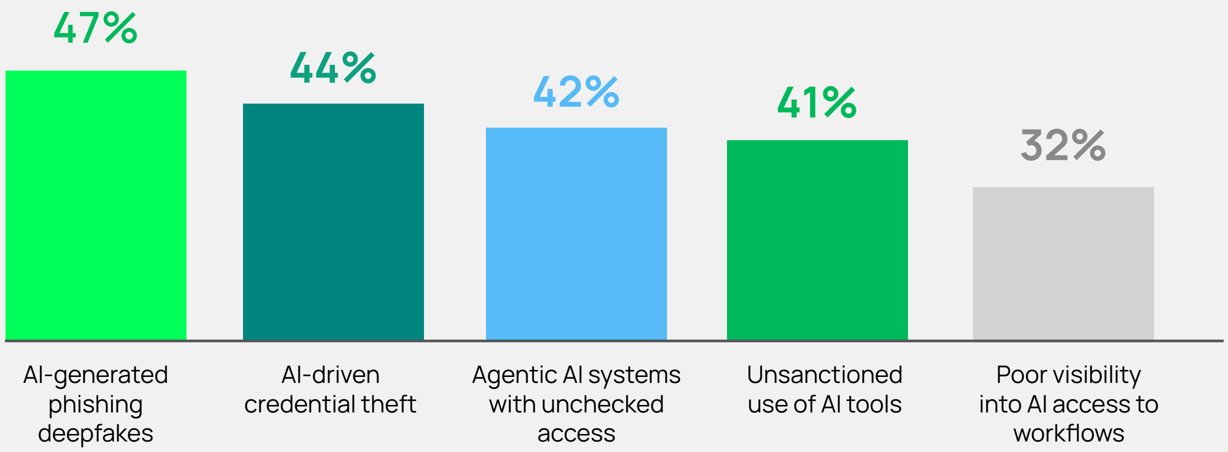

Generative AI and agentic AI bring different risks

Top identity security risk posed by AI, by industry

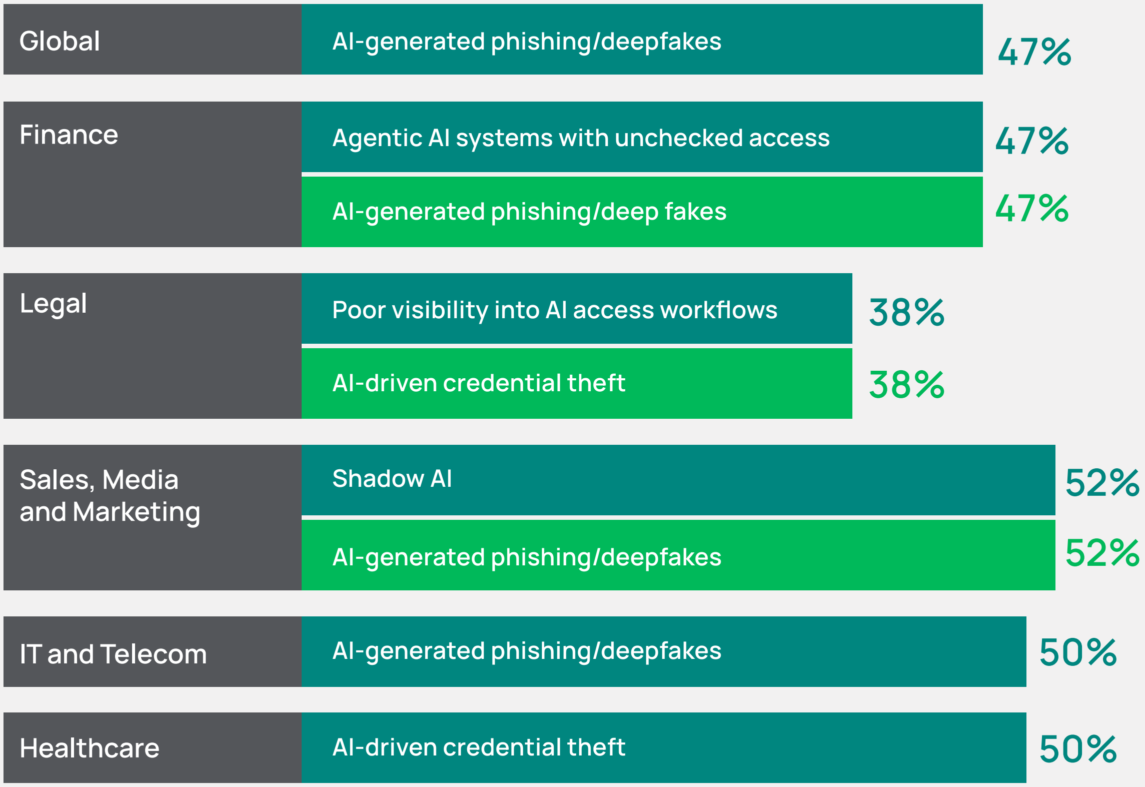

While there are many use cases that deploy a combination of generative AI and agentic AI, decision-makers should understand that each one is a distinct type of artificial intelligence with its own strengths, limitations, and risks. Understanding these differences can help AI policymakers and security architects plan for the right mix of policies and controls that can help their organizations get the most out of use cases in a risk-managed way.

Generative AI is the more prevalent of the two, with **79%** of organizations reporting they use it compared to **66%** who use agentic AI. **60%** actively use both types of technology.

It's important to know the difference between agentic AI vs generative AI to understand the unique risks of each.

Generative AI risks

Generative AI creates content in response to user input. Generative AI technologies are prompt-driven, reacting to what the "human in the loop" tells them to produce. In the context of IT use cases, this could include code generation, production of new knowledge base documents, data analysis, or generation of report summaries based on data. These technologies help speed up human-led activities, but on their own, they don't act independently.

From a business benefit perspective, the limited autonomy of generative technology can be seen as a weakness. The upside is that the risks are much more contained in scope. Generative AI risks usually revolve around data and privacy-related issues. This includes risk of data exposure or theft, errors generated in content due to malicious or flawed input, and potential privacy violations or intellectual property (IP) issues.

Agentic AI risks

Agentic AI is designed to act autonomously to achieve pre-determined goals. AI agents automatically use data to make decisions, manage workflows, and orchestrate multistep tasks. In IT use cases, agentic AI could be used to autonomously triage and route tickets as well as analyze system logs and events for problems and automatically change environments.

In some instances, agentic AI might use generative AI to create content or analysis that feeds a complex chain of end-to-end automation managed by the agent. For example, an IT ops assistive agent designed to help optimize Infrastructure as Code (IaC) may use generative AI to create code for environments based on certain parameters. This then feeds a broader automation of that code deployment. An ongoing analysis of environment status will call for more code generation as the status and business conditions change.

Agentic AI holds tremendous business potential because of its self-sufficiency and scalability. When we asked businesses already actively using agentic AI what the major drivers were for their deployments, the top three answers were increased operational efficiency (**63%**), improved incident response time (**45%**), and cost reduction (**42%**).

The problem with agentic AI—especially in IT use cases—is that it also significantly ups the ante on risk. This is because it typically brings with it all the same data-related risks of generative AI but adds the force multiplier of autonomy to the mix. AI agents are executing actions and decisions at scale without human intervention. This means that unchecked autonomous decisions based on failed data or flawed AI logic could cause unintended harm.

The lack of human oversight, the privileged status of AI agents, and the automation powered by them make these agentic AI systems very juicy targets for threat actors to subvert. These will inevitably be the breeding ground for a host of new and newly reinvigorated cyberattacks to compromise credentials, compromise or sabotage systems, steal data, and automatically initiate fraudulent transactions.

This is why security teams need to rethink their identity security playbook to account for agentic AI.

Emerging AI risk cases

The risk scenarios for both generative AI and agentic AI are still unfolding, but some of the most troubling ones that have come to light through security research and real-world incidents include:

Overprivileged AI agents

Agentic AI with excessive privilege and access to sensitive systems can expose organizations to the visibility and extraction of sensitive data, compliance violations, and broader security breaches.

Real-world example: Security researchers were able to breach McDonald's hiring chatbot, "Olivia." This system had administrative access to 64 million job applicants, and it was protected with the password "123456."

Extreme data exposure

The McDonald's example highlights a broader class of risks posed by both generative and agentic AI. These systems are often given access to tremendous troves of data, and many of them are very sensitive. In some cases, the AI systems themselves are experimental and may not be well-vetted, which is a recipe for flaws and exploitation.

Real-world example: Productivity software maker Asana introduced an experimental model context protocol (MCP) server that was designed to help AI agents query workflow data. Asana discovered that a single line of code introduced a logic flaw that exposed some of the most sensitive types of data for 1,000 organizations, including strategic roadmaps, merger and acquisition discussions, financial data, and customer information.

Automation gone rogue

Whether it's triggered by malicious actions or AI incorrectly interpreting data, agentic AI "going rogue" and autonomously making undesired changes or causing outages is a very real security assurance concern. The more permissions an agentic AI system has and the broader its scope of authority, the more dire the consequences could be.

Real-world example: An incident involving AI coding tool Replit shines a bright light on this risk. The firm's tools made headlines when its AI agent deleted a company's entire production code base during a test run.

Goal hijacking

Adversarial input that targets agentic AI workflows and corrupts AI decision-making paths can help attackers subvert the goals of the agent and have it do an attacker's bidding. This could be used to enable remote code execution, compromise systems, steal data, or even initiate fraudulent financial transactions. Goal hijacking could be achieved by prompt injection, exploiting unsecured prompts, or otherwise exploiting code or logic vulnerabilities in the system.

Real-world example: Princeton University showed how they could trigger malicious behavior from AI agents, such as Mastercard's Agent Pay and PayPal's Agent Toolkit, by implanting fake "memories" into the data they use to make decisions.

AI policies struggle to keep up with shadow AI risks

Woven throughout these AI risk scenarios is the reality that many generative AI and agentic AI deployments still lack the kind of transparency and visibility demanded of traditional automated IT systems. This further exacerbates risk when it comes to auditing the actions of these systems and identifying underlying security issues. These issues grow even more dire when considering the unmanaged shadow AI systems that are increasingly popping up at organizations today.

Shadow AI is the unsanctioned use of AI tools by employees or departments without oversight from IT or security teams. Many organizations with at least some AI usage are struggling to wrap their arms around shadow AI risks. The most commonly cited issues with shadow AI are business units deploying AI solutions without involving IT/security (**44%**) and unauthorized usage of generative AI by employees (**44%**). This is particularly a problem for IT and Telecom, with **59%** of respondents from these industries identifying it as an issue.

The majority of organizations today have implemented at least some form of policies or controls to restrict or monitor AI tool access to sensitive data. Globally, that number stands at **89%**. However, the comprehensiveness of these controls is less universal. Only about half (**52%**) of global organizations say their controls are comprehensive. That number falls dramatically as the organizational size grows smaller. Just **30%** of companies with under 50 employees say they have comprehensive policies and controls in place.

AI controls are lacking

Digging further into the governance and visibility measures most commonly used, our data shows that many companies still have a ways to go before they're able to establish full visibility into their environments. For example, the most common measure is an acceptable use policy for AI tools. This should be table stakes for any organization using AI, but only just over half of organizations have one, indicating that many organizations today are flying blind with regard to AI activity in their digital ecosystems.

AI controls in use

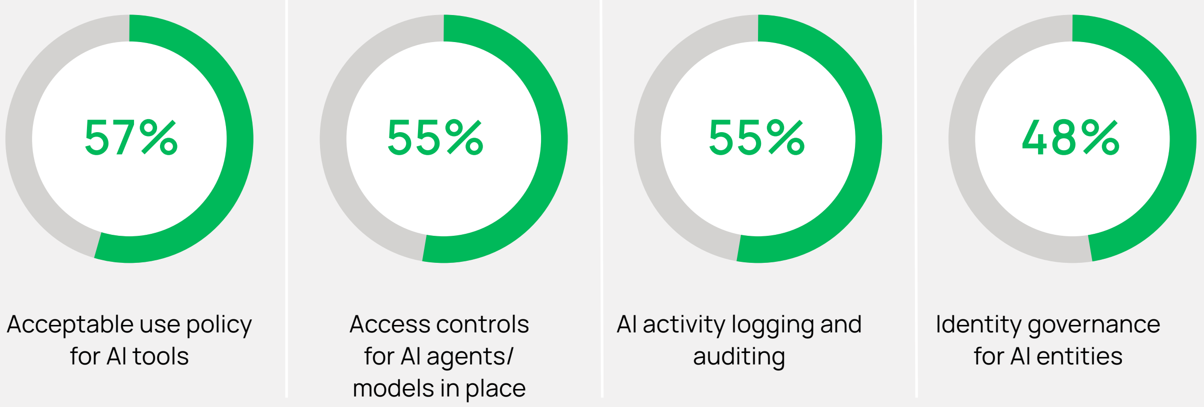

This risk from lack of controls is compounded by the growing prevalence of shadow AI. Just over half of firms surveyed (**56%**) report that they're running into shadow AI issues at least once a month. For a good third of respondents, it occurs multiple times per month.

Shadow AI frequency of occurrence

Frequency of identifying AI tools or agents deployed without IT/security team approval

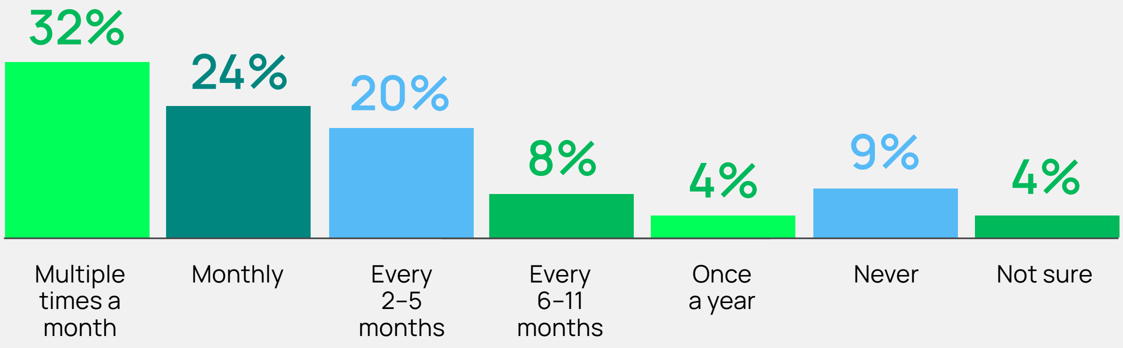

Organizations' confidence in machine identity management may be misplaced

Whether the deployments are shadow AI or fully authorized, organizations will need to up their game in identity security to effectively manage risk in the AI era. Agentic AI adds a whole new dimension to the risk of machine identities, such as large language models (LLMs), as they are given more autonomy and independence to impact critical systems and data. The research shows that organizations feel confident, but their environment tells a different story.

The vast majority of organizations (**93%**) are confident that their machine identity security efforts are keeping pace with emerging threats, like AI manipulation.

> **"Identity is very important in the era of GenAI and Agentic AI. The AI needs to only operate where the identity can be confirmed and validated. This should be paired with audit logs that can track AI activity, which can be ingested in a security platform for detection and monitoring of AI risks."**
>
> – CISO, Healthcare Organization

How firms manage machine identities

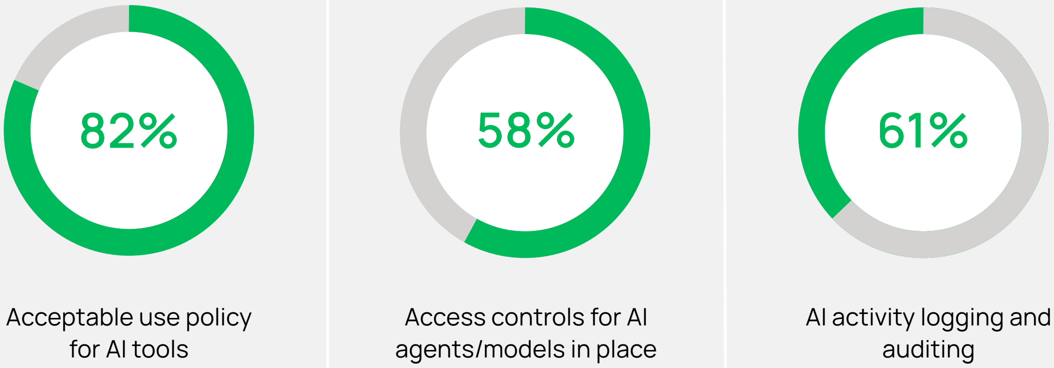

This confidence, however, may be misplaced, given the limitations in AI visibility at many organizations. To meet agentic AI risks head-on—particularly in critical IT systems—organizations must adapt their identity strategies to account for agentic AI risks.

Agentic AI demands agentic security

This growing disconnect reveals a widening governance gap that organizations can't afford to ignore. Agentic AI agents will require more granular and dynamic access controls than traditional role-based access control. New control mechanisms are going to stand as the linchpin to secure AI architectures in the future. Here are some examples:

> **“Agentic AI demands agentic security. Organizations must rethink identity, building adaptive, risk-aware systems that treat every action, whether human or machine-driven, as a decision point that needs to be verified, validated, and secured.”**
>
> Art Gilliland, CEO of Delinea

▶ AI-to-AI credential brokering

AI-to-AI credential brokering uses AI to help automate the machine-to-machine exchange and the verification of digital credentials of all AI agents communicating or acting on behalf of the organization. A smarter automated credential exchange between AI entities—using token-based credentials or digital certificates—will be key for authenticating and authorizing agents.

Additionally, authorization and permissioning will need to be thought through and governed carefully. The more tightly scoped that organizations can make agentic AI tasks or actions, the more they can rein in the consequences when threats emerge.

▶ Visual digital identity mapping

Governance demands will also push cybersecurity and IT leaders to think about how they can make authorization relationships understandable and auditable in the agentic AI age. Organizations will need to find ways to achieve a visual mapping of digital identities, including AI personas, agent IDs, training model metadata, and so on. As well as ways to identify how they differ from human identities.

▶ Strengthen PAM

Organizations will need to strengthen their privileged access management (PAM) models to ensure they can get the level of monitoring and control they need to detect anomalies, like privilege abuse or unusual access patterns, that could indicate agent compromise or failure. For sensitive or high-impact operations, organizations should be able to leverage PAM to require real-time human approval before an agent can proceed with high-risk behavior.

5 steps companies can take today to manage agentic AI risks

1. **Discover and classify AI identities:** Use automated tools to inventory AI agents—scripts, bots, and autonomous models across hybrid and multi-cloud environments. Classify these by sensitivity, privileges, and business impact to align with intelligent privilege management.

2. **Define roles and guardrails:** Set clear operational boundaries for each AI identity classification. Use policy-based access to tie privileges to specific tasks, keeping actions aligned with business intent and risk tolerance.

3. **Enforce least privilege, just-in-time access:** Replace standing privileges with just-in-time access. Grant AI agents only what's needed, when it's needed—then revoke it automatically to reduce risk.

4. **Authenticate and authorize by intent:** Require strong, verifiable identities for AI-to-system and AI-to-AI interactions. Go beyond identity to validate intent, ensuring actions match approved use cases.

5. **Monitor, detect, and continuously improve:** Continuously monitor AI behavior to detect anomalies and misuse. Log actions with cryptographic integrity, enforce encryption, and regularly test workflows to harden identity and access controls.

Leading firms use AI in security operations and identity security

The good news is that, while AI introduces significantly more risks into IT infrastructure and adds a whole new dimension of work in identity governance and access controls, AI can also be used as a force multiplier for security effectiveness.

An overwhelming **91%** of global organizations say that they're either already using, piloting, or planning a deployment in the next 12 months to integrate AI to improve security operations. Encouragingly, a solid 40% of organizations say they're already using agentic AI to improve security operations, and another **24%** say they're in the pilot phase of such a deployment.

For organizations actively using or considering integrating AI into their security solutions, **44%** use AI to fully automate, provision, monitor, and govern all identities. Approximately **63%** use AI to govern some or all identities, and **21%** use AI in an advisory role.

AI is key to governing AI. The top two use cases for AI in security solutions were AI governance and access audit of AI agents (**54%**), and behavioral analytics for anomaly detection (**52%**), which will be essential for early detection of malicious subversion of AI agents and agentic AI failures.

Next steps for stronger identity security in the AI era

Agentic AI demands a new playbook for identity security and cyber risk management. Governance and guardrails are the answer to many of the challenges identified by our survey data. Organizations need to take a programmatic approach to AI risk management that enables human-centric AI governance. We recommend a three-phase crawl, walk, run approach to achieving meaningful AI governance.

PHASE 1: **Foundational**

- Form an AI governance committee and establish initial policies.
- Identify and catalog all generative and agentic AI tools in use.
- Implement baseline access controls, PAM, and logging mechanisms.
- Provide foundational AI security training to key teams.

PHASE 2: **Strengthen controls**

- Expand access management to enforce least privilege and role-based access.
- Deploy advanced identity threat detection and response tools (ITDR) for GenAI environments.
- Enhance supply chain security through continuous vendor assessments.
- Automate compliance monitoring and data protection measures.

PHASE 3: **Full maturity**

- Adopt continuous adaptive risk assessment processes.
- Leverage AI-driven analytics for anomaly detection and real-time response.
- Regularly refine governance policies to align with industry standards.
- Foster a security-first culture with ongoing training and incident simulations.

Discover how Delinea can help you rethink your identity security in the age of AI.

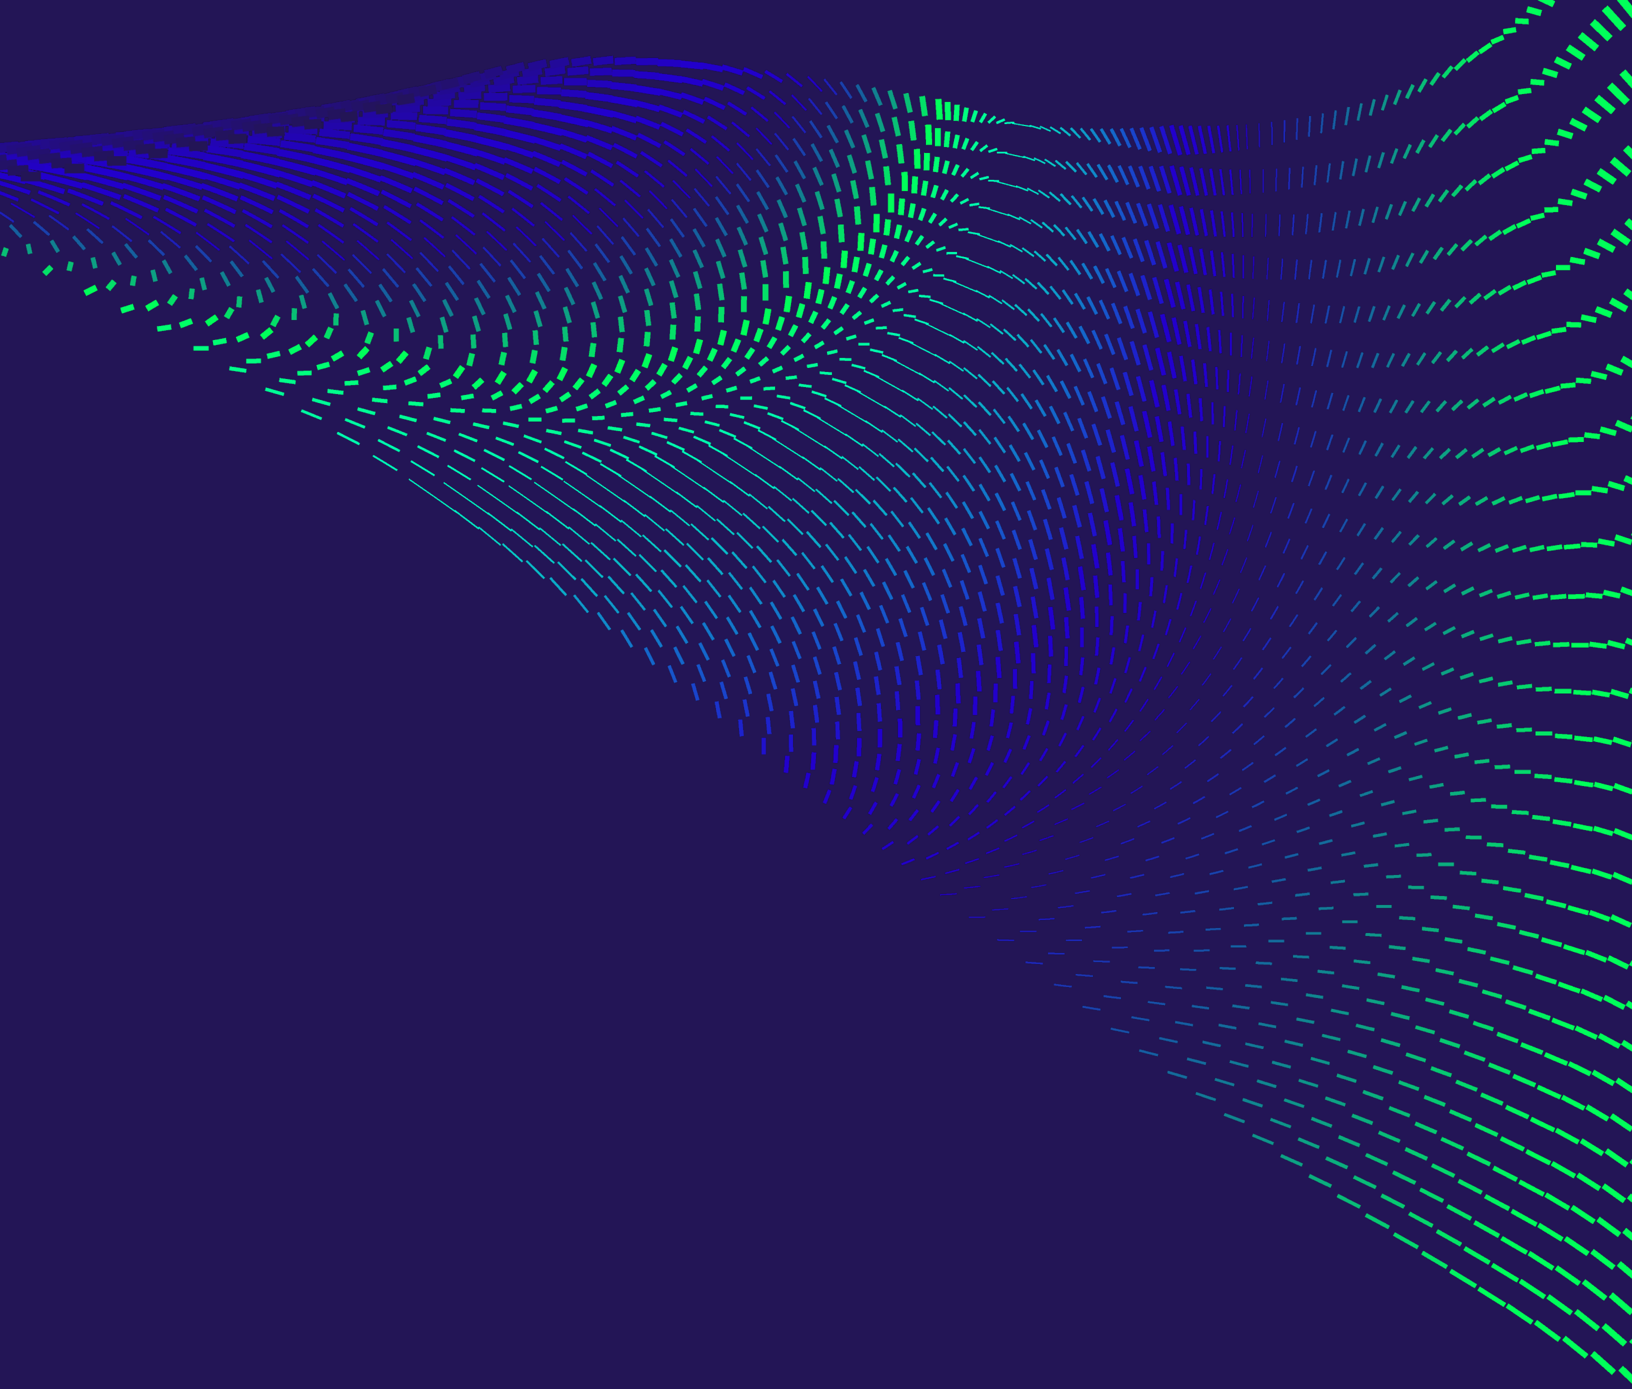

Delinea

Securing identities at every interaction

Delinea is a pioneer in securing human and machine identities through intelligent, centralized authorization, empowering organizations to seamlessly govern their interactions across the modern enterprise. Leveraging AI-powered intelligence, Delinea's leading cloud-native Identity Security Platform applies context throughout the entire identity lifecycle across cloud and traditional infrastructure, data, SaaS applications, and AI. It is the only platform that enables you to discover all identities — including workforce, IT administrator, developers, and machines — assign appropriate access levels, detect irregularities, and respond to threats in real-time. With deployment in weeks, not months, 90% fewer resources to manage than the nearest competitor, and a 99.995% uptime, the Delinea Platform delivers robust security and operational efficiency without complexity. Learn more about Delinea on **Delinea.com**, **LinkedIn**, **X**, and **YouTube**.