digicert®

AI Trust Outlook

AI is scaling faster than trust

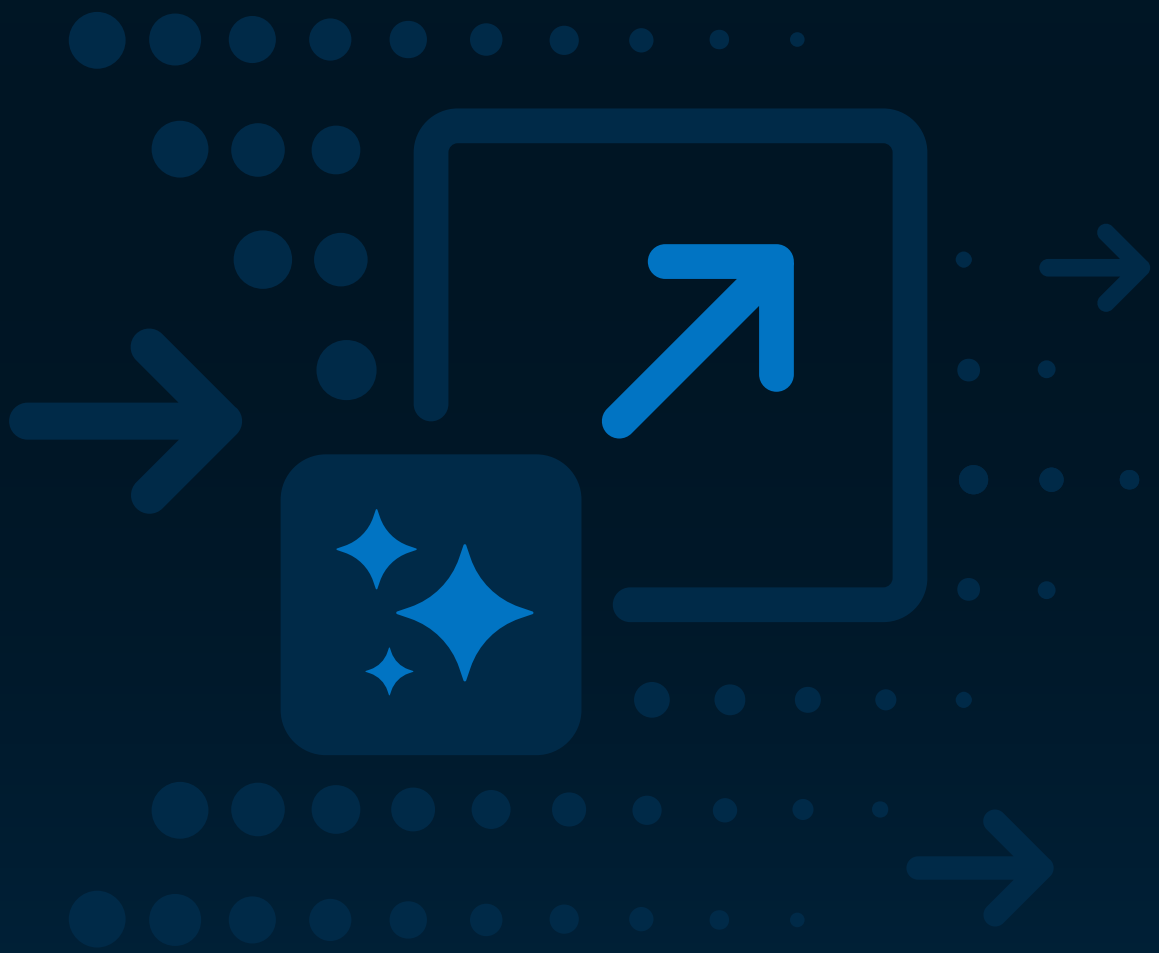

Research Report | 2026

78%

of organizations reported experiencing AI-related security incidents or identifying AI-related vulnerabilities.

AI deployment outpaces AI accountability

The AI conversation has changed. There's no longer a question about whether organizations will use it. They've made that decision. The new question is whether AI can be governed, secured, and trusted at scale.

New research from DigiCert shows that AI has moved well beyond experimentation. Most survey respondents deployed multiple AI-powered systems in just the past six months, and many executives now view AI as a strategic business capability. Meanwhile, AI risk is no longer theoretical: 78% of organizations reported experiencing AI-related security incidents or identifying AI-related vulnerabilities.

With AI embedded in so many everyday operations, it introduces a growing network of non-human actors that require identity, authentication, oversight, and accountability. Organizations are increasingly grappling with how to:

The findings point to a new reality: AI security is becoming an identity and trust challenge. The same principles that secure users, devices, applications, and machine identities are now essential for securing AI.

Read on to discover how more than 1,001 global IT and security leaders are approaching AI adoption, governance, accountability, and the risks shaping the future of trusted AI.

AI security incidents are already here.

Key findings

78%
experienced AI-related incidents or identified AI-related vulnerabilities

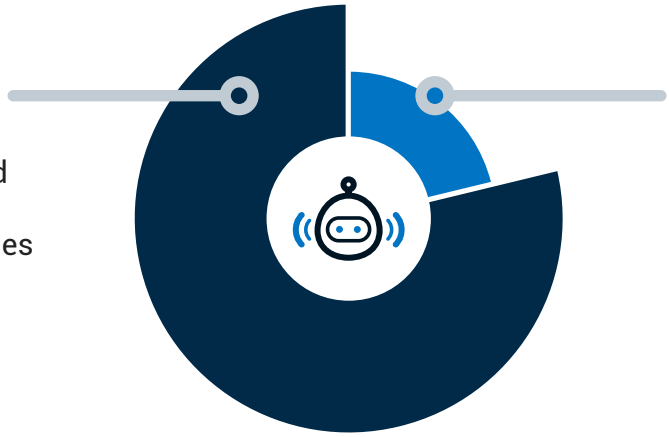

21%
reported no AI-related incidents

Nearly eight in ten respondents have already encountered AI-related security issues. That's a signal that AI has entered a new phase. The technology is creating value, but it's also creating operational risk.

Every AI deployment introduces new identities, credentials, APIs, and machine-to-machine interactions. Most teams still struggle to inventory and govern all of them.

Without visibility, security becomes reactive. Without accountability, trust becomes hard to establish.

The lesson is simple: Securing AI requires securing the identities, credentials, and connections that make AI possible.

Every bot needs a badge.

Key findings

Nearly half of organizations have assigned unique digital identities to all AI agents

Four in ten have assigned identities to at least some AI agents

Nearly **90%** have begun implementing AI identity practices

AI agents are starting to act less like software and more like participants in business processes. They're accessing systems, retrieving information, making decisions, and taking action on behalf of users.

That's why identity has become such a critical control. Every agent needs a verifiable identity, appropriate permissions, and accountability for its actions.

Forward-thinking teams are already applying machine identity principles to AI, using cryptographic credentials and automated controls to verify agent identity and manage access.

AI agents may be new. But identity management isn't.

AI has moved from pilot to production. Security is still playing catch up.

Key findings

75%
of organizations deployed four or more AI-powered systems during the last six months

35%
deployed more than ten systems

22%
deployed three or fewer systems

AI is now embedded across business operations, customer engagement, software development, security operations, and enterprise workflows. That's changing the security conversation.

During the pilot phase, the focus was performance. During the production phase, the focus shifts to governance, accountability, and trust.

AI systems can't be treated as experimental technology once they become part of core business processes. They need the same discipline applied to every other critical system.

Everyone's talking about AI governance. Few are doing it.

Key findings

90%
of organizations have discussed AI governance at the executive or board level

50%
have dedicated budgets and formal governance programs

64%
have taken the first step of identifying and inventorying their AI systems

Executive awareness isn't the problem. Execution is. Leaders understand the risk, but AI introduces legal, regulatory, operational, reputational, and cybersecurity concerns that extend well beyond the IT department.

What's missing is operational maturity. Many teams are still building inventories of AI systems to understand where they operate and establish accountability for their actions. To be effective, governance conversations need to turn into action, starting with cataloguing what already exists.

“Why did it do that?” The question half of enterprises can’t answer.

Key findings

53%
can fully trace AI decisions back to the models and source data that produced them

39%
report only partial traceability into how AI systems arrive at decisions

52%
maintain centralized monitoring with regular executive reporting

This is one of the most important findings in the research. Nearly half of respondents lack full visibility into how AI systems arrive at outcomes.

That becomes a problem the moment an AI system produces an unexpected or controversial result. Customers, executives, and regulators will all ask, “Why did it do that?” If the answer is, “We don't know,” trust will start breaking down.

The ability to establish traceability may soon become as important as the AI systems themselves.

Organizations are preparing for the consequences of AI risk.

Key findings

Nearly 90% have evaluated AI-related liability exposure

57% have dedicated budgets for securing AI

86% have formal or informal processes for revoking compromised AI systems

Security teams aren't carrying this responsibility alone. AI governance now touches legal, compliance, risk management, executive leadership, and cybersecurity.

The conversation has expanded from protecting systems to managing business risk. That's a sign AI is maturing, and that expectations for accountability are rising with it.

No country has solved AI accountability.

Key findings

AI inventories remain a work in progress globally, ranging from 59% to 68% across all regions

Dedicated AI security funding is common but not universal, with roughly half of organizations in every market maintaining AI security budgets

The ability to fully explain AI decisions remain remarkably consistent across regions: 53% in the U.S., 54% in the U.K., and 52% in Australia

For all the discussion about regional differences in AI regulation and adoption, the survey results tell a different story.

Respondents in the U.S., U.K., and Australia reported many of the same challenges when it comes to AI accountability. The clearest example is decision traceability. Despite differences in governance approaches and investment levels, only about half of respondents in each country can fully trace AI decisions back to the models and data that produced them.

The broader patterns are strikingly similar as well. AI inventories are still incomplete, dedicated security funding isn't universal, and many teams are still working to establish the visibility and controls needed to govern AI effectively.

The takeaway isn't that one country is ahead and another behind. It's that everyone is working through the same trust challenges at the same time. Geography doesn't seem to play a role in the problem.

Industry perspectives: Every industry has an AI problem. Not every sector has a plan.

AI-related incidents are widespread across every industry surveyed. What's less consistent is how prepared different sectors are to govern, secure, and explain the AI systems they're deploying.

Key findings

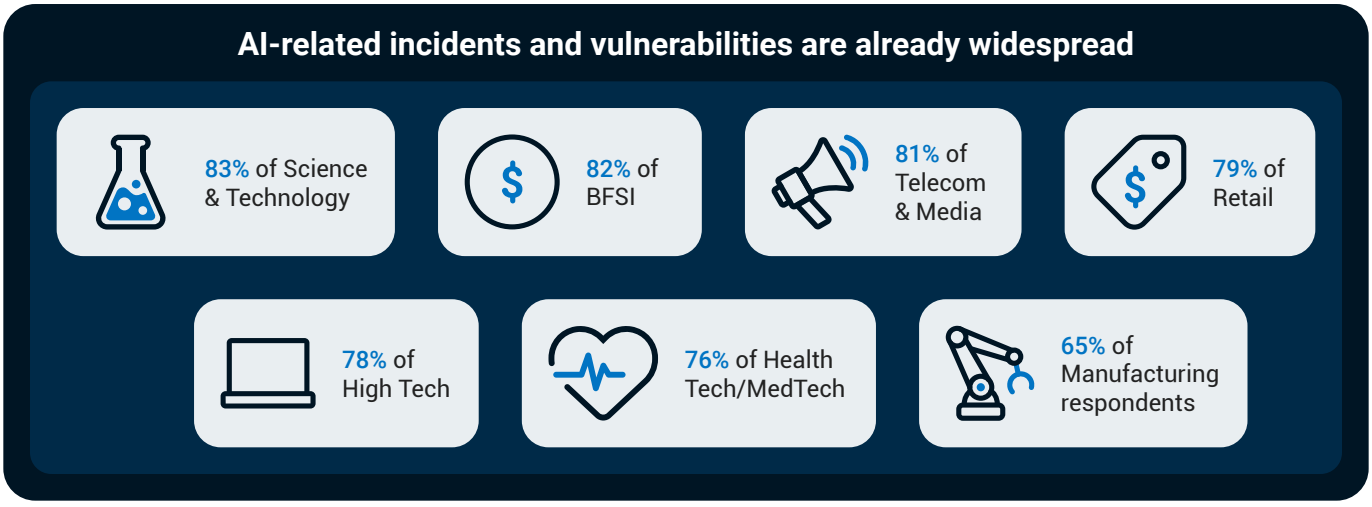

Telecom & Media are furthest along in governing AI

76% maintain complete AI agent inventories, compared to:

- 68% of Financial Services
- 67% of Science & Technology
- 64% of Manufacturing
- 63% of Healthcare
- 59% of Retail

Science & Technology is putting the most money behind AI security

64% maintain a dedicated AI security budget, followed by:

- 62% of Telecom & Media
- 60% of BFSI

Retail has the biggest explainability gap

Only 45% of Retail respondents report full AI traceability, compared to;

- 57% of Science & Technology
- 57% of Telecom & Media

Conclusion: The AI era doesn't belong to the fastest. It belongs to the most trusted.

The AI adoption story has already been told. The new story is what happens next.

Across industries and regions, teams are deploying AI faster than they can govern it. New systems are entering production, new agents are gaining access to data and applications, and new decisions are being made with limited visibility into how they were reached.

Yet, security incidents are already widespread. Accountability is still a work in progress, and many teams still struggle to answer basic questions about what their AI systems are doing and why.

None of this means organizations should be slowing down. It means trust has become a key part of the deployment challenge.

The next phase of AI won't be defined by who deploys the most models, launches the most agents, or automates the most workflows. It will be defined by who can govern those systems, secure their identities, explain their decisions, and earn confidence in their outcomes. Trust is becoming the infrastructure layer for enterprise AI.

Methodology

This report is based on findings from an independent survey conducted by Propeller Insights on behalf of DigiCert in May 2026. The study surveyed 1,001 IT and cybersecurity decision-makers across three markets: the United States (500 respondents), the United Kingdom (251 respondents), and Australia (250 respondents).

Respondents represented organizations across a range of industries and company sizes and included professionals responsible for cybersecurity, IT infrastructure, risk management, compliance, digital transformation, and technology strategy. The survey explored organizational readiness and investment priorities related to artificial intelligence (AI), machine identity security, governance, cryptographic resilience, certificate lifecycle management, and emerging cybersecurity risks.

Data was collected through an online survey and analyzed in aggregate to identify trends, priorities, challenges, and opportunities shaping the future of intelligent trust. Percentages reported throughout this report may not total 100% due to rounding.

About DigiCert

DigiCert is a global leader in intelligent trust. We protect the digital world by ensuring the security, privacy, and authenticity of every interaction. Our AI-powered DigiCert ONE platform unifies PKI, DNS, and certificate lifecycle management to secure infrastructure, software, devices, messages, and AI content, agents, and models. Learn why more than 125,000 organizations, including 90% of the Fortune 500, choose DigiCert to stop today's threats and prepare for a quantum-safe future at www.digicert.com