




Pulse Check Report

AI: The New Superpower And *The New Super-Risk*

How CISOs in U.S. organizations are mitigating AI risks
today and preparing themselves for 2026



Foreword

In January 2025, a Chinese AI startup, DeepSeek, left an unsecured ClickHouse database exposed online. This database housed over 1 million log lines, including user chat histories, backend details, and API keys, as well as notes, letters, and clinical and PHI information from Healthcare organizations in the United States.

The incident dealt a significant blow to DeepSeek. Financially, they had to deal with multi-million USD overhead in incident response, legal, and regulatory areas. Their trust and reputation also suffered a setback because businesses in the U.S. now saw the team behind DeepSeek's flagship AI platform as unable to handle basic cloud configurations. But this was not an isolated incident. Throughout 2025, there have been major AI-related risk and vulnerability incidents across the SaaS & Technology, Financial Services, Healthcare, and Manufacturing industries in the United States.

Large U.S. organizations are at an elevated risk.

AI adoption among U.S. enterprises is high, with AI deeply integrated across multiple processes.

- A single AI-related incident in a SaaS organization or its vendor ecosystem is likely to cost USD 10 million per major customer breach, according to IBM's averages. When you multiply this estimate by dozens or hundreds of affected tenants, the ecosystem-wide cost reaches hundreds of millions of USD.
- For regulated financial institutions, even a relatively small AI data leak could expose sensitive financial data, trigger regulatory investigations, disclosure duties, and potential class-action lawsuits. According to IBM's estimates, an AI-related incident at a U.S. bank or insurer can cost between USD 10 and 20 million, in addition to capital market and reputational damage. Repeated shadow-AI leaks compound this expense and may influence exam findings and capital requirements.
- In healthcare, breach costs are higher due to regulatory penalties and notification requirements. For a U.S. hospital, PHI exposure in an AI-related incident could cost approximately USD 10.22 million and result in an OCR enforcement action. Such incidents also erode patient trust.



CISOS must strike a balance between protection and enablement.

Building an AI Governance Stack that supports incident detection, response, and continuous monitoring becomes a top investment priority for CISOs in 2026.

Table of Contents

1	Methodology: Understanding the audience	01
2	How aware are CISOs about AI-related risks?	03
3	What are CISOs struggling with while mitigating AI risks?	06
4	How are CISOs preparing for 2026?	10
5	Key takeaways for CISOs	12
6	Bonus Giveaway: 7 tenets for CISOs to successfully mitigate AI risks in 2026	14

The AI Pulse Check Report presents insights from a survey of **103 CISOs**, offering a comprehensive understanding of the current state of AI risk management in U.S. organizations. In this report, you will find detailed insights on how prepared CISOs are today to mitigate AI risks and what their top priorities are for 2026.


Methodology: Understanding the audience

01.

Methodology: Understanding the audience


Audience size

103 responses



Country

United States



Job title breakdown

44.94%

Chief Information Security Officer (CISO)


26.97%

VP/Head of Security with CISO responsibilities


28.09%

Other individuals who are primary owners of risk management and compliance strategy in their organization

Organization size breakdown



Industry breakdown



The Current State of AI-Risk Awareness

02.

How aware are CISOs about AI-related risks?


Our study shows that **over 30%** of organizations have experienced a major AI-related security incident within the past 12 months.

Name	Explanation	Cause
Shadow AI or Unapproved AI Usage	Employees use external AI tools, unintentionally exposing internal data.	Lack of governance and safe organizational workflows.
Data Leakage or Model Inversion	Attackers extract sensitive data that models memorize unintentionally.	AI models that learn from personally identifiable data (such as PII or PHI).
API Abuse or Unauthorized Access	Attackers exploit exposed model endpoints to perform computations or data extraction.	Weak API key protection, no rate limits.
Data Poisoning	Attackers inject malicious data into training sets, which causes biased, harmful, or incorrect outputs.	Lack of dataset validation and provenance checks.
Embedded or Integrated AI	AI is embedded by default across most vendors. Most platforms ship AI features (like summaries, copilots, or agents) that process data, often without teams knowing which LLM is used, where data is stored, or what's retained.	CISOs don't consistently treat "daily tools" as AI systems that need identification, controls, and monitoring at scale.

70%

of organizations are aware of AI-related regulations and are actively preparing to comply. This is great! The use of AI carries potential legal, financial, and reputational risks. CISOs who understand them can build compliant controls early, avoid costly rework and fines, and enable organizations to deploy and adopt AI safely and more quickly.

Are you aware of AI-related regulations or standards?



53%

of organizations actively track AI-related risks as a dedicated risk category. This is critical because AI usage creates new, rapidly evolving risks that don't fit into traditional security categories. Tracking AI-related risks as a dedicated category enables CISOs to see adoption, prioritize controls, meet compliance requirements, and prevent blind spots.

How do you actively track AI-related risks?

44.94%


Grouped under broader risks (like third-party or data security)

52.81%

As a dedicated category


2.25%

Do not actively track AI-related risks



U.S. CISOs are increasingly prioritizing AI-related risks because pressure is coming from multiple directions at once: emerging regulatory expectations, internal teams and functions flagging gaps, and boards and executives demanding more transparent oversight. At the same time, customer and prospect demands, often surfacing through RFPs, security questionnaires, and vendor due diligence, are turning AI risk management into a sales and renewal requirement, not just a security initiative. Here are a few motivations that were most commonly mentioned in the survey:

What are the top motivations for CISOs in U.S. organizations to focus on AI-related risks?



Most CISOs in U.S. organizations recognize the risk AI carries, are aware of the regulations or standards that exist, and treat AI risks as a dedicated category.

Now, let's dig deeper to understand how well-equipped CISOs are to mitigate AI risks.


The Current State of AI-Risk Mitigation

03.

What are CISOs struggling with while mitigating AI risks?

30% of organizations are less prepared to mitigate AI-related risks.

How prepared are you to handle AI-related risks compared to other security risks?




Risk windows for AI usage should be measured in minutes because new model features, plugins, or workflows are adopted across teams very rapidly (in a matter of hours), and slow responses turn governance into an after-the-fact cleanup.

2 in 3

organizations take between a week and up to 6 months to implement controls or policy changes in response to AI-related risks. When it takes too long to implement controls or policy updates for AI-related risks, GRC teams end up either stalling promising tools or witnessing employees use them in the shadows anyway. And, both outcomes increase risk.


How long does your organization take to implement controls or policy changes in response to AI-related risks?



27%


of organizations have NOT yet automated risk management for AI threats. AI threats and usage change too quickly for manual processes. Automation enables continuous monitoring, faster detection/response, and scalable governance. Automated AI governance can detect risky prompts, unapproved tools, and sensitive data flows in near real time, helping compliance and risk management keep pace with adoption.

How does your organization primarily manage AI-related risks today?



39% of organizations do not consistently enforce their AI usage policies. This is a red flag because preventing shadow AI usage, reducing security and privacy incidents, and ensuring compliance across teams becomes extremely difficult to manage. Consistent enforcement ensures expectations are precise and repeatable, thereby lowering the risk of human error. It also creates auditable evidence of governance for regulators, customers, and insurers.


How mature is your organization's AI usage policy?



One of the most common AI risks in U.S. organizations is associated with employees uploading sensitive information to AI-native platforms or platforms with built-in or embedded AI capabilities.

21% of organizations have controls in place to prevent uploads of sensitive information to publicly available AI platforms. This is a significant area of concern. Having controls in place to prevent the upload of sensitive information helps avoid data leakage and loss of confidentiality/IP. It also reduces regulatory and contractual exposure. Once sensitive data is shared externally, it may be retained, reused, or exposed.

What AI governance practices do you follow at your organization?



Most CISOs in U.S. organizations are well aware of AI-related risks and how to mitigate them, but are unable to do so efficiently due to limitations in their current AI Governance Stack.


Future-readiness of AI governance and risk mitigation

04.

How are CISOs preparing for 2026?


25% of organizations possess the compliance maturity to mitigate AI-related risks

How would you rate your organization's overall AI risk governance maturity?




69% of organizations have already allocated budgets to mitigate AI-related risks in 2026, while 17% plan to do so in the upcoming budget cycle. This is a clear signal that AI governance is moving from “nice-to-have” to a funded, operational priority. Budget is where intent becomes execution. It enables dedicated ownership, the right tooling, training, and continuous monitoring so controls can keep pace with how fast AI adoption changes day to day. This level of resourcing also typically reflects board and senior leadership alignment, because sustained spend and cross-functional buy-in occur only when leaders agree that AI risk is worth proactively managing.

Do you have budgets for 2026 that can help you manage AI-related risks?



What are the top priorities of CISOs to enable them to mitigate AI-related risks in 2026?



Most U.S. organizations have not yet achieved the GRC maturity needed to mitigate AI risks proactively, but they recognize this and are allocating additional budgets to better prepare for 2026.

Conclusion and Key Takeaways

05.

Conclusion and Key Takeaways

The AI Risk Pulse Check report reveals a clear but uncomfortable truth about AI risk management in U.S. organizations: **awareness is high, but operational readiness is uneven.**

Most organizations now recognize AI as a material security and compliance concern. Nearly 70% of respondents report that they are actively following AI-related regulations or standards and preparing to comply, and more than half (53%) have elevated AI to a dedicated risk category, rather than incorporating it into broader third-party or data security programs.

Unfortunately, awareness has not translated into consistent control execution.

Over 30% of organizations report experiencing a major AI-related security incident in the past 12 months, and the most common incident patterns are precisely the ones that thrive in weak governance environments: shadow AI usage, data leakage/model inversion, API abuse, and data poisoning. These are not “future” risks. They are already creeping into day-to-day operations, often faster than policies and processes can keep pace.

The most significant gaps aren’t about intent. They’re about enforcement and speed.

Nearly 39% of organizations have an AI usage policy that exists but is not consistently enforced, making it challenging to reduce shadow usage, prove compliance, or reliably influence employee behavior. Even more concerning, only 21% report having controls in place to prevent sensitive data from being uploaded to publicly available AI platforms. In this area, a single employee action can result in irreversible exposure of IP, confidential data, PII, financial information, or regulated information.

The path forward is becoming clearer. **AI risk management cannot scale as a set of documents and periodic reviews. It must be a continuous practice.**

Looking ahead to 2026, organizations are investing in AI risk mitigation. Most organizations are still building foundational capabilities, such as implementing technical controls for AI usage, conducting recurring AI risk assessments, providing workforce training, gamifying it, formalizing policies, and increasing automation. 69% have already allocated budget to manage AI risks next year, and another 17% plan to do so in the next cycle.



In the era of AI risks, CISOs need to ensure that speed and safety coexist.

CISOs need governance that is *built for the speed and variability of AI.*

CISOs need Sprinto.

Sprinto is an Autonomous Trust Platform that centralizes trust requirements across security frameworks, security reviews, and vendor due diligence.

Sprinto autonomously executes tasks needed to maintain trust across compliance, audits, risk management, vendor risk, privacy, and AI governance to enable organizations to maintain a strong, reliable trust posture without draining operational bandwidth and resources on repetitive tasks.

Trusted by over 3,000 organizations across 75 countries, Sprinto helps organizations stay audit-ready, manage real-time risks, and scale fearlessly with 300+ integrations and AI-driven automation. Sprinto supports 200+ global security standards, including SOC 2, ISO 27001, GDPR, HIPAA, PCI-DSS, and more.

Founded in 2020 by second-time founders Girish Redekar and Raghuveer Kancharla, Sprinto powers compliance for organizations like Whatfix, Encora, Anaconda, Whatnot, Ultrahuman, WeWork, Everstage, AI Foundation, HackerRank, and many more.

From fast-growing startups chasing their first certification to mature enterprises driving proactive risk management, Sprinto enables trust and resilience at every stage of a company's growth.

950 Million

continuous compliance checks every month

6.5 Million

data sync operations per month

30 Million

entities processed monthly

Speak to a Sprinto expert today!

Book a demo

Bonus Giveaway

7 tenets for CISOs to
successfully mitigate
AI risks in 2026

7 tenets for CISOs to successfully mitigate AI risks in 2026

1. Treat AI as a first-class risk domain, not a subset.

Tracking AI as a dedicated category enhances visibility into adoption, accelerates prioritization, and reduces blind spots. This is particularly useful, as AI simultaneously impacts data, identity, third-party relationships, and engineering workflows.

2. Policy maturity is meaningless without enforcement.

A policy that isn't consistently enforced effectively institutionalizes exceptions. Enforcement is what reduces shadow AI, lowers incident frequency, and creates defensible evidence for customers, boards, and regulators. However, it is crucial to ensure that enforcement does not curb innovation.

3. Prevent sensitive data exposure at multiple levels.

Training alone won't stop accidental uploads. Controls like DLP, extensive browser/app risk assessments, approved-tool allowlists, redaction, and tenant-level configuration are critical because external sharing can be irreversible.

4. Shift from manual governance to continuous, autonomous controls.

Manual tracking can't keep pace with model updates, new tools, and evolving attack techniques. Autonomous AI governance enables continuous monitoring, faster detection/response, and scalable compliance operations.

5. Speed is now a security control.

If it takes weeks to implement mitigations, the organization is effectively operating with known gaps. Build playbooks, pre-approved guardrails, and lightweight change paths so AI risk fixes can ship in hours, not weeks.

6. Focus 2026 spend on the highest-leverage foundations.

The most impactful priorities align with what respondents already indicate: technical controls for AI usage and recurring risk assessments, followed by training, policy formalization, and automation to sustain them all.

7. Design for your environment and avoid one-size-fits-all automation.

AI governance must reflect your framework, system, vendor, and risk thresholds. Flexible automation that adapts to custom environments is what prevents programs from stalling in survival mode.