CISO priorities in the AI era: Balancing security and innovation for a resilient, AI-native enterprise

A report on top themes from Kyndryl's CISO Expert Exchange Program

kyndryl

Table of Contents

01
Foreword

CISOs are entering a phase where AI is no longer something to prepare for—it is already reshaping enterprise risk in real time.

AI moves faster than most organizations' ability to govern and secure their digital environments. Autonomous agents are being introduced into live environments. Applications and infrastructure are becoming more dynamic and interconnected. At the same time, compliance managers and executive boards and regulators demand clearer, more defensible proof of resilience—often before organizations feel confident that they can provide it.

CISOs around the globe are navigating the tension between speed and responsibility. How can they effectively communicate business consequences of cyber risk to their executive board?

Who is accountable when non-human actors are compromised? How can they govern identities that operate at machine speed, and how can they manage risk from ever-expanding vendor and third-party ecosystems?

These questions point to a broader concern that AI adoption may amplify fragility across already complex estates. What becomes clear through these discussions is that AI does not stress one control domain in isolation—it exposes fragmentation across identity, networks, data, infrastructure, applications and security operations. This fragmentation is worsened by the fact that AI agents operate under hidden logic structures, making it difficult to understand and defend their decisions.

Incremental fixes and isolated modernization efforts are not sufficient at AI scale. Compliance, modernization, governance, security and resilience now must move together by design.

Through customer conversations and Kyndryl's invite-only CISO Expert Exchange program, I hear first-hand accounts about the opportunities and challenges AI presents. The following report details conversations with more than 200 CISOs across the US, Canada and Europe. I believe it's fair to suggest that AI is changing the game. CISOs who act now will be the ones who define the terms.

— **Paul Savill**
Global Practice Leader
Cyber Resilience & Connectivity

02

Introduction

As AI moves from experiment to essential, Chief Information Security Officers (CISOs) are no longer just gatekeepers of security. CISOs have become architects of innovation and resilience. They increasingly occupy seats at the senior leadership team table, and their role has morphed into that of a strategic partner. The CISO now guides the executive team on how to proceed with business priorities fueled by AI and autonomous agents without introducing undue risk. CISOs are helping shape strategy, inform investments and communicate risk to the board.

Since 2022, Kyndryl has hosted 24 CISO roundtables in the U.S., Canada and Europe to provide CISOs a forum to connect with peers in the role and share thinking about today's most pressing security topics.

In those meetings, three themes have surfaced repeatedly as priorities for building a cyber-resilient, AI-native enterprise:

- **Risk and compliance as board-level priorities.** Once back-office functions, governance and risk management now dynamically connect cyber risk, AI enablement and business performance. CISOs must translate risk metrics into strategic insight about profitability, resilience and growth.
- **Identity and access management for AI agents.** Identity and access management is no longer static. CISOs must prepare to manage AI agents in accordance with employee access. Doing so will require least-privilege access backed by real-time, context-aware checks at each access attempt, enabled by automated identity governance to keep pace with business transformation.
- **Vendor risk and regulatory resiliency for the AI-native enterprise.** Trust-based assurance is not enough in an AI-native ecosystem of cloud, edge, SaaS, open source and third parties. In a complex environment of downstream dependencies, enterprises must prioritize security-by-design procurement and continuous validation.

This report explores these themes in detail, with perspective and recommendations from Kyndryl's team of cyber-resilience experts.

CISO Insight:

> **How do we use AI to be more productive while being safe? I'm not looking to block it or prevent it. Instead, I'm looking to use it and see how we can do that most efficiently and effectively without the security implications.**

03

Theme 1: Risk and compliance as board-level priorities

As enterprises accelerate AI adoption, organizations must innovate and grow while managing unprecedented levels of risk. To act as strategic partners to the board in this context, CISOs must now do more than report risks and controls. CISOs must be active partners in helping board members understand how cyber and AI risks affect business outcomes, such as profitability, resilience and growth.

With increasing implementation of AI and the advent of agents taking over task execution, organizational risk is expanding to new dimensions. It now includes a larger attack surface, higher attack velocity, AI-triggered operational disruptions, model-driven data leakage and reputational exposure from AI-based actions. Further, AI introduces emerging threats such as prompt based manipulation, AI-augmented social engineering, model drift that weakens safeguards over time and unmanaged shadow AI agents. Governance and compliance approaches must evolve to account for the complexity of these risks.

Risk and compliance management is now firmly a strategic priority situated where cyber risk, AI enablement and business performance meet. And boards are demanding greater clarity and seeking stronger guidance, insights and informed judgment from security leaders.

CISO Insight:

> **Where we have had the most challenge is translating data that goes into a risk register into actionable, insightful information consistently used in setting our strategy.**

How to translate risk into board-relevant insight

The technical findings captured in risk reports are rarely compelling for board members, since this data doesn't influence capital allocation, operating decisions or strategic tradeoffs. As a result, boards need guidance on how managing cyber risk is a strategic necessity for ensuring the business's overall health and success.

3 in 5 business and IT leaders feel their IT is not ready to manage future risks.

(Source: 2025 Kyndryl Readiness Report)

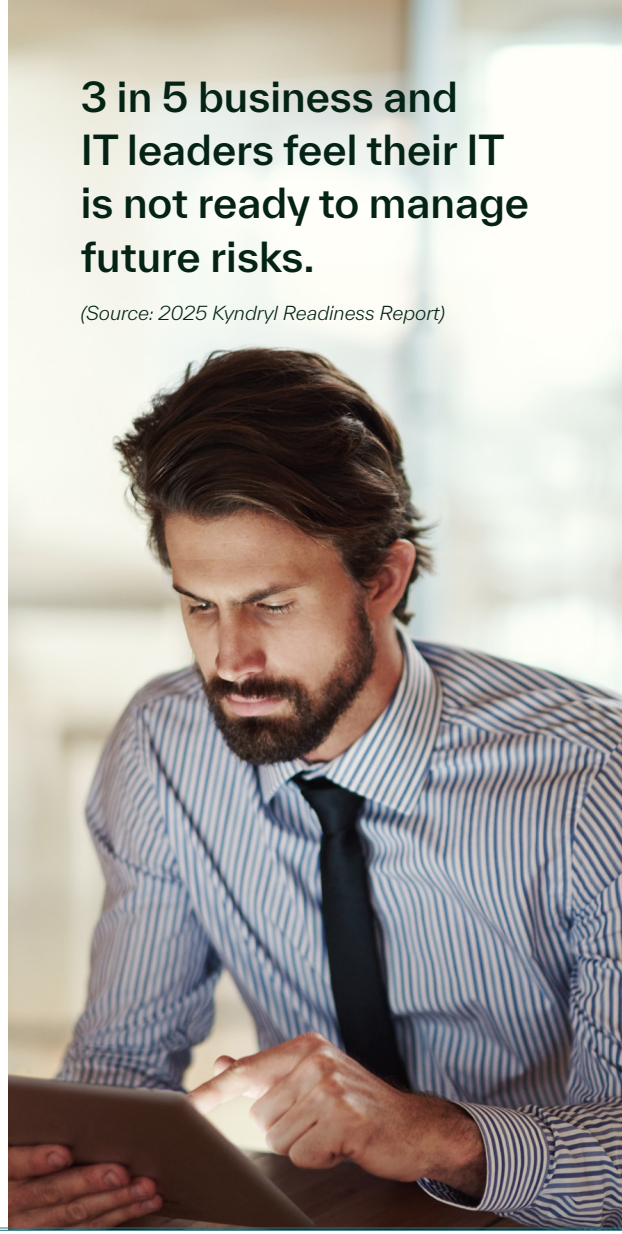

CISO Insight:

> **It's hard for a board member to think about cyber because A, most of them are not technologists and B, they tend to think about it from a financial context. So, one of the things you have to begin to do is force them to think about risk in a different way.**

CISOs' job now includes educating board members about how unmanaged cyber risk can lead to potential harm related to financial health, operational functioning, regulatory exposures, reputation and IP management. They must shift the conversation away from technical information and toward business consequences, making clear what is at stake if security is not a holistic priority for the organization.

These consequences include potential revenue loss, operational downtime, regulatory sanctions and loss of customer trust.

There are several ways to shift security language to better engage and educate board members:

- **Anchor risk and compliance in business priorities:** Decisions about how and where to prioritize cybersecurity investments, what risks to accept and how to accelerate innovation are best informed by aligning risk and compliance with top-down business priorities rather than bottom-up control inventories.
- **Express risk in financial and operational terms:** CISOs can make risk more visible in business language by translating cyber-security threats into financial terms via cyber-risk quantification (CRQ), a methodology that enables the measurement and prioritization of risks based on monetary impact. This shows board members that management of risk and compliance can shift from being a defensive cost center to an enabler of strategic profits and growth.

CISO Insight:

> **I think one of the main challenges is you need to be able to showcase evidence from the top down that everything you have done and designed is based on your business needs.**

- **Provide proof of supporting business priorities:** CISOs can showcase evidence that the organization's evolving governance frameworks, risk-reduction activities and compliance regimens are designed to further business priorities.
- **Orchestrate risk-appetite discussions:** Boards can engage with security teams on setting the organization's tolerance for risk by discussing scenarios based on key risk indicators (KRIs). Conversations with this kind of specificity help board members understand how risk and governance relate to business goals and outcomes.

Clear, board-level security conversations lay the groundwork for safer innovation.

CISO Insight:

> **When it comes down to resource planning and making sure that we have US dollars and cents allocated, the business is much more motivated now to help us work the plan."**

04

Theme 2: Identity and access management for the age of machines

Identity has become the enterprise's primary control plane, displacing reliance on the network perimeter. Yet most identity and access management (IAM) programs were designed for human users accessing web applications. While still necessary, that model is no longer sufficient. IAM must now include machine identities, such as services, workloads, bots, APIs and AI agents, and enforce continuous, context-aware authorization across a far more fluid environment.

AI agents require identity and access security and monitoring, just like human workers do. But unlike humans, they do not possess inherent ethical judgment, which instead must be deliberately engineered and enforced from inception. Their access must be tightly governed and limited to the specific data and actions for the tasks at hand, with continuous controls and monitoring to ensure safe and compliant behavior.

Machines behave and scale differently than humans, introducing new trust boundaries that traditional IAM structures aren't equipped to accommodate. Enterprises that fail to evolve their IAM approaches will continue to find that identity management remains their most vulnerable point, even after adopting zero trust.

3 in 4 IT and business leaders report they are investing in AI for cybersecurity (such as fraud detection and risk management). This is more investment than in any other AI capability.

(Source: 2025 Kyndryl Readiness Report)

CISO Insight:

"These agents are simply digital employees, and we should be securing them the same way we would our employees. We must make sure that each agent has an identity that's been restricted to only access the data that it is authorized to have.

Why machines identities break traditional IAM

Several fundamental shifts strain the boundaries of legacy IAM:

- **Volume and velocity:** With certificates, tokens, API keys, service principals and AI agents proliferating, traditional identity management processes are overwhelmed by the rapid growth and complexity of digital identities.
- **Autonomy and authority:** AI agents increasingly make their own authorization decisions, shifting identity management from a simple access function to a real-time operational and execution interface. As agents learn from execution, they may develop unanticipated behaviors or drift beyond intended access paths, making continuous, runtime validation essential.
- **Opaque risk:** The proliferation of nonhuman identities opens new avenues of risk that are easy to overlook. Model choice, prompt flows, tool access, data provenance and output integrity are all now identity questions.

CISO Insight:

> **There are a million different models, and hyper-specific ones. How do we assess those models' risk and drive toward more consistent use of vetted models?**

How to evolve IAM

Enterprises must evolve their IAM or risk leaving identity management as their enterprise's weakest link. Managing digital workers with the same rigor as human employees requires rethinking IAM entirely: it is no longer simply a gateway but a real-time control plane for dynamic access and continuous oversight. CISOs will benefit from considering the following guidance in this process.

- **Manage machine's identity lifecycles:** Restrict machine's access to the minimum privileges they require, continuously adjusting those privileges based on context and role changes. Enforce real-time verification at execution and revoke credentials immediately upon decommissioning.
- **Adopt policy-as-code frameworks:** Implement frameworks that authorize agents to access data only within defined parameters under approved tickets. Incorporate human-in-the-loop review and strict tool-level guardrails, particularly for high-risk decision points where unauthorized behavior would have outsized operational or business consequences.
- **Apply strict data management controls:** Enforce consistent, permanent data classification and labeling so agents can reliably recognize sensitivity levels. Identities, human or non-human, that are not authorized for a given dataset should have no possible path to retrieve or infer it.
- **Enact model risk governance:** Govern models with the same approach as agents, with each attached to its own metadata. Vet models and use authorization policies to match models to tasks.
- **Derisk automation:** Define and implement strict controls that dictate when an agent may act autonomously, when it must pause for human approval and when it should limit itself to offering recommendations.
- **Maintain continuous governance:** Prevent entities from maintaining access or privileges that no longer apply to their purposes and tasks with a verifiable chain of custody listing prompts, access, authorizations and outputs.

Modern identity governance for both human and machines demands tighter lifecycle management, clearer authorization boundaries and ongoing oversight that keeps pace with rapid operational change. Enterprises can make identity a source of resilience by treating IAM as a context-sensitive, continuously adaptive system rather than a static locus of control.

05

Theme 3: A new imperative for managing vendor risk and regulatory compliance

The boundaries between organizations and their vendor ecosystems have become extremely porous, and even nonexistent. Organizations must be able to show that they actively manage ecosystem risk as an extension of their enterprise rather than simply relying on contractual agreements with vendors.

Regulators around the world now expect organizations to demonstrate continuous, evidence-based oversight of their entire vendor ecosystem and are increasingly holding firms directly accountable for failures anywhere in their supply chain.

CISO Insight:

> **It's important to know how your critical third-party providers are running their operations.**

Oversight without control in an expanding ecosystem

Organizations may have thousands of direct suppliers, each with their own third- and fourth-party relationships. A single vulnerability such as an exposed API, insecure software update or compromised AI agent can trigger a domino effect of damage throughout the ecosystem.

Enterprises remain accountable for downstream breaches. AI compounds this dilemma, as vendors increasingly embed AI into their offerings without always revealing their model training methods or data management practices. Even tools designed to increase security can inadvertently introduce new vulnerabilities.

Smaller vendors, in particular, can introduce risk into the ecosystem, as many play critical roles in application development but aren't equipped to implement enterprise-level controls and security practices. The result is an ecosystem in which risk is unevenly distributed but universally shared.

CISO Insight:

> **The expectation that small companies should conduct in-depth security audits of major global technology suppliers is unrealistic, highlighting a disconnect between regulatory requirements and practical capability.**

How to enact secure-by-design procurement and continuous validation

Verifiable, ongoing assurance is now non-negotiable. Core practices such as secure-by-design software development practices, identity-based access controls and AI governance standards must be built into contracts, measured over time and continuously validated.

- **Secure-by-design procurement:** Organizations must shift to engineering security and resilience by design in the vendor-selection process. This approach embeds security, privacy and AI governance directly into vendor selection, contracting and onboarding to set expectations and vet providers up front.
- **Point-in-time assessments:** Traditional annual questionnaires cannot keep pace with the speed and volatility of AI-driven vendor risk. Leading enterprises are now pairing data-driven external monitoring with risk-based tiering instead of treating them all uniformly. The old model assessed all vendors but monitored only the critical ones; the new model continuously monitors all vendors and dynamically deep dives those that present elevated or emerging risk.
- **Continuous risk intelligence and telemetry:** Static, audit-heavy assurance models do not provide a scalable path forward. Enterprises can instead tap nonintrusive monitoring, open-source data, curated threat intelligence feeds and AI-driven analytics to get real-time insight into vendor security practices and threat trends.
- **Data sovereignty and AI governance:** Organizations need to embed security and sovereignty criteria into procurement processes and continuously validate vendor performance. This requires clear control over where sensitive data and AI models reside, how they are used and moved across jurisdictions, and who can access or influence them across the ecosystem. It also extends to operational sovereignty, ensuring that critical business processes delivered or supported by vendors can remain secure, resilient and compliant even when third-party systems or infrastructure are involved.

The vendor ecosystem is now part of the enterprise itself, and a secure-by-design procurement approach is now the only viable operating model for a resilient, AI-native enterprise. Regulators expect continuous control and measurable oversight throughout the extended enterprise.

82% of enterprises experienced significant outages within the last year. 25% of these outages were caused by a third-party vendor or supplier.

(Source: 2025 Kyndryl Readiness Report)

06

Priority recommendations for CISOs

Across the three themes in this paper, the guidance points to three priority recommendations for CISOs:

- **Step into a strategic role.** CISOs should embrace a strategic leadership role, influencing senior executive decisions with a focus on risk reduction. Moving in this direction requires communicating cyber risk in terms that resonate with the board. This approach will help other C-suite leaders understand the financial impact and business value of securely managing AI.
- **Focus on identity and access.** CISOs should approach the management of digital workers with the same oversight and rigor they bring to human employees. Identity and access management is now an integral part of an overarching system that governs cyber security and impacts business decision-making.
- **Rethink oversight of vendors.** CISOs should apply measurable oversight to the sprawling vendor ecosystem in which they are embedded. Resilient, AI-native enterprises must continuously validate the performance of third- and fourth-party vendors rather than rely on trust-based assurances and contractual language.

These recommendations underscore that CISOs have evolved from security guardians to strategic leaders shaping how their organizations adopt and govern AI. Their influence is now a differentiator that signals trust to customers, investors and employees. By embedding responsible AI governance and security practices, they enable their leaders to build confidence, accelerate AI adoption and innovation, and ensure durable resilience.

07

How Kyndryl helps

Kyndryl helps organizations transition from their current infrastructure to a high-performance, AI-native future. Our security and resiliency services are designed to accelerate growth and innovation, offering scalable solutions that protect data and AI assets, ensure compliance, and stay ahead of emerging threats. We secure mission-critical infrastructures and develop AI-native capabilities by combining strategic advice with hands-on execution, empowering leaders to confidently pursue growth and innovation.

08

About the CISO Expert Exchange

The Kyndryl CISO Expert Exchange is a curated, peer-driven program that brings together CISOs and senior IT leaders for meaningful collaboration and shared learning. Select members of the community convene quarterly in virtual sessions designed to foster open dialogue around today's most pressing cybersecurity challenges, enabling CISOs to exchange real-world insights, perspectives and strategies with trusted peers. There are CISO Expert Exchange communities in the U.S, Canada and Europe. Participation is open to CISOs and senior security executives from large enterprises. Join to share, learn and help make a difference in the security and resiliency community.

kyndryl