kyndryl

CISO Expert Exchange

November 25, 2025

Executive Summary

Hosts:

Denis Villeneuve
Cybersecurity and Resilience Practice Leader, Kyndryl Canada

Tony De Bos
Vice President, Governance, Risk and Compliance, Kyndryl

Overview

In this CISO peer exchange, security leaders from various Canadian organizations convened to share challenges and opportunities around various cybersecurity and resiliency topics. The conversation focused on three core areas: enhancing data security measures, implementing robust artificial intelligence governance frameworks, and managing the emerging risks associated with autonomous AI agents. The consensus emphasized that successful management in these areas requires strong corporate leadership, continuous monitoring, and structured processes to reduce organizational risk and maximize efficiency.

Expert exchange themes

- Data stewardship and attack surface reduction
- Establishing effective AI governance and responsible use
- Managing emerging agentic AI risks and the need for oversight

Data stewardship and attack surface reduction

- Organizations often struggle to manage decades of accumulated legacy data, making strong corporate buy-in and clear data ownership essential for effective risk management and compliance efforts.
- Members widely agreed that deleting unnecessary data reduces organizational risk, noting that long-living data complicates legal and regulatory compliance. Participants actively seek data minimization strategies to ensure organizations retain only the information strictly necessary to support business operations.
- Gaining complete corporate support for data governance presents a significant challenge, especially regarding data ownership and determining who possesses the authority to make retention and deletion decisions. Without decisive mandates from senior leadership, security and information technology teams often must operate based on conjecture.
- The application of automated labeling solutions is crucial for adequately protecting information, although participants noted this remains challenging for complex engineering files versus standard office documents. One executive shared a best practice of using scanning tools to identify users who possess disproportionately large amounts of sensitive files, allowing focused cleanup campaigns.
- Securing outbound data transfer remains a universal struggle, as users often circumvent secure corporate solutions by utilizing various unmonitored external transfer services. Best practices include ensuring double encryption (at rest and in transit) and implementing continuous processes to verify that third-party recipients comply with agreed-upon deletion schedules.

Establishing effective AI governance and responsible use

- Implementing robust governance requires clear leadership and structuring corporate oversight to prevent duplication across committees, ensuring that AI adoption is simultaneously innovative and responsible.
- Organizations universally reported challenges with governance structure, where key corporate functions — such as legal, privacy, and risk management — must staff multiple governance councils, significantly slowing the process for reviewing and approving AI initiatives. One organization addresses this by focusing the governance council on specific risks, such as intellectual property infringement, to increase efficiency.
- One firm promotes innovation through establishing a center of excellence composed of various leaders, monitoring user activity on external AI platforms, and producing regular reports for leadership review. This centralized approach ensures that new AI tools are vetted for business necessity before widespread implementation, preventing tool proliferation.

- A key guardrail adopted by one participant mandates that no fully automated AI solution should interact directly with customers, requiring a human operator to remain involved in the loop for sensitive processes. Another participant highlighted the difficulty of monitoring the massive number of new and embedded AI functionalities constantly being introduced into vendor products.
- The discussion also focused on commercial risk: executives must scrutinize vendor contracts to ensure a common understanding of how AI is defined, as vendors may try to use ambiguous definitions to bypass security controls and gain greater access to organizational information. Members agreed that focusing on data value rather than contract value is essential when assessing third-party risk.

Managing emerging agentic AI risks and the need for oversight

- Autonomous AI agents introduce a parallel operational world that demands equivalent identity management, logging, monitoring, and predefined behavioral guardrails, with human oversight implemented for exceptions that cross established risk thresholds.
- The group discussed treating autonomous agents as non-human entities operating in a parallel world, requiring the deployment of similar controls used for people, including access management, logging, and the capability to stop or contain agent behavior. Testing agent behavior using digital twin environments is recommended before launching them into production.
- One participant raised the concern that initial plans for human supervision often degrade over time as confidence grows, potentially leading to unsupervised agents making critical decisions. Therefore, continuous monitoring of behavioral guardrails and presenting performance metrics that are relevant and tangible to business outcomes are vital to maintaining trust.
- The complexity of maintaining human oversight in automated processes where multiple agents communicate and interact with internal and external parties poses a major difficulty. The suggested best practice is to conduct a business risk assessment for each agent, defining thresholds, and routing exceptions directly to human reviewers for intervention.
- Participants also discussed subjective risks, such as one executive sharing a use case where an AI coaching bot for management could inadvertently alter the firm's culture or values if adopted without fully assessing the long-term impact of its recommendations. This example underscores that governance must address cultural and human capital risks, not just technical vulnerabilities.

Summary

The conversation highlights that managing modern digital risk requires organizations to treat data, AI, and autonomous agents not just as technical problems, but as core business assets and operational entities requiring dedicated, top-down governance. Failing to manage decades of data sprawl is akin to leaving the front door unlocked, while failing to govern AI is like granting unsupervised access to the keys once inside.

To learn more about the Kyndryl Canada CISO Expert Exchange or to become a member of this community, please visit this website.

kyndryl

This document is current as of the initial date of publication and may be changed by Kyndryl at any time without notice. Not all offerings are available in every country in which Kyndryl operates. Kyndryl products and services are warranted according to the terms and conditions of the agreements under which they are provided.