CISO Cross-Industry

Expert Exchange

Q4 Executive summary
October 21, 2025

kyndryl

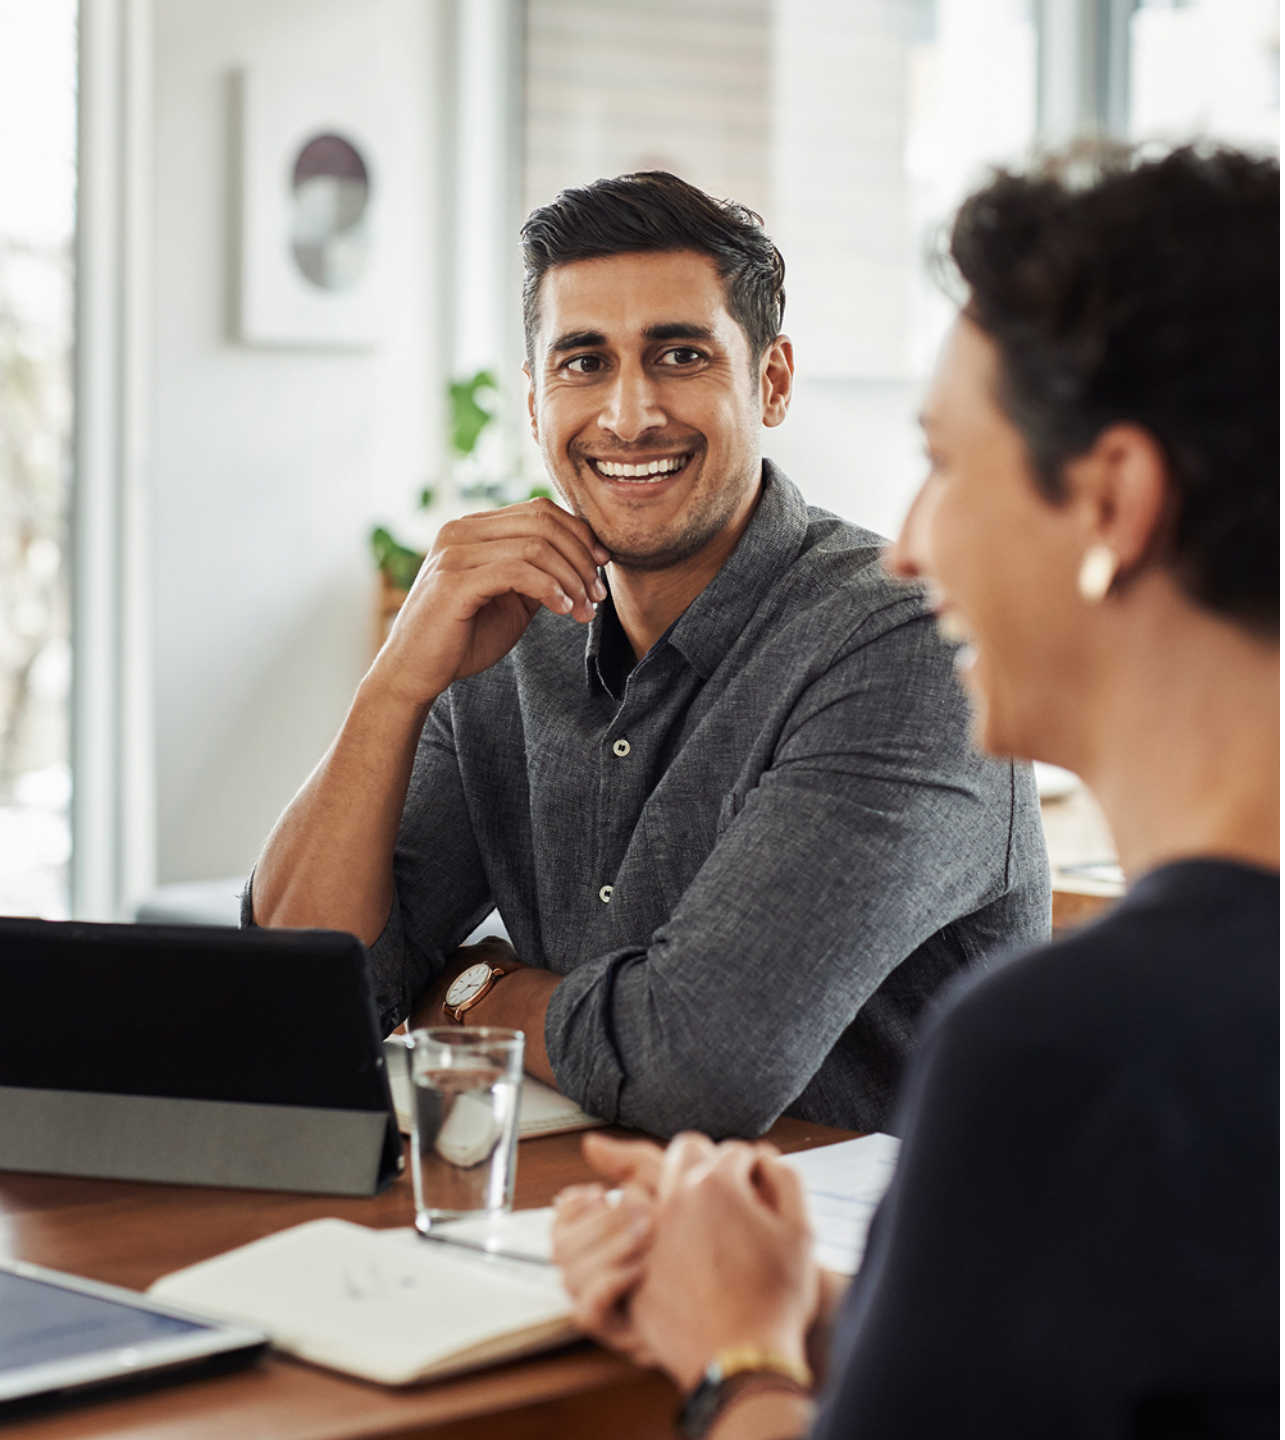

Overview

Security executives from various industries convened to discuss pressing cybersecurity issues, and four primary themes dominated the discussion: implementing a mature response to the recent F5 breach; establishing secure identities for autonomous AI agents; formalizing AI governance and risk assessment frameworks; and strategically utilizing AI to enhance cyber defense capabilities.

Host

Aaron Severance– US Security & Resilience Practice Lead, Kyndryl

SME

Jimmy Nilsson– VP of Professional Service (Zero Trust), Kyndryl

Key topics

Managing Risk from the F5 Breach

- Security leaders stressed that the F5 breach affecting widespread network equipment necessitates a measured, analytical response over panic-driven reaction. The immediate goal is to stabilize the environment and minimize risk, allowing time for thoughtful modernization decisions.

- A technology provider recommended that organizations immediately prioritize patching all affected devices and urgently upgrading any end-of-life infrastructure. They developed a set of 28 compensating controls designed to be applied upstream and downstream of the impacted systems to further reduce exposure during the interim period.

- Executives noted that organizations should leverage relationships with threat intelligence providers, as specific reports offer non-public intelligence and Indicators of Compromise that aid in internal threat hunting efforts and deep inspection.

- A key learning shared by participants was the importance of applying not only patches but also additional hardening best practices recommended by the vendor. Furthermore, leaders emphasized examining fourth-party risk, as many critical digital service vendors also rely on the compromised infrastructure.

"We advise our clients... to make sure that we patch what can be patched and develop a package for compensating controls, rather than running and panicking and immediately replacing all equipment."

— Jimmy Nilsson, VP Professional Service (Zero Trust)

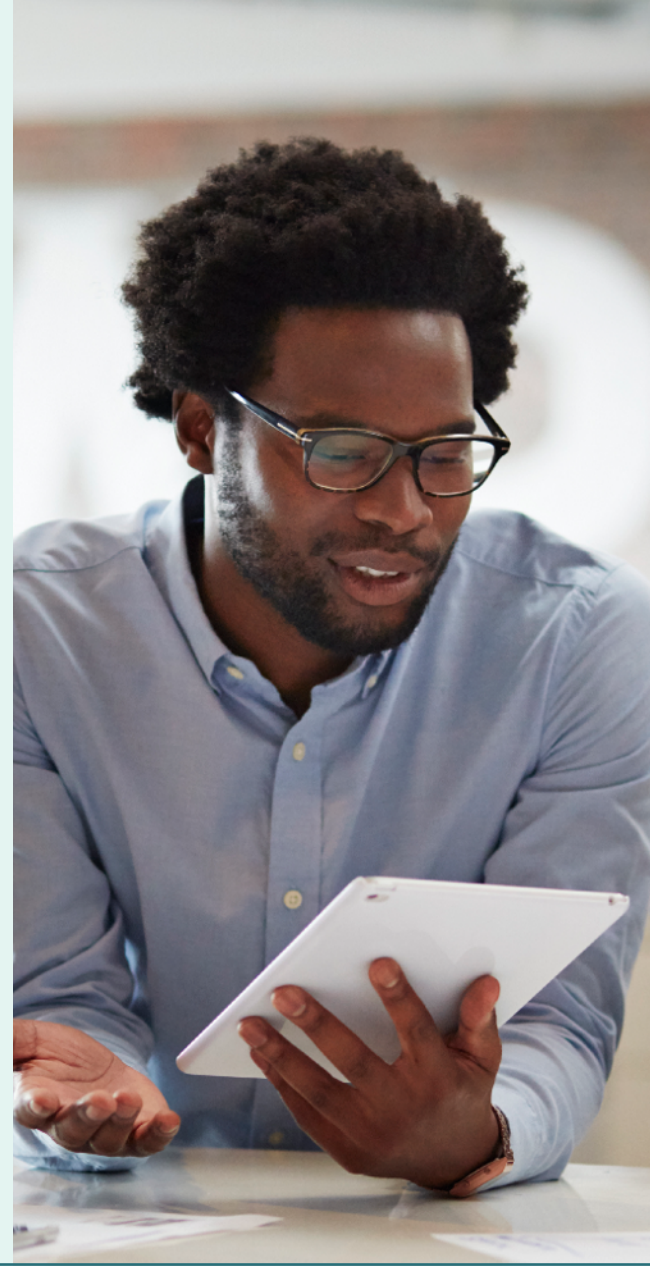

Securing Agentic AI Identities and Access

- Security leaders agree that AI agents effectively function as “digital employees” and securing them demands treating their access and identity management as a core priority. The primary challenge is preventing privilege escalation and unauthorized data access when these agents move from individual use to team-wide deployment.

- Executives widely agreed that for agents used strictly for individual productivity (e.g., general AI assistants), the risk is lower because the agent uses the identity and existing access privileges of the human user.

- The risk profile dramatically increases when these agents are shared across teams or the organization. Security leaders are concerned that a shared agent might retain the sensitive access rights of its developer, thus granting unauthorized access to other users.

- Many organizations are actively addressing this by establishing distinct “silicon identities” for autonomous agents, ensuring their access is strictly defined and limited to only the data necessary for their function. One leader emphasized that allowing the agent to inherit the builder’s identity leads to high-risk situations. A technology provider suggested that applying principles from the Zero Trust framework provides a foundational lens for AI security, focusing on managing the agent’s identity, governing the application it runs on, and strictly controlling data access.

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

— CISO Expert Exchange Member

Establishing Governance and Risk Assessment

- The rapid adoption of AI solutions is outpacing the maturity of foundational data security programs in many organizations. Security leaders are working to formalize governance processes and risk assessment tiers quickly to manage this gap and ensure safe, productive adoption.

- Participants agreed that AI readiness directly relies on a strong enterprise data management program, noting that weak data inventory, classification, and security posture severely restrict the speed and safety with which organizations can deploy AI.

- An executive emphasized the importance of establishing explicit governance policies early to set expectations and create regulatory frameworks for a large enterprise. When evaluating new AI-powered vendors, leaders prioritize solutions that allow the organization to disable the AI function or provide application programming interfaces for signal integration and data monitoring.

- Participants stressed that the vetting process for third-party vendors must evolve to include continuous monitoring and deep inquiry into how the vendor uses data and where agents run. One security leader strongly advised hosting vendor AI agents (e.g., Docker containers) within the organizational network to ensure monitoring of traffic flow and third-party communications.

- A strong argument was made for adjusting the operating model to move away from siloed security thinking. Security leaders agreed that cross-functional teams, which include business, technology, and security representation, are necessary to build the most optimal and practical security architecture for AI solutions.

> **"We are only able to move at a certain pace because there's a lot about our data and the exposure of that data, both internally as well, that we just don't know."**
>
> **— CISO Expert Exchange Member**

Leveraging AI to Enhance Cyber Defence Capabilities

- While still in the early stages, security professionals are actively investigating how AI can improve cyber defense, focusing mainly on productivity, enrichment, and analysis, rather than autonomous action. The goal is to evolve traditional automation into agentic AI capabilities for low-risk, high-efficiency security tasks.

- The most immediate application of AI is within the Security Operations Center to accelerate alert triage and incident detection. AI also offers significant value in identity and access management by transforming static, noisy logs into contextual, predictive data for better decision-making.

- Technology providers highlighted that advanced behavior analytics is an area where AI provides substantial benefits, using massive data ingestion and signal correlation across multiple security tools to map normal behavior and proactively suggest policy changes based on emerging anomalies.

- A key challenge is the relative immaturity of tools designed specifically to monitor and detect AI-unique risks, such as data poisoning, prompt injection, and model drift. Security architects are currently vetting emerging security stacks that promise capabilities in these critical areas.

> "The core function in Zero Trust is the decision point, where you want to take signals from different capability and apply a policy to form your decisions; this is an area that I think we as an industry should look more at."
>
> — Jimmy Nilsson, VP Professional Service (Zero Trust)

kyndryl

The CISO Expert Exchange is hosted by Kyndryl. Please contact Justin Haney with any questions about Kyndryl or this Exchange.