Canada CIO

Expert Exchange

Q2 Executive summary
May 28, 2026

kyndryl

Overview

The recent roundtable explored how frontier AI is reshaping vulnerability management and resilience, with offensive capabilities compressing remediation timelines and driving changes in architecture, supply chains, and risk ownership.

Key themes included faster patching, open source and AI-driven vulnerability risks, legacy modernization toward agentic systems, and stronger shared accountability for resilience.

Hosts

Karen Cheng
CTO, Kyndryl Canada

Brian Medeiros
President, Kyndryl Canada

Special Guests

Denis Villeneuve
Cyber Resilience and Connectivity Practice Leader, Kyndryl Canada

Cory Musselman
CISO, Kyndryl

Key topics

Accelerating Patch Velocity in a Frontier AI Threat Landscape

Exploit Chaining: Participants discussed how advanced AI models are now capable of chaining together previously unrelated software flaws and reverse engineering compiled binaries back into readable source code. These capabilities create exposure paths that traditional defensive programs were never designed to address. Leaders noted that Microsoft reported, across forty technology vendors, a projected thirty to forty percent increase in patch volume across the next four consecutive quarters, with comparable surges expected from other major suppliers participating in Project Glasswing.

Continuous Patching: Executives agreed that historical monthly patching cadences are functionally obsolete and that organizations must transition toward aggressive twenty-four to forty-eight hour service level agreements for critical remediation. To absorb this dramatic acceleration, leaders described a three horizon strategy combining near term resource reallocation through dedicated response teams, short term automation of vulnerability triage and deployment workflows, and longer term agentic enablement across the broader operating model. Foundational practices, including hardened identity controls, network segmentation, accurate asset inventory, and multifactor authentication, were repeatedly emphasized as more important than ever, though participants stressed that these fundamentals must now operate at machine speed to be meaningful.

Speed to Response: Several participants raised concerns that publicly available open source AI models are likely no more than six months behind leading frontier capabilities, indicating that the population of capable offensive tools available to threat actors will expand rapidly. Leaders observed that the time required for an adversary to understand and exploit a newly disclosed vulnerability has compressed from weeks and months down to days, with some scenarios now potentially measured in hours. Consequently, executives are systematically revisiting risk decisions previously made under fundamentally different threat assumptions and rewriting their internal risk algorithms to reflect this dramatically altered landscape.

> "The patch velocity is becoming a board-level risk, and the historical monthly cycles are simply no longer fast enough to be considered safe."
>
> — CIO Expert Exchange Member

AI Vulnerability Economics & Open Source Exposure

Vendor Exposure: Executives observed that Project Glasswing is actively running advanced AI models against the source code of leading commercial platforms, but small to mid sized software vendors, SaaS providers, and open source maintainers lack the resources to participate equally. Participants identified these smaller producers as the most likely sources of unaddressed vulnerabilities across the near to medium term, with one leader noting that organizations heavily dependent on internally developed applications and open source frameworks face a particularly difficult position.

AI Economics: The token economics of running frontier AI against enterprise codebases are emerging as a strategic challenge for technology leaders. One participant shared field intelligence indicating that a major vendor expended approximately ten to twenty million dollars in token costs over a sixty day window, analyzing its own software, surfacing tens of thousands of new vulnerabilities, of which ten to fifteen percent were rated severe at a false positive rate under five percent. Another executive observed that comparable analysis using publicly available leading frontier models can run several thousand dollars per code repository, while organizations in Project Glasswing are reportedly spending seven to eight figures continuously.

Industry Certifications: Leaders explored the prospect of a forthcoming frontier AI model certification for commercial software, conceptually parallel to existing assurance frameworks but specifically affirming that a vendor has subjected its source code to advanced model analysis and remediated identified vulnerabilities. Participants agreed that procurement leverage will ultimately be essential to drive industry adoption of such standards, with executives prepared to use vendor selection decisions to enforce minimum expectations. The conversation also surfaced an emerging discipline blending financial operations with artificial intelligence operations, where executives must continuously balance the cost of inference against the operational value generated by the agentic workloads consuming those tokens.

> "Companies are blowing through their entire annual AI budget within the first two or three months of the year because they did not anticipate how quickly the costs would compound."
>
> — CIO Expert Exchange Member

Modernizing Legacy Applications Toward Agentic Architectures

Legacy Drag: Executives detailed substantial legacy application portfolios that complicate the acceleration challenge. Rather than incrementally patching deteriorating systems, several leaders are using the current threat environment as justification to accelerate previously deferred modernization roadmaps. Organizations are actively weighing whether to rebuild applications like for like or to consolidate multiple legacy systems into fewer, more capable agentic experiences that fundamentally rethink how work flows through the enterprise.

Experience Transformation: Participants shared striking examples of cost compression achieved through agentic delivery models, with one public sector technology leader citing a single rebuild that delivered a previously fifty four million dollar scoped project for approximately two million dollars in actual delivery cost. Executives emphasized that the user experience implications are equally transformational, describing card deck interfaces where business users swipe through outlier cases requiring human judgment while AI agents process the high volume baseline work autonomously. Leaders observed that once internal stakeholders experience these next generation interfaces, demand for the legacy equivalents collapses almost immediately and accelerates the broader transition.

Change Constraints: The conversation underscored that organizational change management is now the binding constraint rather than engineering velocity. Leaders noted that customers and internal stakeholders accustomed to traditional menu driven workflows often struggle to adapt to non equivalent agentic experiences, which limits how aggressively organizations can consolidate. Several participants acknowledged that organizations with predominantly internal commercial off the shelf application portfolios face a particularly difficult transition because they must simultaneously develop new build competencies while retiring purchased software estates that were never designed to be replaced internally.

> "Once stakeholders see an agentic front end on legacy data, nobody wants the old applications anymore — it changes the conversation entirely."
>
> — CIO Expert Exchange Member

Distributing Risk Ownership and Strengthening Operational Resilience

Risk Distribution: A consistent theme across the discussion was that centralized security functions cannot absorb the volume of remediation decisions implied by the new threat velocity and must therefore explicitly transfer accountability to operational leaders. Participants described emerging governance models where security publishes remediation requirements and individual business units formally accept residual risk if they cannot meet defined patching timelines. One financial services participant shared a structural change embedding dedicated patching windows into every product release cycle, signaling unambiguously to business stakeholders that feature delivery and security remediation must coexist permanently.

Resilience Planning: Executives discussed how operational resilience programs are being substantially expanded to accommodate both planned and unplanned outages during this elevated threat window. Several leaders described tiered application redundancy strategies, where tier one critical systems are engineered for true active-active failover, secondary tiers for cross-site failover, and lower priority systems for planned downtime. One participant described an operational resilience program built around critical business service failover capable of running on secondary infrastructure for up to two weeks while primary systems are remediated, while another shared national cybersecurity guidance recommending organizations prepare to isolate systems for up to three months and develop credible plans to rebuild from a severe cyber incident.

Incident Response: Executives emphasized that incident response playbooks must now anticipate multiple major simultaneous events, including mixed internal and third party incidents that historically would not have overlapped. Leaders are revisiting their definitions of minimum viable operations to identify the absolute smallest set of systems required to sustain the enterprise during prolonged disruption. The conversation also examined autonomous and continuous "red teaming" (ethical hacking) as the strategic successor to point in time penetration testing, with participants exploring both fully agentic platforms and human supervised offerings while acknowledging that current models are not yet accurate enough to authorize unsupervised production changes. Several leaders summarized this posture as keeping "humans on the loop" rather than "in the loop", where automation runs continuously, but qualified engineers retain authority over any consequential action.

> “Every release now leaves a window for patching because we essentially told the business that we no longer have a choice.”
>
> — CIO Expert Exchange Member

kyndryl

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