

Frontier AI and the Need for Risk-based, Automated Resilience

Insights from the June 2026 CISO Expert Exchange

kyndryl®



Contents

01

**A new look at
vulnerability and
resilience**

02

**The frontier AI threat
and the coming patch
surge**

03

**Emergency patching
and risk-based
prioritization**

04

**Maintenance windows
and regulatory
expectations**

05

**Resilience across
legacy, operational
technology & Zero Trust**

06

**How to shift to risk-based,
automated resilience**

07

**What is the Kyndryl
CISO Expert Exchange?**

08

**Kyndryl cyber
resilience services**

01

A new look at vulnerability and resilience

Frontier artificial intelligence is producing a step change in vulnerability discovery across banking, financial services, manufacturing, and industrial environments.

AI-driven discovery is triggering an imminent patch surge, rendering traditional patching cadences obsolete. Organizations see the need to renegotiate the trade-off between service availability and disruption with regulators while rethinking resilience across legacy systems and operational technology environments.

On June 3, 2026, Kyndryl hosted a Security & Resilience Expert Exchange to explore questions and ideas about how organizations can best navigate AI readiness, governance, and data sovereignty in a post-Mythos world. Seventeen participants joined the wide-ranging conversation. The following pages present highlights from the session and related notes for consideration.



02

The frontier AI threat and the coming patch surge

Anthropic's new LLM, Claude Mythos, represents a step change in cyber threat dynamics, making a significant surge in patching demand likely imminent. The most critical development is its ability to chain lower-severity vulnerabilities into high-impact exposures. More than 10,000 high and critical severity vulnerabilities have been disclosed across tested platforms. **Patch volume is expected to rise 40-50% over the coming months**, with Microsoft projecting roughly 40% more patches over the next four quarters. As various models surface different flaws, the same code may require multiple patches, driving added cost and complexity. General availability is expected within three to six months.

Kyndryl is doubling its reserves for cyber-incident recovery in preparation for breaches. Longer term, LLMs like Mythos and other forthcoming models are expected to shift the advantage to defenders, who are positioned to provide the context required to find, fix, and strengthen code.

"When we re-ranked every vulnerability from the worst all the way down, the new chained vulnerabilities clustered right at the very top of the list."

—Participant, Technology Partner Organization



Emergency patching and risk-based prioritization

Traditional monthly patching cadences are no longer viable. Emergency patching is becoming the everyday norm, forcing a shift to risk-based prioritization and far greater automation of testing and deployment.

Newly discovered vulnerabilities are already outpacing remediation capacity. **Vulnerability discovery is outpacing remediation by 60%**, rendering a “patch everything” approach impractical. Organizations must instead aggressively prioritize focusing on risks that are exposed, critical, or actively exploited. Backlogs must be segmented by risk: prioritizing critical, internet-facing, high-severity, and actively exploited systems.

A single individual can now mount a sophisticated attack using an AI model. Dedicated task forces that automate testing and deployment of patches—including self-created or virtual patches—must be able to respond in hours or even minutes. Decommissioning end-of-life systems and hardening easy exposures are essential complements to faster patching.

“We can no longer test everything — users will have to accept that occasionally a minor feature breaks, because we simply have to keep patching.”

—Participant, Manufacturing Sector Organization



04

Maintenance windows and regulatory expectations

Faster patching is colliding with finite maintenance windows and established regulatory expectations, forcing a renegotiation of the trade-offs between service availability and disruption, with both internal stakeholders and with regulators.

While senior leaders and regulators are engaged, current expectations taken literally would require patching everything. This means it is essential to educate regulators and colleagues involved in risk assessment and audit. CISOs warn that faster patching will come with trade-offs, including more frequent maintenance windows, reduced testing, and lower online availability of services, requiring regulators to adapt accordingly.

Approaches vary by sector: Some manufacturers patch during lunch breaks tied to local operating hours, while banks with 24/7 availability requirements patch at least weekly. With modern, high-availability architecture, organizations are questioning the need for maintenance windows. Whether it is safer for systems to shift or gracefully degrade rather than go offline is unclear. Change processes built for a monthly cadence cannot survive the new speed, and vendors must deliver patching that avoids downtime and reboots.

“We told the regulator directly that patching this fast may force more maintenance windows and reduce the availability of our online services.”

—Participant, Banking Sector Organization



05

Resilience across legacy, operational technology

Beyond patching, resilience must be multifaceted as threats move toward machine speed: assume breach, isolate legacy and operational technology, harden third-party dependencies, and revisit Zero Trust.

Detection, response, regular testing, and compensating controls matter as much as remediation. Operational technology is especially difficult — controllers, industrial control systems, robots, and production machines cannot easily be taken offline, they often run unsupported legacy systems, and may be patched only once or twice a year. Isolation and perimeter security are therefore their primary defenses. As a result, some organizations are choosing to delay new third-party library updates during their first week of release.

At scale, strong automation paired with an AI-supported platform is the only viable way to operate at the required speed. Concrete measures include hot standby and ready-to-promote backups for critical cloud systems, and curated, mutually authenticated access for internet-facing services.

“The human factor is moving out of the decision and into the design, because our time to respond is heading toward zero.”

—Participant, Financial Services Organization



06

How to shift to risk-based, automated resilience

These shifts point to one imperative: Stop trying to patch everything, and instead operate on risk. Steps for doing so include:

- Segment vulnerability backlogs by risk and make emergency patching routine, compressing response times to hours or minutes.
- Include self-created or virtual patches where no vendor patch yet exists.
- Decommission out-of-support and end-of-life systems, and strengthen easily addressed exposures rather than relying on patching alone.
- Renegotiate maintenance windows and regulatory expectations openly, assume breach, isolate legacy and operational technology, and renew investment in Zero Trust.

Above all, pair strong automation with an AI-supported platform. It is the only viable way to operate at machine speed.



Benchmarks for AI-era technology leaders:

- Resilience
- Transparency
- Recovery at speed

07

What is the Kyndryl 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. There are CISO Expert Exchange communities in the U.S., Canada and Europe. Participation is open to CISO level executives at large enterprises.

The June 3, 2026 conversation was moderated by Jorgen Floes, VP, Kyndryl's Resiliency & Security Practice and Engineering Lead. Its Kyndryl subject-matter expert was Paul Savill Global Practice Leader Cyber Resilience & Connectivity.

Visit Kyndryl CISO Expert Exchange -->



Jorgen Floes

VP, Kyndryl's Resiliency & Security Practice and Engineering Lead



08

Kyndryl cyber resilience services

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. Learn more at Kyndryl.com.





© Copyright Kyndryl Inc. 2026. All rights reserved.

Kyndryl is a trademark or registered trademark of Kyndryl, Inc. in the United States and/or other countries. Other product and service names may be trademarks of Kyndryl, Inc. or other companies.

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.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions. Kyndryl products and services are warranted according to the terms and conditions of the agreements under which they are provided.