P O Box 427 Pretoria 0001 South Africa
370 Helen Joseph Street Pretoria 0002
+27 12 313 3911 / 0861 12 7272
www.resbank.co.za

South African Reserve Bank
Prudential Authority

Prudential Authority Communication to the Financial Sector: Preparing for Frontier AI and AI-Accelerated Cyber Risk

Purpose and application

This communication applies to all supervised financial institutions, and it is issued to raise awareness of a material shift in the cyber-risk landscape arising from rapid advances in artificial intelligence (AI).

Background

Recent advances in frontier AI are changing the cyber risk landscape by dramatically increasing the speed and scale of attacks, weakening traditional security approaches that rely on periodic scanning, patching cycles, and manual intervention. Many institutions will have seen recent coverage relating to *Anthropic's Claude "Mythos" Preview*, which demonstrated the ability of advanced models to autonomously discover high-impact software vulnerabilities, generate working exploits, and materially compress the time between vulnerability existence and exploitation. While AI-assisted vulnerability discovery and attack automation are not new, frontier AI, such as Mythos-class capabilities represent a significant acceleration in speed, scale, and autonomy. These developments reinforce a trajectory already underway, in which attackers increasingly identify and weaponise weaknesses at machine speed across complex technology environments.

What Frontier AI means for the Financial Sector

Frontier AI materially compresses the time between vulnerability discovery and exploitation and increases the likelihood of correlated cyber events across institutions with shared technologies or dependencies. However, these capabilities can be applied defensively by institutions to strengthen detection, response, and resilience.

It is important to note that these developments do not signal an immediate crisis for the South African financial sector; however, it requires a measured, forward-looking approach for institutions to be prepared. Therefore, the Prudential Authority (PA)'s supervisory focus is shifting from awareness of AI-accelerated cyber risk to execution, operational resilience, and effective decision-making under compressed timeframes.

Supervisory Expectations and Practical Readiness

Institutions are expected to prioritise cyber risk based on exploitability rather than volume, with continuous validation of exposure across applications, dependencies, third-party connections, systems and applications; and automated service accounts. Detection and response capabilities should operate at machine speed through appropriate automation, supported by strong and ongoing identity and access controls. Institutions are expected to apply AI intentionally and securely to enhance investigation, triage, containment, and remediation, as well as to treat N-day vulnerabilities as urgent, particularly in legacy and third-party environments. AI usage across development and operational environments should be visible, governed, and monitored to detect unsafe behaviour early.

Expectations for Boards and Senior Executives

AI-accelerated cyber risk is a Board-level and executive responsibility requiring clear ownership and oversight. Boards and senior management are expected to ensure that decision-rights for containment and recovery are pre-defined; risk appetite is reviewed against compressed attack timelines; and management is capable of detecting and responding to cyber events at the required speed and scale.

Conclusion

The capabilities demonstrated by frontier AI models such as Mythos are real and evolving rapidly. They reinforce a future in which speed, automation, and coordination increasingly determine outcomes. However, the factors that will ultimately determine how institutions withstand these developments remain unchanged; they include visibility across the environment, disciplined cybersecurity hygiene, strong governance, and the ability to detect and respond at speed. Establishing and

maintaining this foundation remains within each institution’s control. It is also the most effective way to remain resilient as AI capabilities continue to evolve.

These developments do not replace or reduce existing regulatory obligations. They reinforce the continued relevance of the Joint Cybersecurity and Cyber Resilience Standard, which remains the cornerstone of regulatory expectations for governance, detection, response, recovery, and timely reporting of incidents. It is important to note that institutions that are significantly aligned with the Joint Standard’s intent and principles will be better positioned to sustain operational continuity, confidence, and financial stability.

The PA will continue to engage with the sector on these developments and will increasingly focus supervisory attention on preparedness, execution under stress, and resilience outcomes, rather than awareness alone.