

Supervisory expectations on cyber risk associated with frontier AI

Remarks delivered to financial services industry leaders in Auckland on 15 October 2026

By Angus McGregor, Assistant Governor Financial Stability (Reserve Bank of New Zealand) and Louise Unger, Executive Director Response and Enforcement (Financial Markets Authority)

With special thanks to Sarah Nash, Deepika Ichaporia and Chelsea Pankhurst



Key messages

- **AI presents significant opportunities as well as risks for the financial sector.** Advances in AI are changing the operating environment for regulated entities, amplifying existing risks and introducing new challenges. Entities should establish a clear and proportionate approach to managing AI-related risks.
- **Improving cyber fundamentals is key.** Regulated entities should maintain strong foundational cyber controls. Frontier AI increases the speed and scale at which vulnerabilities are identified and exploited, particularly in legacy or unsupported technology environments.
- **AI-related risks are testing the effectiveness of governance, oversight and risk management settings.** Cyber security strategy and frameworks should remain fit for purpose. Regulated entities should actively identify, assess and monitor material changes in their cyber and AI risk profile. Boards and senior management should set risk appetite and accountabilities and maintain sufficient visibility.
- **Third-party and concentration risks should be assessed and appropriately managed.** Entities should understand their third-party threat exposure and how arrangements could introduce or amplify threats from AI, particularly in instances of concentration risk. Accountability for managing operational risks remains with regulated entities.
- **Entities should assume incidents will occur and should prepare accordingly.** Incident preparedness should include maintaining response and recovery protocols for AI-enabled attacks and disruptions, supported by severe but plausible scenario testing.
- **This will remain an active supervisory focus for both the Reserve Bank and FMA.** Sector-wide coordination and information sharing will be increasingly important.

Angus McGregor, Assistant Governor Financial Stability (Reserve Bank)

Tēnā koutou katoa.

Thank you for being here this afternoon. We appreciate you taking the time to connect with us today.

The emergence of frontier AI raises important opportunities and challenges for the New Zealand financial sector, so it's fantastic to be here with colleagues from National Cyber Security Centre, Department of the Prime Minister and Cabinet, and the Financial Markets Authority, including Executive Director, Response and Enforcement, Louise Unger, who I will pass to shortly.

As we all know, cyber risk has long been seen as a key operational risk for regulated entities and a focus for both the Reserve Bank of New Zealand (Reserve Bank) from a financial stability perspective and the Financial Markets Authority (FMA) from a conduct perspective. However, the latest releases of frontier AI models, including those released by Anthropic and OpenAI, have

marked a “new phase” in the cyber security threat landscape¹. Frontier AI represents the most advanced models of AI software² and has fundamentally changed the capability, speed and accessibility of the tools that can be used to identify and potentially exploit weaknesses in an increasingly complex technology environment. The NCSC’s Cyber Threat Report 2026³ is clear that New Zealand organisations are operating in a faster-moving and more complex threat environment, with emerging AI capabilities reshaping the cyber landscape.

Internationally, frontier AI has been noted as an “immediate concern” for the global financial system⁴. In June, the Five Eyes security agencies⁵ issued a rare joint “Call to Action” which noted that cyber resilience is integral to advancing business continuity, market confidence and long-term value. One of the reasons for our joint roundtable today with our colleagues from other agencies⁶ is to make it clear that we all agree on the significance of the risk and the need to take action. Cyber risk is a technical issue that has become a core business risk⁷ for regulated entities and a system-wide issue for the financial sector in New Zealand. It also raises the risk for consumers of financial services, including through disruption to essential services, loss of access to accounts, scams and fraud, as well as data breaches. As such, we want to signal that we expect regulated entities to work together, and with us, to share insights to uplift cyber and AI capability across the financial sector and with third-party suppliers to improve resilience across New Zealand. We hope the discussion that follows helps to facilitate this.

AI capabilities are continuing to advance quickly, with recent industry assessments indicating that leading open-weight models are now only months behind frontier AI systems on a range of tasks⁸. As these capabilities become more widely available, and with models improving so rapidly, we expect all regulated entities to be taking steps now to shore up cyber resilience.

We have long expected regulated entities to maintain cyber-security strategies and frameworks that adequately address the cyber threats they face⁹. Today, we would like to take the opportunity to outline our supervisory expectations specifically on the cyber risk associated with the advances in frontier AI. These reflect existing prudential, conduct, cyber and operational resilience obligations and aim to support resilience for regulated entities and the broader financial system. This is an area of growing importance for both Reserve Bank and the FMA, and staying closely connected is critical as we take a system-wide view of these risks and what they mean for consumers as well as New Zealand’s financial sector.

We will focus on how entities should be responding to three key risks arising from frontier AI developments: increased cyber risk, evolving governance and oversight challenges, and growing third-party and concentration risk. While these developments present significant opportunities for the financial sector, they also reinforce the need for entities to be prepared for incidents, whether arising from AI-enabled attacks, model failures or supplier disruptions. As highlighted in recent

¹ [FSI OP 28 When machines attack: frontier AI cyber threats and policy responses in the financial sector](#)

² [Frontier AI: Managing the increasing risks from vulnerabilities](#)

³ [Cyber Threat Report 2026](#)

⁴ [FSB Chair’s letter to August 2026 G20 FMCBG](#)

⁵ [Leaders of Five Eyes Cyber Security Agencies Call to Action on AI Preparedness](#)

⁶ Agencies in attendance: RBNZ, FMA, Department of the Prime Minister and Cabinet (DPMC), National Cyber Security Centre (NCSC)

⁷ [Leaders of Five Eyes Cyber Security Agencies Call to Action on AI Preparedness](#)

⁸ [How Far Behind the Frontier are Leading Open Weight Models on Cyber? | AISI Work](#)

⁹ [Guidance on cyber resilience](#) by RBNZ, [Resources on operational resilience](#) by FMA

NCSC guidance, organisations should assume compromise¹⁰ and ensure they maintain robust response, recovery and testing arrangements.

This will lead into our discussions later on about how we as public sector agencies and the private sector can coordinate and work together in response to a crisis.

Improving cyber fundamentals and AI-enabled defence

In the face of this dynamic risk landscape, we expect regulated entities to maintain strong foundational cyber controls and security posture. This means focussing on getting the basics right: ensuring strong identity and access management, strong asset lifecycle management and timely patching. Our colleagues at the NCSC have published some excellent guidance¹¹ and we expect you all to be considering this in the context of your business.

Frontier AI is making it easier to discover new vulnerabilities and has reduced the time between their discovery and exploitation¹². Given the speed of development, we encourage regulated entities to accelerate their risk-based vulnerability identification, patching and remediation to reflect the speed and scale of AI-enabled attacks. However, we do acknowledge that increasing the pace of patching to respond to recently identified vulnerabilities and reduce any backlog of lower impact vulnerabilities comes with its own risks, including the potential for incidents. Entities should therefore consider the operational risks associated with accelerated remediation and ensure these are appropriately managed.

As part of this, it is worth assessing whether predominantly manual monitoring and triaging processes remain adequate against machine-speed threats or whether automated tools could be deployed to respond faster and reduce the cost and impact of incidents. If considering the use of AI and automation for threat detection or vulnerability management, it is crucial to acknowledge that the effectiveness of AI to support cyber resilience depends on the environment, controls and processes around an AI model that makes its outputs useful, safe and reliable. Of course, for any AI implementation, we would expect regulated entities to establish effective oversight, accountability and control arrangements proportionate to their business model and risk profile.

It is also important here to note that frontier AI increases the risks associated with legacy technology, particularly where legacy systems are out of support. We encourage regulated entities to consider this, not only as part of their overall security posture, but as part of strategic discussions around the scope and timelines of technology transformation programmes.

While cyber fundamentals remain critically important, resilience in an AI-enabled environment is also increasingly influenced by governance, risk management and third-party dependencies. I will now hand over to Louise Unger from the FMA to discuss these areas in more detail.

Louise Unger, Executive Director Response & Enforcement (FMA)

Thank you, Angus.

As Angus has outlined, developments in AI are changing the cyber threat environment and increasing the pace and complexity of risks facing financial institutions.

¹⁰ [Frontier AI: Managing the increasing risks from vulnerabilities](#)

¹¹ [Cyber readiness in the Frontier AI era ; The implications of Frontier AI models for cyber defence](#)

¹² [Frontier AI: Managing the increasing risks from vulnerabilities](#)

From the FMA's perspective, our focus is not only on technology risk itself, but on what those risks mean for consumers, market confidence and the continued delivery of critical services. With that in mind, I'd like to focus on three areas: governance and risk management, third-party risk, and incident preparedness.

Reviewing AI-related governance and risk management settings

Getting the cyber basics right also means reviewing your cyber security strategy and frameworks to ensure these are fit for purpose in the current AI risk environment. Risk from AI is revealing weaknesses not only in technical cyber controls, but also in governance, oversight and risk management settings. In some cases, existing risk management practices may not be designed to operate at the speed and scale required to respond to rapidly evolving threats.

Regulated entities should establish a clear approach to managing AI-related risk that is proportionate and aligned with their business profile and strategy, risk appetite, tolerance levels and operational resilience objectives. They should also identify, assess and monitor material changes in their cyber and AI risk profile, including control weaknesses, incidents, third-party dependencies and emerging exposures. Boards and senior management should maintain sufficient understanding and visibility of material AI-related risks and dependencies to set risk appetite and assign clear accountability for material AI risks. This includes establishing and maintaining board capability to provide effective challenge and ensuring the board have the appropriate expertise and skills to respond to the evolving risk. They can be supported by appropriate independent assurance to confirm that controls are appropriately configured and operating as intended.

We expect that increasingly AI adoption will extend beyond operational processes and cyber defence into areas such as investment research, portfolio management, client communications and financial advice. Used well, these technologies have the potential to drive innovation, improve productivity and support better outcomes for consumers and organisations alike. As these tools become more widely used, boards and senior management should ensure they have sufficient visibility of where AI is being used, the risks it may introduce, and whether governance and control frameworks remain effective. This includes considering potential impacts on operational resilience, consumer outcomes, and the delivery of important products and services.

Understanding and managing third-party risk

Another area we expect regulated entities to assess and appropriately manage is the AI-related risks associated with third-party arrangements. Given the interconnectedness of the financial system and the increasing reliance on a small number of AI and cloud service providers, disruption, outages, model failures or operational issues at these providers could affect multiple institutions simultaneously. As noted in the Reserve Bank's Future of Banking Study, such exposures could increase the risk of operational failures and amplify system-wide vulnerabilities¹³.

Entities should look at third-party threat exposure and how arrangements could introduce or amplify threats from AI, particularly in instances of concentration risk. This could include understanding how third parties use AI within their environments and how that use may affect the regulated entity's threat exposure.

¹³ [future-of-banking-study.pdf](#) pages 54-55

Assessing the resilience of key suppliers, ensuring appropriate controls are in place and establishing contractual arrangements that provide for timely notification of material changes, vulnerabilities and incidents are all important elements of this. Where responsibilities are shared with a provider or group entity, they should be clearly defined, understood and subject to effective oversight. To be clear, accountability for managing operational risks remains with regulated entities.

Incident response and recovery

As well as these regulatory expectations, regulated entities should assume incidents will occur and prepare accordingly. Increasingly, entities should be looking out for AI-enabled fraud and scam activity, including deepfake impersonation, synthetic identities and increasingly sophisticated social engineering attacks. These threats can result in harm to consumers, but may also expose entities to financial loss, operational disruption, reputational damage and erosion of trust. Entities should consider whether existing fraud controls, customer and employee verification processes and awareness programmes remain effective against these evolving threats.

Threat horizon preparations should include maintaining response and recovery protocols for AI-enabled attacks and disruptions, as well as utilising severe but plausible disruption scenarios, such as high-speed attacks, zero-days, model failure and third-party outages to test manual workarounds, backups, failover and recovery.

Given the third-party risk already mentioned, critical third parties could be included where appropriate to ensure any weaknesses are identified and remediated through testing. Entities should also understand their dependencies on key providers and how they would operate if those providers experienced a significant outage, particularly where concentration risk has been identified.

Critically, frontier AI may expose risks that require coordinated engagement and information sharing across the financial sector and with government agencies. We will outline today how we collaborate as regulators and government agencies to both support preparedness and respond effectively during live incidents as well as the role we see the financial sector playing. We're very keen to hear from you on how we, as NZ Inc., can better work together in a crisis to ensure the resilience of New Zealand's financial sector and are looking forward to that part of our discussion today.

Looking to the future

We see this as the beginning of an ongoing conversation. The FMA and Reserve Bank will continue to build on this work together, drawing on insights from our supervisory engagement, thematic reviews and broader industry discussions. We will continue to work with other agencies to understand emerging risks and support the resilience of New Zealand's financial system as AI evolves. We will also draw from the experience of overseas regulators and, where appropriate, consider how their approaches could be applied within New Zealand's context.

We are also seeking to deepen our understanding of how AI is being adopted across the financial sector. This includes the FMA's thematic review of AI in financial advice, which is exploring how entities are using AI-enabled tools, the governance and risk-management arrangements

supporting their use, and the implications for consumer outcomes.¹⁴ It also complements our broader engagement with entities through ongoing operational resilience thematic work, which continues to provide insights into how organisations are managing technology, outsourcing and resilience risks in an increasingly complex environment.¹⁵

The insights gained through these initiatives will help inform our future work and engagement with industry. As part of this, we will continue discussing with regulated entities how they are identifying, assessing and managing AI-related cyber, operational resilience and third-party risk.

Thank you. This is an area that will continue to evolve, and we look forward to ongoing engagement with industry. I'll now hand back to Angus to share the Reserve Bank's perspective on the road ahead.

Angus McGregor, Assistant Governor Financial Stability (Reserve Bank)

Thank you, Louise.

As regulators, our ongoing work and future expectations will continue to support trust and resilience of the financial system in the face of operational risks and cyber threats. We will continue to monitor our regulatory settings to ensure that they remain appropriate for the risks posed by AI.

Frontier AI is continuing to rapidly evolve and is being increasingly widely adopted by the financial sector and malicious actors alike. This brings new opportunities and efficiencies for businesses but also amplifies existing risks and introduces new risks for the sector. While much of today's discussion focuses on cyber threats, entities should also be considering broader implications of AI, including model risk, fraud and scams, concentration risk, consumer outcomes and operational resilience impacts.

The core message today is simple: advances in AI are changing the operating environment for financial institutions. Entities are responsible for ensuring that they have the systems and controls in place to manage the risks arising from AI. The organisations best placed to respond will be those that continue to strengthen their governance, resilience, risk management and preparedness arrangements as technology and threats evolve. We encourage our regulated entities to remain vigilant of the cyber risks amplified by developments in frontier AI and advisories issued by our security agencies. For our part, operational resilience is and remains a key supervisory focus area for both the Reserve Bank and the FMA, so you can expect us to continue to engage with regulated entities on this topic, so that we can be assured you are managing the risks appropriately.

We look forward to continuing this conversation with industry and building our collective understanding of these issues over time.

Ngā mihi. Tēnā koutou katoa.

Thank you.

¹⁴ [Thematic review: Artificial intelligence in financial advice | Financial Markets Authority](#)

¹⁵ [Operational resilience thematic: findings and insights | Financial Markets Authority](#)