MAPPING BRIEF

# Using the Zscaler Zero Trust Exchange to Meet APRA's Supervisory Expectations on Artificial Intelligence

©2026 Zscaler, Inc. All rights reserved.

# Executive Summary

### Responding to APRA's Call for a Step-Change in AI Risk Management

On 30 April 2026, APRA published a letter to all regulated entities setting out the findings of a targeted supervisory engagement with selected large banks, insurers and superannuation trustees conducted in late 2025. APRA found that every entity it engaged with was actively adopting AI, but that governance, risk management, assurance and operational resilience practices were not keeping pace with the scale, speed and complexity of the technology.

The letter is APRA's first published set of AI-specific expectations. APRA's prudential framework remains principle-based and technology and vendor agnostic, yet the letter sets out, observation area by observation area, what APRA expects of Boards and accountable executives across information security, governance, supplier risk, and change management and assurance. APRA also warned that high-capability frontier AI models are expected to increase the probability, speed and scale of cyber attacks.

> "While we are not proposing to introduce additional requirements at this stage, we expect to see a significant improvement in how entities are closing the gaps between the power of the technology they are using and their ability to monitor and control it."
>
> Therese McCarthy Hockey, APRA Member, 30 April 2026

Regulators across the region are aligned. ASIC's open letter of 8 May 2026 urged licensees to act now to strengthen cyber resilience against AI-enabled threats. On 15 October 2026, the Reserve Bank of New Zealand and the Financial Markets Authority outlined joint supervisory expectations on cyber risk associated with frontier AI, focused on increased cyber risk, governance and oversight, and third-party and concentration risk.

Many of the gaps APRA identified are, at their core, visibility and control problems: AI tools adopted outside approved frameworks, controls that detect rather than prevent, identity models that have not adjusted to AI agents, and AI supply chains that are opaque. These are problems the Zscaler Zero Trust Exchange was designed to address. Because it sits inline between users, devices, workloads and agents and the applications and AI services they use, it can discover AI usage, enforce preventative policy, protect data in prompts and uploads, broker agent access on a least-privilege basis, and generate continuous evidence for assurance.

©2026 Zscaler, Inc. All rights reserved.

This guide is designed for risk, compliance, security and technology leaders at APRA-regulated entities. It summarises the regulatory expectations, maps each APRA expectation to supporting Zscaler capabilities, and makes explicit the responsibilities that remain with the regulated entity.

## What APRA Observed

APRA's letter is accompanied by an executive debrief for accountable executives, including Chief Risk Officers, Chief Technology Officers and Chief Information Security Officers. It is organised around four observation areas.

### 1. Information Security

AI threats are increasing, but information security practices are struggling to keep pace. APRA highlighted new attack pathways such as prompt injection, data leakage, insecure integrations and exploit injection, and noted that:

- Identity and access management has not yet adjusted to non-human actors such as AI agents
- The volume and speed of AI-assisted coding is straining change and release controls
- Security testing has gaps in scope and coverage
- Shadow AI is managed through policy direction and detective, after-the-fact measures rather than preventative technical controls

### 2. Governance

AI adoption is moving fast, but governance maturity is lagging. Use is shifting from internal productivity to customer-facing processes such as claims triage and loan processing, yet few entities have operationalised AI governance, and AI is often treated as "just another technology".

### 3. Supplier Risk

Supplier risk management is in place, but concentration and opacity present challenges. Some entities rely heavily on a single provider for many AI use cases, with weak contingency, exit and substitution planning. AI embedded in software platforms and developer tools obscures upstream dependencies on foundation models, training data and fourth parties.

### 4. Change Management and Assurance

Point-in-time, sample-based assurance is not sufficient for probabilistic models that learn, adapt and degrade. Few entities continuously monitor for drift, bias, failure modes or control breakdowns, and second-line risk and internal audit often lack the specialist skills and tools to assess agentic behaviour and AI-assisted code.

### Board Expectations

At a minimum, APRA expects Boards to maintain sufficient AI understanding and literacy to set strategic direction and provide effective challenge, and to oversee an AI strategy consistent with risk appetite and tolerances, supported by monitoring and reporting (including on third-party dependencies) with clearly defined triggers aligned to resilience objectives.

©2026 Zscaler, Inc. All rights reserved.

### What Comes Next

APRA is finalising a forward plan for supervising AI risk, combining proportionate entity reviews, thematic activities and engagement with AI suppliers, and will consider whether further policy action is needed. Legal commentators, including Clayton Utz and MinterEllison, read the letter as a meaningful shift in supervisory posture, with stronger supervisory action and enforcement in reserve for entities that fail to manage AI risk proportionately.

## A Trans-Tasman View

APRA's letter is part of a consistent regional message. Regulators in Australia and New Zealand agree that frontier AI changes the threat environment, but not the fundamentals of good risk management.

### ASIC: Act Now on Cyber Resilience

ASIC's open letter of 8 May 2026 to AFS licensees and market participants called for urgency, focus and accountability, framing cyber resilience as a core licensing obligation rather than an IT issue. ASIC's suggested steps include reassessing cyber plans and governance, protecting critical assets, strengthening fundamentals such as access control and patching, actively managing third-party risk, preparing for incident response and using AI for defence. ASIC asked licensees to table the letter at Board and risk committees.

### RBNZ and FMA: Getting the Basics Right

In remarks delivered on 15 October 2026, the RBNZ and FMA set out expectations on three risks arising from frontier AI: increased cyber risk, evolving governance and oversight challenges, and growing third-party and concentration risk. Key messages include:

- Strong foundational controls: identity and access management, asset lifecycle management and timely patching
- Accelerated, risk-based vulnerability remediation, and consideration of whether manual monitoring remains adequate against machine-speed threats
- Heightened risk from legacy and out-of-support technology
- Visibility of where AI is used, the risks it brings and who is accountable
- Understanding how third parties use AI and how that affects the entity's threat exposure
- Assume incidents will occur, and test severe but plausible scenarios such as high-speed attacks, zero-days, model failure and third-party outages
- Vigilance against AI-enabled fraud, including deepfake impersonation and synthetic identities

> "Boards and senior management need to know where AI is being used in their business, understand the risks it brings, and be clear about who is accountable for managing them."
>
> Louise Unger, Executive Director, FMA, 15 October 2026

©2026 Zscaler, Inc. All rights reserved.

# How Zscaler Secures AI

The Zscaler Zero Trust Exchange is a cloud-native security platform that connects users, devices, workloads and AI agents directly to applications based on identity and policy, without exposing the network. Zscaler extends the same inline architecture to AI across four layers.

### Secure the Use of AI

- Discovery and control of public, sanctioned and embedded GenAI applications
- Prompt and upload inspection with inline Data Loss Prevention, including TLS and WebSocket inspection
- Coaching, isolation and tenant restrictions for safe enablement rather than blanket blocking

### Secure the AI You Build and Run

- AI Asset Management (AI-SPM) and AI Bill of Materials across SaaS, cloud AI platforms, models and MCP servers
- AI Red Teaming for automated adversarial testing of AI applications and agents
- AI Guard runtime guardrails against prompt injection, jailbreaks and data leakage

### Secure AI Agents

- AI Broker and Agent Registry for least-privilege brokering of MCP and agent-to-agent traffic
- AI Access Graph mapping human and agent identities to data
- Endpoint AI Security for AI coding tools, local models and agent runtime activity

### Defend at Machine Speed

- Zero trust access that removes applications from the internet attack surface
- Deception, Agentic SOC, Unified Vulnerability Management and Risk360 for detection, containment and Board-level risk reporting

| | |
|---|---|
| 23% | of enterprise AI/ML transactions came from the finance and insurance sector, the largest of any industry |
| +83% | year-on-year growth in enterprise AI/ML transactions in 2025 |
| 100% | of enterprise AI systems tested by Zscaler red teams had critical flaws |

Source: Zscaler ThreatLabz 2026 AI Security Report (2025 data)

©2026 Zscaler, Inc. All rights reserved.

## How To Use This Guide

The following pages map each expectation in APRA's April 2026 letter to supporting capabilities in the Zscaler Zero Trust Exchange. Expectations are grouped by APRA's observation areas, with a final section on frontier AI and cyber fundamentals that also reflects the ASIC and RBNZ/FMA messages.

Zscaler provides technical controls, visibility, telemetry and architectural enforcement that support these expectations. Strategy, risk appetite, accountability, human oversight of high-risk decisions, contractual arrangements and regulatory engagement remain with the APRA-regulated entity.

Some capabilities described were announced in 2026 and are rolling out progressively. Availability and licensing vary by capability; confirm the current status with your Zscaler account team.

## Zscaler Product Terminology

| Industry Terminology | Zscaler Capability |
|---|---|
| Security Service Edge (SSE) | Zscaler Zero Trust Exchange |
| Secure Web Gateway | Zscaler Internet Access (ZIA) |
| Zero Trust Network Access | Zscaler Private Access (ZPA) |
| GenAI app control and prompt DLP | Secure Access to AI |
| AI security posture management | AI Asset Management (AI-SPM) |
| LLM runtime guardrails | Zscaler AI Guard |
| AI red teaming | Zscaler AI Red Teaming |
| MCP / agent gateway | AI Broker and Agent Registry |
| Digital experience monitoring | Zscaler Digital Experience (ZDX) |
| Cyber risk quantification | Zscaler Risk360 |
| Deception technology | Zscaler Deception |
| AI-driven security operations | Zscaler Agentic SOC |

©2026 Zscaler, Inc. All rights reserved.

# Regional Supervisory Themes at a Glance

## Australian and New Zealand Regulator Expectations on AI and Cyber Risk

| Theme | APRA<br>Letter to Industry, 30 Apr 2026 | ASIC<br>Open letter, 8 May 2026 | RBNZ and FMA<br>Supervisory expectations, 15 Oct 2026 |
|---|---|---|---|
| **Cyber fundamentals** | Step-change in cyber practices; PAM, patching, hardened configuration, testing and controls over agentic workflows | Back-to-basics resilience: critical assets, access reviews, patching, defence in depth | "Get the basics right": IAM, asset lifecycle and timely patching; legacy technology risk |
| **Governance and Board oversight** | Board AI literacy; AI strategy aligned to risk appetite; inventory, lifecycle ownership and human oversight | Boards to engage directly and seek evidence; letter tabled at Board and risk committees | Clear, proportionate approach; Board visibility, risk appetite and accountability for AI risk |
| **Third-party and concentration** | Map the full AI supply chain; transparency and auditability; credible exit and substitution | Actively manage third-party risk | Understand third-party AI use and threat exposure; concentration risk; timely notification |
| **Assurance and monitoring** | Continuous, integrated assurance; skilled second and third lines; lifecycle risk assessment | Evidence over assurances: testing, audit findings and lessons from incidents | Identify, assess and monitor material changes; independent assurance |
| **Incident preparedness** | Credible fallback where AI supports critical operations | Prepare for incident response; use AI for defence | Assume incidents will occur; severe but plausible scenario testing including model failure |

©2026 Zscaler, Inc. All rights reserved.

# APRA Letter to Industry on Artificial Intelligence

## Expectation Mapping to Zscaler Zero Trust Exchange

### APRA AI Requirement Mapping

The following pages map the expectations in APRA's April 2026 letter to supporting capabilities in the Zscaler Zero Trust Exchange. Expectations are paraphrased; refer to the APRA letter for the authoritative wording.

- **APRA Expectation:** A reference and summary of the expectation or observation in APRA's letter.
- **Entity Responsibility:** Strategy, risk appetite, accountability, oversight and decisions that remain with the regulated entity.
- **Zscaler Alignment:** Technical controls, visibility, telemetry and enforcement provided by the Zero Trust Exchange.

This guide is intended to support control design and evidence discussions. It does not replace an entity's own legal, regulatory or prudential assessment.

| APRA Expectation | | Shared Responsibility & Zscaler Alignment | |
|---|---|---|---|
| **Ref** | **Expectation** | **Entity Responsibility** | **Zscaler Alignment** |
| **BOARD OVERSIGHT** | | | |
| BRD-01 | Board maintains sufficient AI understanding and literacy to set strategic direction and provide effective challenge and oversight. | Board education programs, director capability and independent challenge of management remain entity-owned. | APRA observed an over-reliance on vendor presentations. Zscaler helps Boards ground challenge in the entity's own evidence rather than vendor narratives:<br>• **AI usage visibility:** Dashboards drawn from inline telemetry show which AI applications are in use, by whom, with what data, and what was blocked or coached.<br>• **Risk360:** Quantifies cyber risk across Zscaler and third-party data, including estimated financial exposure, and exports Board-ready reporting.<br>• **Threat research:** Zscaler ThreatLabz research on AI usage and AI-enabled attacks can inform Board education on the evolving threat landscape. |
| BRD-02 | Board oversees an AI strategy consistent with risk appetite and tolerances, supported by monitoring and reporting (including third-party dependencies) and triggers aligned to resilience objectives. | AI strategy, risk appetite statements, tolerance levels and trigger definitions remain entity-owned. | **Turning Risk Appetite into Enforceable Policy**<br>Zscaler translates appetite decisions into technical policy that is enforced consistently for every user, device and location:<br>• **Graduated controls:** Allow, coach, isolate or block AI applications by risk attribute, user group and data sensitivity.<br>• **Key risk indicators:** AI adoption, data volumes sent to AI applications, DLP violations and policy exceptions can be reported continuously and streamed to SIEM and GRC tooling.<br>• **Dependency monitoring:** Zscaler Digital Experience (ZDX) measures the availability and performance of critical SaaS and AI services from the user's perspective, supporting resilience triggers. |

©2026 Zscaler, Inc. All rights reserved.

| APRA Expectation | | Shared Responsibility & Zscaler Alignment | |
|---|---|---|---|
| **Ref** | **Expectation** | **Entity Responsibility** | **Zscaler Alignment** |
| **INFORMATION SECURITY** | | | |
| SEC-01 | Assess the implications of AI reliance for operational resilience and business continuity, with credible fallback processes where AI supports critical operations. | Business impact analysis, fallback processes, tolerance levels and BCP testing remain entity-owned. | • **Dependency mapping:** AI Asset Management builds an AI Bill of Materials showing which AI services, models and integrations are in use, helping identify where critical operations depend on AI.<br>• **Early detection of disruption:** ZDX helps isolate whether a disruption originates in a device, network, ISP or the AI or SaaS provider.<br>• **Policy-based fallback:** Access policy can redirect users from an impaired AI service to an approved alternative or a non-AI process.<br>• **Resilient security controls:** Zscaler's business continuity architecture is designed to keep security enforced, rather than bypassed, during disruption. See the Zscaler CPS 230 companion guide. |
| SEC-02 | Security controls address AI-specific threats, including privileged access management, timely patching, hardened configurations, automated vulnerability discovery, penetration testing, and controls over agentic and autonomous workflows. | Vulnerability scanning, patching, configuration hardening, penetration testing and PAM tooling remain entity-owned. | **Reducing the Attack Surface**<br>• **Zscaler Private Access (ZPA):** Applications are not exposed to the internet and users connect to applications rather than networks, removing a key target of AI-accelerated reconnaissance and limiting lateral movement.<br>• **Unified Vulnerability Management:** Prioritises remediation using exposure context, such as whether an asset is already protected behind ZPA.<br>**Controls for Non-Human Actors**<br>APRA noted that identity and access management has not adjusted to AI agents. Zscaler extends zero trust to agents:<br>• **AI Broker and Agent Registry:** Brokers MCP and agent-to-agent traffic with identity-bound, least-privilege policy and logging (announced June 2026).<br>• **AI Access Graph:** Maps which human and agent identities can access which models, applications and data, surfacing over-privileged access.<br>• **Endpoint AI Security:** Inventories AI agents, MCP servers and skills on devices, monitors agent runtime activity and blocks dangerous actions (rolling out).<br>• **Workload egress control:** Restricts and inspects outbound traffic from cloud AI workloads and agents. |
| SEC-03 | Robust security testing of AI-generated code, software components and libraries. | Secure SDLC, code review, application security testing and release controls remain entity-owned. | Zscaler complements, rather than replaces, application security testing tools:<br>• **AI code and agent scanning:** AI Asset Management scans code repositories and cloud environments for risky AI usage, agents and MCP servers.<br>• **AI Red Teaming:** Automated adversarial testing of AI applications, agents and MCP servers across thousands of attack scenarios, before deployment and continuously.<br>• **Developer environment protection:** Inline DLP prevents source code and secrets being sent to unsanctioned coding assistants, and Endpoint AI Security inventories packages, plugins and model files on developer devices. |
| SEC-04 | Shadow AI managed through preventative technical controls, not only policy direction or detective, after-the-fact measures. | Acceptable use policy, approved tool lists and the calibration of experimentation to risk appetite remain entity-owned. | **Preventative Control at the Point of Use**<br>Because Zscaler sits inline, it can enforce policy before data leaves the organisation:<br>• **Discovery:** Identifies thousands of GenAI applications, including AI features embedded in sanctioned SaaS platforms. |

©2026 Zscaler, Inc. All rights reserved.

| APRA Expectation | | Shared Responsibility & Zscaler Alignment | |
|---|---|---|---|
| **Ref** | **Expectation** | **Entity Responsibility** | **Zscaler Alignment** |
| | | | • **Prompt-level DLP:** Inspects prompts and uploads, including over WebSockets, and blocks sensitive data such as customer information, source code and credentials.<br>• **Safe enablement:** Coaching pages steer users to approved tools; Browser Isolation allows read-only use without copy, paste or upload; tenant restrictions separate corporate from personal accounts.<br>• **Out-of-band coverage:** Integration with enterprise AI compliance APIs extends visibility to sanctioned AI tenants. |
| SEC-05 | Controls for AI-specific attack pathways such as prompt injection, data leakage, insecure integrations and exploit injection. | Application design, model selection and integration architecture remain entity-owned. | • **AI Guard:** Runtime guardrails between AI applications and models, with detectors for prompt injection, jailbreaks, sensitive data, malicious URLs and toxic content, deployed inline or via API.<br>• **Closed-loop hardening:** AI Red Teaming findings can be converted into AI Guard policy, closing the gap between testing and enforcement.<br>• **Inline threat prevention:** TLS inspection, sandboxing and threat protection apply to traffic between AI integrations and the internet. |
| **GOVERNANCE** | | | |
| GOV-01 | Frameworks (policy, standard, guidance) and reporting lines for safe, responsible and sustainable AI adoption. | AI policy frameworks, reporting lines and Board approvals remain entity-owned. | Zscaler makes AI policy enforceable and measurable. Acceptable use rules become technical policy applied consistently across head office, branches, remote staff and cloud workloads, with reporting that shows whether the policy is working in practice. |
| GOV-02 | Ownership and accountability across the AI lifecycle, from design and deployment through to monitoring and decommissioning. | Assignment of owners and accountable persons, including under the Financial Accountability Regime, remains entity-owned. | • **Lifecycle visibility:** Continuous discovery surfaces new AI tools, agents and integrations as they appear, so owners can be assigned before they become embedded.<br>• **Decommissioning:** Retired or prohibited AI tools can be blocked centrally, confirming decommissioning is effective.<br>• **Agent accountability:** The Agent Registry associates agents with defined permissions, supporting clear ownership of autonomous workflows. |
| GOV-03 | An inventory of AI tooling and use cases. | Use-case registers, risk classification and approval decisions remain entity-owned. | **A Continuously Updated AI Inventory**<br>AI Asset Management provides a living inventory rather than a point-in-time survey, covering:<br>• GenAI SaaS and AI features embedded in sanctioned applications<br>• Cloud AI platforms such as Amazon Bedrock, Azure AI Foundry and Google Vertex AI<br>• Models, MCP servers, agents and AI pipelines, with an AI Bill of Materials<br>• Local AI tools, coding assistants and models on endpoints, via Endpoint AI Security |
| GOV-04 | Human involvement and accountability for high-risk decisions. | Design of human-in-the-loop controls and decision accountability remains entity-owned. | Zscaler supports human oversight by constraining and recording what AI can do:<br>• **Bounded autonomy:** AI Broker policy can limit which tools and data an agent may use, so high-risk actions require a human.<br>• **Audit trail:** Full conversation views, prompt and response replay, and agent-run reconstruction on endpoints provide evidence of how AI was used in a decision. |
| GOV-05 | Structured training and education of staff on AI use, misuse, limitations | Training content, delivery and completion tracking | Zscaler reinforces training at the moment of risk. Coaching pages explain policy when a user reaches an unsanctioned AI tool or attempts to share sensitive data, and usage analytics help target training to the teams and |

©2026 Zscaler, Inc. All rights reserved.

| APRA Expectation | | Shared Responsibility & Zscaler Alignment | |
|---|---|---|---|
| **Ref** | **Expectation** | **Entity Responsibility** | **Zscaler Alignment** |
| | and secure practices. | remain entity-owned. | behaviours that need it most. |
| **SUPPLIER RISK** | | | |
| SUP-01 | Map and maintain visibility of the full AI supply chain, including material third-party and fourth-party dependencies. | Supplier registers, materiality assessments and CPS 230 service provider management remain entity-owned. | • **Embedded AI detection:** Identifies AI functionality inside sanctioned SaaS platforms, one of the opacity issues APRA highlighted.<br>• **AI supply-chain risk:** AI-SPM assesses models and AI platforms for posture and supply-chain risk.<br>• **SaaS posture:** SaaS Security Posture Management detects AI features enabled in SaaS tenants and configuration drift. |
| SUP-02 | Contractual and governance arrangements provide transparency, auditability and assurance. | Contract negotiation, audit rights and notification terms remain entity-owned. | Where AI providers offer limited transparency, entity-controlled telemetry can act as a compensating control. Zscaler records what data flows to each AI provider, by whom and under which policy, independent of the provider's own reporting.<br>As a provider, Zscaler supports assurance through independent certifications and attestations, including IRAP assessment, ISO/IEC 27001 and SOC 2 Type II, available via the Zscaler Trust Portal. |
| SUP-03 | Ability to understand model behaviour, material changes, performance issues and outcomes across the service lifecycle. | Model risk management and outcome monitoring remain entity-owned. | • **Continuous red teaming:** Re-testing after provider model updates helps detect behavioural change that was not notified.<br>• **Runtime monitoring:** AI Guard detector trends show changes in model outputs over time.<br>• **Performance:** ZDX tracks availability and responsiveness of AI and SaaS services. |
| SUP-04 | Actively manage concentration risk, including plausible and systemic failure scenarios, and the credibility of substitution, portability or exit for critical AI providers. | Concentration risk assessment, exit planning and scenario testing remain entity-owned. | • **Measuring concentration:** Usage telemetry shows how dependent the organisation is on each AI provider across teams and use cases.<br>• **Supporting substitution:** A single policy layer across all AI providers means switching or adding an approved provider does not require rebuilding security controls.<br>• **Zscaler as a provider:** Where Zscaler is a material service provider, its resilience architecture and exit support are described in the Zscaler CPS 230 companion guide. |
| **CHANGE MANAGEMENT AND ASSURANCE** | | | |
| ASR-01 | Globally recognised control frameworks, control libraries and change control for AI. | Framework selection and control library design remain entity-owned. | Zscaler AI Red Teaming includes built-in testing aligned to frameworks such as the NIST AI Risk Management Framework, the OWASP Top 10 for LLM Applications and the EU AI Act, with compliance reporting to support ISO/IEC 42001-style management systems. |
| ASR-02 | Integrated assurance across cyber, data governance, model performance, operational resilience, privacy and conduct. | Assurance planning and the three lines model remain entity-owned. | One platform generates correlated evidence across security, data protection, AI usage and digital experience, reducing the fragmentation APRA observed. Telemetry can be streamed to SIEM, SOAR and GRC platforms to support an integrated assurance view. |
| ASR-03 | Second-line risk and internal audit have the technical skills and tools to independently assess AI, including probabilistic models and agentic workflows. | Capability development, resourcing and independence of second and third lines remain entity-owned. | Zscaler gives assurance functions practical tools rather than sample-based snapshots:<br>• Read-only, role-based access to AI usage, DLP and policy evidence<br>• Prompt and response replay and full conversation views<br>• Agent-run reconstruction showing tools, data and actions<br>• Red team reports and posture findings for AI applications |

©2026 Zscaler, Inc. All rights reserved.

| APRA Expectation | | Shared Responsibility & Zscaler Alignment | |
|---|---|---|---|
| **Ref** | **Expectation** | **Entity Responsibility** | **Zscaler Alignment** |
| ASR-04 | Comprehensive risk and security assessments before deployment and throughout the lifecycle, with continuous monitoring proportionate to criticality. | Risk assessment methodology and criticality ratings remain entity-owned. | • **Before deployment:** AI Red Teaming and posture assessment evaluate new AI applications and agents.<br>• **In production:** AI Guard, continuous AI-SPM posture monitoring and inline DLP provide ongoing control and detection.<br>• **Proportionality:** Policy can apply stricter controls and monitoring to higher-criticality use cases. |
| **FRONTIER AI AND CYBER FUNDAMENTALS** | | | |
| FRT-01 | A step-change in cyber security practices in response to high-capability frontier models, including faster identification and patching of vulnerabilities. | Patching cadence, remediation of legacy technology and transformation planning remain entity-owned. | APRA, ASIC and the RBNZ/FMA all emphasise strengthening fundamentals as frontier AI shortens the time between discovery and exploitation:<br>• **Remove what attackers can find:** ZPA hides private applications from the internet and replaces exposed VPN concentrators, reducing exposure while patching catches up.<br>• **Prioritise what matters:** Unified Vulnerability Management focuses remediation on exposed, high-impact assets.<br>• **Detect machine-speed attackers:** Deception places decoys on likely attack paths; automated AI-driven reconnaissance is likely to touch them, producing high-fidelity alerts.<br>• **Contain quickly:** Agentic SOC capabilities support investigation and inline containment at machine speed. |
| FRT-02 | Assume incidents will occur, maintain response and recovery protocols for AI-enabled attacks, and test severe but plausible scenarios (RBNZ/FMA). | Incident response plans, crisis management, regulatory notification and scenario design remain entity-owned. | • **Forensic evidence:** Detailed transaction and session logs support investigation and regulatory reporting.<br>• **AI-enabled fraud and phishing:** Inline threat prevention and Browser Isolation reduce exposure to AI-generated phishing and impersonation lures.<br>• **Scenario testing:** Business continuity failover, AI provider outage and agent compromise scenarios can be exercised using Zscaler policy and telemetry. |

## References


1. APRA, Letter to Industry on Artificial Intelligence (AI), 30 April 2026 – apra.gov.au
2. APRA, Media release: APRA calls for a step-change in AI-related risk management and governance, 30 April 2026
3. ASIC, Open letter to AFS licensees and market participants on cyber resilience, 8 May 2026
4. RBNZ and FMA, Supervisory expectations on cyber risk associated with frontier AI, 15 October 2026
5. Clayton Utz, APRA's AI letter: a shift from framework to targeted expectations, 5 May 2026
6. MinterEllison, APRA's AI letter: A wake-up call for managing your third-party suppliers, 11 May 2026
7. Mallesons, ASIC and APRA issue call to action on artificial intelligence, 19 May 2026
8. Zscaler ThreatLabz, 2026 AI Security Report, January 2026


©2026 Zscaler, Inc. All rights reserved.

# Next Steps

Zscaler works with APRA-regulated entities across Australia and New Zealand to secure AI adoption with zero trust. To discuss how Zscaler can support your response to APRA's expectations on AI, contact your Zscaler account team or visit zscaler.com/compliance.

### AI Visibility Baseline

Use existing Zero Trust Exchange telemetry to establish which AI applications are in use, by whom and with what data, supporting APRA's inventory and shadow AI expectations.

### AI Security Architecture Workshop

Map your AI use cases and agentic roadmap to preventative controls across users, applications, agents and workloads.

### Additional Resources

- **Zscaler Trust Site:** Real-time platform availability and performance metrics — trust.zscaler.com
- **CPS 230 and CPS 234 Companion Guides:** Zscaler mapping briefs for Operational Risk Management and Information Security
- **Zscaler ThreatLabz 2026 AI Security Report:** Research on enterprise AI usage and AI-enabled threats

**CONTRIBUTORS**

Thomas Ferguson and Paul Krasey – Zscaler FSI Specialist team

**DISCLAIMER**

This document is provided for informational purposes only and does not constitute legal, regulatory, or prudential advice. It summarises publicly available regulatory publications and third-party commentary; refer to the original publications for authoritative wording. APRA-regulated entities remain responsible for determining how they meet APRA's expectations and their obligations under the prudential framework. Some Zscaler capabilities described were announced in 2026 and availability and licensing may vary; confirm current status with your Zscaler account team.

zscaler™ | Experience your world, secured.™

**About Zscaler**

Zscaler (NASDAQ: ZS) accelerates digital transformation so that customers can be more agile, efficient, resilient, and secure. The Zscaler Zero Trust Exchange protects thousands of customers from cyberattacks and data loss by securely connecting users, devices, and applications in any location. Distributed across more than 150 data centers globally, the SASE-based Zero Trust Exchange is the world's largest inline cloud security platform. Learn more at **zscaler.com** or follow us on Twitter **@zscaler**.

© 2026 Zscaler, Inc. All rights reserved. Zscaler™, Zero Trust Exchange™, Zscaler Internet Access™, ZIA™, Zscaler Private Access™, ZPA™, Zscaler Digital Experience, and ZDX™,and other trademarks listed at **zscaler.com/legal/trademarks** are either (i) registered trademarks or service marks or (ii) trademarks or service marks of Zscaler, Inc. in the United States and/or other countries. Any other trademarks are the properties of their respective owners.