INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-08 19:05 UTC

ASIC Issues AI-Driven Cyber Resilience Warning to Financial Sector

GOVERNANCE | HIGH

SCC Item ID	SCC-GOV-2026-0031
Type	Governance
Severity	HIGH
Affected Products	Financial firms operating under ASIC jurisdiction (Australia); broadly applicable to global financial services sector
Published	2026-05-08
Discovery Source	Gemini

Executive Summary

Australia's financial sector regulator, ASIC, has formally warned regulated firms that AI-enabled attack techniques, including deepfake social engineering, AI-assisted phishing, and automated vulnerability exploitation, have materially raised the risk profile for financial institutions. The advisory applies directly to ASIC-regulated entities and signals a broader regulatory expectation that existing cyber resilience plans must be updated to address AI-augmented threats. Firms that do not revisit their controls and resilience frameworks face both increased breach exposure and potential regulatory scrutiny for failing to meet current supervisory expectations.

Technical Analysis

This is a sector-wide governance advisory, not a CVE-linked vulnerability disclosure. No specific software flaw, affected version, or patch is cited. ASIC's warning targets AI-augmented attack techniques. Relevant MITRE ATT&CK correlations include: T1589 (Gather Victim Identity Information), T1650 (Acquire Access), T1566 (Phishing), and T1588 (Obtain Capabilities). The threat model includes: AI-assisted spear phishing with high-fidelity lure generation; deepfake audio and video for business email compromise and voice-based social engineering; adversarial use of large language models to accelerate reconnaissance and initial access; and automated scanning and exploitation pipelines that compress the window between vulnerability disclosure and active exploitation. No CVE, CWE, or CVSS score applies. The risk is systemic and process-level, not tied to a patchable software flaw. Foundational controls cited implicitly include MFA, privileged access management, network segmentation, and incident response plan currency.

Action Checklist

1. Step 1: Governance Review, Pull your current cyber resilience plan and mark every section that references threat actor capabilities or attack techniques. Flag any section written before 2024 that does not account for AI-augmented phishing, deepfake social engineering, or automated exploitation. Schedule a review cycle within 30 days.
2. Step 2: Detection, Audit your email security gateway logs for signs of high-volume, low-variance phishing campaigns that suggest AI-generated lures. Review identity and access logs for anomalous authentication patterns consistent with credential harvesting (T1589). If you run voice-based authentication or executive wire-transfer approval processes, assess whether deepfake audio detection controls are in place.
3. Step 3: Eradication, Identify and close process gaps that rely on voice or email confirmation alone for high-value transactions. Enforce phishing-resistant MFA (hardware tokens or passkeys) on all privileged and finance-adjacent accounts. Remove or restrict access paths that have no compensating control against automated credential acquisition (T1650).
4. Step 4: Recovery, Validate that incident response playbooks include decision trees for suspected deepfake-facilitated fraud and AI-assisted phishing campaigns. Confirm contact trees for executive impersonation scenarios are current and tested. Run a tabletop exercise against an AI-augmented BEC scenario within 60 days.
5. Step 5: Post-Incident, Document control gaps identified during the review. For ASIC-regulated firms, ensure documentation of your review and any remediation actions is retained as evidence of regulatory responsiveness. Align updated controls to APRA CPS 234 requirements. For additional framework guidance, consider mapping to NIST CSF 2.0 Govern and Protect functions and relevant NIST SP 800-53 controls (SA-11, SI-3, AC-2).

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to executive leadership, legal counsel, and (for ASIC-regulated entities) compliance officers if: (1) any detection audit in Step 2 surfaces evidence of an active or completed AI-phishing or deepfake-facilitated wire transfer fraud event, (2) identity logs reveal successful credential compromise on finance or executive accounts, or (3) a review gap is identified that constitutes a material deficiency under ASIC's stated regulatory expectations — triggering mandatory regulatory notification obligations under Australian Corporations Act cybersecurity provisions and ACSC reporting thresholds.
Recovery Notes	Following remediation of process gaps and MFA enforcement, monitor IdP authentication logs daily for a minimum of 30 days for re-attempted credential stuffing or anomalous authentication patterns on newly hardened accounts, as threat actors using automated tooling will retry against the same target population. Verify that all wire-transfer approval workflows have been tested end-to-end with the new out-of-band verification procedures before re-enabling any previously restricted transaction paths. Retain all remediation evidence (MFA enforcement records, playbook versions, tabletop outputs) in a durable, access-controlled location for a minimum of 7 years in line with ASIC's record-keeping expectations under RG 255 and the Corporations Act, or per your documented data retention policy if stricter.

Forensic Artifacts	Email security gateway disposition logs (90-day lookback) filtered for campaigns exhibiting AI-generation signatures: high body-text similarity across recipients, low DKIM/SPF pass rates, lookalike domains targeting executive or finance staff, and click telemetry on lures referencing wire transfers or executive directives — the specific artifact pattern of AI-assisted phishing at scale targeting financial institutions | Identity provider or Active Directory authentication logs — specifically Event IDs 4625 (Failed Logon), 4648 (Explicit Credential Use), 4768/4771 (Kerberos pre-auth failure), and Azure AD Risky Sign-In alerts — filtered to accounts in finance, treasury, and executive assistant roles, which are the primary credential harvesting targets in AI-augmented BEC campaigns (T1589) | Voice and telephony platform call records (Microsoft Teams call records, Cisco CUCM CDRs, or carrier-level telephony logs) for inbound calls to executives and finance staff in the 30-day window preceding any anomalous wire transfer approvals — the specific artifact class that would evidence deepfake audio social engineering attempts | Wire transfer and payment system transaction logs showing approval chain metadata (who approved, via what channel, at what time) for all high-value transactions above your organization's defined threshold in the past 90 days — the primary forensic artifact for confirming whether any AI-augmented BEC attempt succeeded in reaching the payment authorization stage | Cyber resilience plan and IR playbook document metadata (file creation date, last-modified date, version history from SharePoint or Git) capturing the pre-review state as of the ASIC advisory date — the regulatory evidence artifact demonstrating your organization's posture at the time the advisory was issued and the timeline of your response

Per-Action IR Details

Step 1: Governance Review — Pull your current cyber resilience plan and mark every section that references threat actor capabilities or attack techniques. Flag any section written before 2024 that does not account for AI-augmented phishing, deepfake social engineering, or automated exploitation. Schedule a review cycle within 30 days.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR capability, policies, and ensuring plans reflect current threat landscape including AI-augmented adversary techniques

Controls: NIST IR-8 (Incident Response Plan) — plan must be updated to reflect AI-enabled BEC, deepfake social engineering, and automated credential acquisition as named threat vectors, NIST IR-2 (Incident Response Training) — training materials referencing pre-2024 threat actor TTPs must be revised to include AI-assisted phishing and voice deepfake scenarios, NIST RA-3 (Risk Assessment) — risk assessment must be re-scoped to account for AI-augmented attack surface across voice, email, and automated exploitation channels, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — review must confirm vulnerability management process accounts for AI-accelerated exploitation timelines, CIS 7.2 (Establish and Maintain a Remediation Process) — remediation process must include prioritization criteria for gaps exposed by AI-augmented threat capabilities

Compensating: Export your cyber resilience plan to PDF and use a keyword search (grep on Linux or Ctrl+F in any PDF viewer) to locate and flag sections containing terms like 'phishing,' 'social engineering,' 'authentication,' and 'wire transfer.' Create a simple spreadsheet tracking each flagged section, its last-revised date, and whether it references AI-augmented techniques. Use CISA's free Cyber Resilience Review (CRR) self-assessment tool (available at cisa.gov) to benchmark current maturity against financial sector expectations without requiring enterprise GRC software.

Evidence: Before modifying any documentation, capture version-controlled snapshots of your current cyber resilience plan, incident response playbooks, and board-approved risk assessments — specifically sections governing executive communication approval workflows, wire transfer authorization procedures, and phishing response. These pre-review artifacts establish your baseline posture for regulatory evidence purposes under ASIC's expectation of documented responsiveness. Note creation and last-modified timestamps on all documents using OS file metadata (ls -la on Linux, Get-Item in PowerShell) before any edits begin.

Step 2: Detection — Audit your email security gateway logs for signs of high-volume, low-variance phishing campaigns that suggest AI-generated lures. Review identity and access logs for anomalous authentication patterns consistent with credential harvesting (T1589). If you run voice-based authentication or executive wire-transfer approval processes, assess whether deepfake audio detection controls are in place.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Correlating indicators across email, identity, and voice channels to detect AI-augmented credential harvesting and social engineering consistent with T1589 (Gather Victim Identity Information)

Controls: NIST SI-4 (System Monitoring) — monitoring must extend to email gateway telemetry for AI-generated lure patterns and to identity provider logs for anomalous authentication velocity or geolocation anomalies consistent with credential harvesting, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — periodic review of email security gateway logs and IdP/SSO authentication logs specifically for low-variance, high-volume phishing characteristics indicative of AI-generated campaigns, NIST AU-2 (Event Logging) — confirm event logging is enabled for authentication failures, MFA bypass attempts, and email gateway disposition events relevant to AI-phishing detection, CIS 8.2 (Collect Audit Logs) — verify audit log collection is active across email gateway, IdP, and any voice authentication platform used in wire-transfer approval workflows, MITRE ATT&CK T1589 (Gather Victim Identity Information) — detection focus: unusual reconnaissance patterns against executive directories, LinkedIn scraping indicators in web proxy logs, and credential stuffing attempts against financial application login portals

Compensating: For email gateway analysis without a SIEM: export MTA or email security gateway logs (e.g., Microsoft Exchange message tracking logs, Proofpoint/Mimecast export, or Postfix logs) and run a Python one-liner or awk command to identify sender domains with high message volume but low recipient diversity — a statistical signature of AI-generated phishing at scale. For identity log review: pull Azure AD Sign-In logs or on-prem Active Directory Event ID 4625 (Failed Logon) and 4648 (Explicit Credential Use) via PowerShell (Get-WinEvent -LogName Security -FilterXPath) and pivot on accounts associated with finance, treasury, or executive functions. For deepfake audio controls assessment: if no commercial deepfake detection tool is deployed, implement a manual out-of-band call-back procedure using a pre-registered number from a known-good source — document this as a compensating control.

Evidence: Capture the following before any remediation action: (1) Email gateway logs showing sender IP reputation, DKIM/SPF/DMARC disposition, subject line patterns, and click/open telemetry for the past 90 days — specifically filter for campaigns where message body similarity is high across many recipients but sender domains vary, a hallmark of AI-generated lure rotation. (2) IdP or SSO authentication logs (Azure AD, Okta, or on-prem AD) for Event IDs 4625, 4648, 4768, 4771 on accounts in finance, treasury, and executive assistant roles, filtered for off-hours or geographically implausible login attempts consistent with T1589 credential harvesting outcomes. (3) Any voice platform call records (Teams, Cisco CUCM CDRs, or telephony provider logs) for inbound calls to executives or finance staff that preceded unusual wire-transfer requests or authorization approvals.

Step 3: Eradication — Identify and close process gaps that rely on voice or email confirmation alone for high-value transactions. Enforce phishing-resistant MFA (hardware tokens or passkeys) on all privileged and finance-adjacent accounts. Remove or restrict access paths that have no compensating control against automated credential acquisition (T1650).

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication: Removing the conditions that enabled AI-augmented BEC and deepfake fraud by eliminating single-channel transaction approval paths and enforcing phishing-resistant authentication on accounts targeted by automated credential acquisition (T1650)

Controls: NIST IA-5 (Authenticator Management) — replace SMS or voice OTP authenticators with FIDO2 hardware tokens or passkeys on all accounts in finance, treasury, executive, and privileged administrator roles to eliminate phishing-susceptible MFA factors exploitable by AI-generated lures, NIST AC-3 (Access Enforcement) — enforce access restrictions on financial transaction approval systems such that no single voice or email channel can authorize high-value transfers without a phishing-resistant second factor, NIST SI-2 (Flaw Remediation) — treat single-channel voice/email approval as a process flaw requiring documented remediation with a defined timeline, consistent with SI-2

flaw tracking requirements, CIS 6.3 (Require MFA for Externally-Exposed Applications) — enforce phishing-resistant MFA on all externally-exposed financial applications, portals, and VPN entry points immediately, CIS 6.5 (Require MFA for Administrative Access) — extend phishing-resistant MFA enforcement to all administrative accounts, with specific priority on accounts that can initiate, approve, or modify wire-transfer or payment workflows, MITRE ATT&CK T1650 (Acquire Access) — eradication target: access paths (shared credentials, legacy authentication protocols like NTLM or basic auth on financial APIs) that enable automated credential acquisition must be identified via Active Directory authentication logs and disabled or restricted

Compensating: To identify legacy authentication paths without an enterprise PAM tool: run the following PowerShell command against Azure AD sign-in logs to surface accounts using legacy protocols — Get-AzureADAuditSignInLogs | Where-Object {$_.clientAppUsed -notmatch 'Browser|Mobile Apps'} | Select UserPrincipalName, ClientAppUsed, CreatedDateTime. For on-prem AD, query Event ID 4776 (NTLM authentication) in the Security event log filtered to finance and executive accounts. To enforce process controls without technology spend: implement a documented dual-authorization callback procedure for wire transfers above a defined threshold, using a pre-registered direct-dial number verified out-of-band — this is a recognized compensating control in financial sector audit frameworks and costs nothing to implement.

Evidence: Before closing access paths, document the following as evidence of pre-remediation state: (1) A list of all accounts in finance, treasury, executive, and privileged roles and their current MFA method — exportable from Azure AD (Get-MsolUser), Okta Admin API, or on-prem AD — to establish what was phishing-susceptible prior to remediation. (2) Any transaction approval workflow documentation showing voice or email as the sole authorization channel, timestamped and version-controlled, to demonstrate the process gap that existed. (3) Authentication logs showing use of legacy protocols (NTLM, Basic Auth) on accounts adjacent to financial systems, exported from AD Security Event Log or Azure AD sign-in logs, as evidence of the automated credential acquisition surface being closed.

Step 4: Recovery — Validate that incident response playbooks include decision trees for suspected deepfake-facilitated fraud and AI-assisted phishing campaigns. Confirm contact trees for executive impersonation scenarios are current and tested. Run a tabletop exercise against an AI-augmented BEC scenario within 60 days.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery: Restoring operational confidence and process integrity after AI-augmented BEC or deepfake fraud by validating that playbooks, contact trees, and tabletop exercises are calibrated to AI-enabled adversary capabilities

Controls: NIST IR-4 (Incident Handling) — incident handling capability must explicitly include decision logic for AI-augmented BEC: criteria for declaring a suspected deepfake-facilitated fraud event, escalation triggers, and containment actions specific to voice and email impersonation scenarios, NIST IR-3 (Incident Response Testing) — tabletop exercise against an AI-augmented BEC scenario (executive impersonation via deepfake audio leading to fraudulent wire transfer) must be conducted and results documented to validate playbook adequacy, NIST IR-8 (Incident Response Plan) — contact trees for executive impersonation scenarios must be reviewed for currency, including verification that out-of-band contact numbers for executives and finance approvers are up to date and have been tested within the past cycle, CIS 7.2 (Establish and Maintain a Remediation Process) — tabletop findings must feed into the remediation process as documented gaps requiring prioritized closure, NIST CP-4 (Contingency Plan Testing) — the tabletop exercise serves double duty as contingency plan validation for financial transaction continuity in the event of a confirmed BEC or deepfake fraud incident

Compensating: For organizations without a dedicated IR simulation platform: design a low-cost tabletop using a written scenario inject document (freely templatable from CISA's Tabletop Exercise Packages at cisa.gov/tabletop-exercise-packages) with an AI-augmented BEC narrative — specifically, a scenario where the CFO receives a deepfake voicemail from the CEO directing an urgent wire transfer, followed by a convincing AI-phishing email with a lookalike domain. Walk finance, IT, and legal stakeholders through the decision tree verbally. Document inject responses, gaps identified, and action owners in a simple spreadsheet. No simulation software required — structured facilitation and documentation are the control.

Evidence: Before the tabletop, collect and preserve: (1) Current versions of all IR playbooks and contact trees, timestamped, to establish pre-exercise baseline — use file metadata (Get-Item in PowerShell or ls -la on Linux) to record last-modified dates. (2) Records of any prior BEC or wire-fraud incidents (even near-misses) from your ticketing

system or email, which should inform scenario design and serve as historical evidence of the threat's relevance to your organization. (3) Post-tabletop: capture all gap findings, inject responses, and identified decision points where playbooks lacked specific guidance on deepfake or AI-phishing scenarios — this documentation is both an IR improvement artifact and ASIC regulatory evidence.

Step 5: Post-Incident — Document control gaps identified during the review. Map gaps to NIST CSF 2.0 Govern and Protect functions. If your organization is ASIC-regulated, retain documentation of your review and any remediation actions as evidence of regulatory responsiveness. Align updated controls to CISA's financial sector guidance and NIST SP 800-53 controls SA-11, SI-3, and AC-2 as relevant.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Capturing lessons learned, documenting control gaps against AI-augmented threat scenarios, and aligning remediation to regulatory expectations under the ASIC advisory and CSF 2.0 Govern/Protect functions

Controls: NIST IR-5 (Incident Monitoring) — maintain documented tracking of all control gaps identified during the AI-resilience review, including gap owner, remediation timeline, and current status, to satisfy ongoing incident monitoring requirements, NIST IR-6 (Incident Reporting) — if any identified gap corresponds to a confirmed exploitation event (e.g., a completed deepfake-facilitated wire fraud), ensure reporting obligations to ASIC and relevant national authorities (ACSC in Australia) are met within required timeframes, NIST SI-3 (Malicious Code Protection) — review and update malicious code protection coverage to include AI-generated phishing payload delivery mechanisms, ensuring email security gateway signatures and behavioral detection rules are tuned to AI-lure characteristics identified during the detection audit in Step 2, NIST SA-11 (Developer Testing and Evaluation) — apply to any internally developed or vendor-integrated AI-assisted tools used in financial workflows: confirm security testing has been performed against adversarial AI input scenarios, including prompt injection and model manipulation relevant to financial decision support systems, NIST AC-2 (Account Management) — post-review account management documentation must reflect the enforcement of phishing-resistant MFA on finance and privileged accounts as remediated in Step 3, with evidence of account audit and access review retained for regulatory inspection, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — update vulnerability management process documentation to formally include AI-augmented threat scenarios as a reviewed category in each vulnerability management cycle, CIS 5.1 (Establish and Maintain an Inventory of Accounts) — retain updated account inventory showing MFA method per account as part of post-review documentation package for ASIC regulatory evidence

Compensating: For gap mapping without a GRC platform: create a simple spreadsheet with columns for control gap description, CSF 2.0 function (Govern/Protect/Detect/Respond/Recover), responsible owner, target remediation date, and evidence of completion. Export this as a dated PDF after each review cycle and store in a version-controlled location (SharePoint, Git repo, or even a dated folder on a secured file share) to constitute your ASIC regulatory evidence package. For SI-3 tuning without a commercial email security platform: if using open-source tools, deploy a SpamAssassin or Rspamd rule update to flag emails with characteristics of AI-generated lures (unusually high linguistic coherence combined with lookalike sender domains) and log dispositions to a local syslog target for manual review.

Evidence: Retain the following as the post-review evidence package: (1) The annotated cyber resilience plan from Step 1 showing pre-review gaps, with original file timestamps preserved as evidence of the state at the time of the ASIC advisory. (2) Email gateway and IdP log extracts from Step 2 showing the detection audit was conducted, including query parameters used, date range reviewed, and findings summary — even a null finding must be documented. (3) Account MFA audit output from Step 3 (exported account list with authenticator types, dated) showing pre- and post-remediation state. (4) Tabletop exercise documentation from Step 4 including scenario injects, participant responses, gap findings, and action items. (5) The completed CSF 2.0 gap mapping spreadsheet linking each identified control gap to a Govern or Protect function with remediation status — this constitutes the core ASIC regulatory responsiveness artifact.

Detection Guidance

There is no single IOC or log signature for this advisory, the threat is technique-based, not artifact-based. Focus detection effort on behavioral indicators across three areas. First, email and communication channels: flag sudden spikes in lookalike domain registrations targeting your brand (T1588), high-volume inbound phishing campaigns with near-identical but not identical lure text (indicative of AI-generated variance), and any executive impersonation attempts in email headers or display names. Second, identity and access: monitor for credential stuffing patterns (T1589), anomalous after-hours access by privileged accounts, and any access from geographic locations inconsistent with user baselines. Third, voice and video channels: if your organization uses voice verification for transaction authorization, log all such sessions and flag any that deviate from known voice baselines. Cross-reference your SIEM against MITRE ATT&CK techniques T1566.001 (Spearphishing Attachment), T1566.002 (Spearphishing Link), and T1589.002 (Email Addresses) for campaign indicators.

Framework Mappings

MITRE-ATTACK

- **T1589** — Gather Victim Identity Information
- **T1650** — Acquire Access
- **T1566** — Phishing
- **T1588** — Obtain Capabilities

NIST-800-53R5

- **AT-2** — Literacy Training and Awareness
- **CA-7** — Continuous Monitoring
- **SC-7** — Boundary Protection
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **SI-8** — Spam Protection

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.34** — Privacy and protection of personal information

CIS-V8

- **14.2** — Train Workforce Members to Recognize Social Engineering Attacks

HIPAA-SECURITY

- **164.308(a)(5)(i)** — Security Awareness and Training

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1589	Gather Victim Identity Information	Reconnaissance
T1650	Acquire Access	Resource-Development

Technique ID	Technique Name	Tactic
T1566	Phishing	Initial-Access
T1588	Obtain Capabilities	Resource-Development

Sources

Source	URL	Tier
gemini	https://fintech.global/2026/05/08/ai-threats-prompt-asic-cyber-resi...	T3
[PDF] Cyber Vulnerabilities at Large US Financial Institutions and Their ...	https://www.federalreserve.gov/econres/feds/files/2025103pap.pdf	T1
The 6 Biggest Cyber Threats for Financial Services in 2026 - UpGuard	https://www.upguard.com/blog/biggest-cyber-threats-for-financial-se...	T3
Financial Services Sector | Cybersecurity and Infrastructure ... - CISA	https://www.cisa.gov/topics/critical-infrastructure-security-and-re...	T1
Financial Cybersecurity Best Practices - HITRUST Alliance	https://hitrustalliance.net/blog/financial-cybersecurity-best-pract...	T3

Generated 2026-05-08 19:05 UTC by TJS Security Command Center