

Boardroom Comprehensive Abstracts
Study, Mitigation Strategy, and Dialog Session Summaries

AI Cyber Security Research Boardroom

2026-05-31

Contents

| | | |
|----------|---|----------|
| 1 | Boardroom Comprehensive Abstracts | 3 |
| 1.1 | Study Mitigation Dialog Session Summaries | 3 |
| 1.2 | Table of Contents | 4 |
| 2 | 1. Study Abstract (AICSR-STUDY-2026-001) | 5 |
| 2.1 | Purpose | 5 |
| 2.2 | Topics Covered | 5 |
| 2.3 | Principal Conclusions | 6 |
| 3 | 2. Mitigation Strategy Abstract (AICSR-MIT-2026-001) | 7 |
| 3.1 | Purpose | 7 |
| 3.2 | Topics Covered | 7 |
| 3.3 | Principal Conclusions | 7 |
| 4 | 3. Dialog Session Abstracts | 9 |
| 4.1 | 3.1 Introductions (Simulated 2026-05-01) | 9 |
| 4.1.1 | Topics Covered | 9 |
| 4.1.2 | Conclusions | 9 |
| 4.2 | 3.2 Rounds 120 (Simulated 2026-05-02) | 9 |
| 4.2.1 | Topics Covered | 10 |
| 4.2.2 | Conclusions | 10 |
| 4.3 | 3.3 Pillar 2+4 Red/Blue Exercise (Simulated 2026-05-04) | 10 |
| 4.3.1 | Topics Covered | 10 |
| 4.3.2 | Conclusions | 10 |
| 4.4 | 3.4 Remediation Pass 1 (Simulated 2026-05-05) | 11 |
| 4.4.1 | Topics Covered | 11 |
| 4.4.2 | Conclusions | 11 |
| 4.5 | 3.5 L2 Maturity Review (Simulated 2026-05-08) | 11 |
| 4.5.1 | Topics Covered | 11 |
| 4.5.2 | Conclusions | 11 |
| 4.6 | 3.6 Remediation Pass 2 (Simulated 2026-05-09) | 11 |
| 4.6.1 | Topics Covered | 11 |
| 4.6.2 | Conclusions | 12 |
| 4.7 | 3.7 MVAP v1.1 Adoption (Simulated 2026-05-12) | 12 |
| 4.7.1 | Topics Covered | 12 |
| 4.7.2 | Conclusions | 12 |
| 4.8 | 3.8 P7-05 Tabletop (Simulated 2026-05-15) | 12 |
| 4.8.1 | Topics Covered | 12 |

| | | |
|----------|--|-----------|
| 4.8.2 | Conclusions | 13 |
| 4.9 | 3.9 Government Risk Review (Sim. Q1) (Simulated 2026-05-16) | 13 |
| 4.9.1 | Topics Covered | 13 |
| 4.9.2 | Conclusions | 13 |
| 4.10 | 3.10 MVAP v1.2 Adoption (Simulated 2026-05-20) | 13 |
| 4.10.1 | Topics Covered | 13 |
| 4.10.2 | Conclusions | 13 |
| 4.11 | 3.11 Government Risk Review (Sim. Q2) (Simulated 2026-05-21) | 14 |
| 4.11.1 | Topics Covered | 14 |
| 4.11.2 | Conclusions | 14 |
| 5 | 4. Verification Ledger Abstract | 15 |
| 5.1 | Topics Covered | 15 |
| 5.2 | Conclusions | 15 |
| 6 | Footnotes and Reference Bibliography | 16 |
| 6.1 | Markdown Footnote Anchors | 18 |

Chapter 1

Boardroom Comprehensive Abstracts

1.1 Study Mitigation Dialog Session Summaries

| | |
|--------------------------|---|
| Document ID | AICSR-ABS-2026-001 |
| Version | 1.0 |
| Publication Date | 2026-05-31 |
| Parent Documents | AICSR-STUDY-2026-001, AICSR-MIT-2026-001, AICSR-DLG-2026-001 |
| Simulation Window | May 2026 |

*This document synthesizes a **thorough multi-agent boardroom simulation** conducted in **May 2026**. No physical meeting occurred. Participant voices, votes, exercises, and outcomes were produced by configured agent profiles under moderated debate with court-reporter verification. **All source materials remain available for inspection** in the repository: session transcripts (`sessions/`), MVAP specifications (`mvap/`), participant profiles (`participants/`), the verification ledger, and continuation manifests (`output/*-MANIFEST.yaml`, `output/CONTINUE*.md`). External citations reference real published sources (footnotes); speculative claims are labeled in `sessions/verification-ledger.md`.*

1.2 Table of Contents

1. Study Abstract
2. Mitigation Strategy Abstract
3. Dialog Session Abstracts
4. Verification Ledger Abstract
5. Footnotes and Reference Bibliography

Chapter 2

1. Study Abstract (AICSR-STUDY-2026-001)

Source: output/Cyber-Security-AI-Diligence-Research-Study.md Version 1.2 Publication 2026-05-31

2.1 Purpose

The comprehensive study synthesizes eleven boardroom session transcripts, thirty-one participant profiles, MVAP specifications v1.1v1.2, and a living government/zero-day risk register into a single reference document for practitioners, compliance officers, and researchers examining **Cyber-Security and AI Diligence**.

2.2 Topics Covered

- **Methodology:** Multi-agent simulation framework, governance roles, debate protocol, May 2026 compressed chronology, source inspectability (Section 2.42.5)
- **MVAP framework:** Seven pillars, maturity levels L1L3, CI/CD pipeline gates, v1.2 adopted controls (P7-09P7-11, P6 tier-2)
- **Pillar analysis:** Governance registry, OWASP LLM Top 10, SBOM/cosign, IR-AI-02 SOAR, human deepfake training, firmware, zero-day evaluation
- **Zero-day risk:** Historical OSS timeline, OS/application targets, LiteLLM KEV chain case study, AI-accelerated offensive chain, source evaluation matrix
- **Government risk:** CISA governance collapse, GAO 815 contractor violations, Salt Typhoon persistence, EO 14409 implementation gap, GOV-01GOV-15 register
- **Deliberation outcomes:** Twenty-round arc, pillar votes, material dissent, post-round session votes
- **Exercises:** Red/Blue 72-hour, remediation passes, P7-05 tabletop, simulated government reviews
- **Participant profiles:** All thirty-one agents with boardroom citation counts
- **Appendices:** Verification ledger summary, expanded path index (27 artifacts), alphabetical index, full bibliography

2.3 Principal Conclusions

1. **MVAP is necessary but not sufficient** for regulated, OT, or nation-state threat profiles (22/27 adoption).
 2. **Enterprises cannot rely on federal cybersecurity backstop** — CISA capacity loss, MS-ISAC defunding, and NVD backlog require P7-08 redundant intel and self-sufficient KEV operations.
 3. **AI gateways are the new edge appliances** — LiteLLM-class KEV chains demand P7-11 hardening and P7-10 4-hour SLA.
 4. **Application controls lagged detection in Red/Blue exercise** — P2-03/P2-04 failures drove remediation; production MTTC reached 14 minutes post-sprint.
 5. **Classified-adjacent risk is structural** — GOV-02 and GOV-13 remain Critical+ despite v1.2 mitigations for gateway exploitation.
 6. **Simulation transparency** — all claims traceable to `sessions/`, `mvap/`, `participants/`, and footnoted external sources.
-

Chapter 3

2. Mitigation Strategy Abstract (AICSR-MIT-2026-001)

Source: output/MVAP-Complete-Mitigation-Strategy.md Version 1.1 Derived from AICSR-STUDY-2026-001

3.1 Purpose

Translates boardroom findings into an actionable **enterprise mitigation program** for production LLM and AI inference deployments. Maps every major threat from the study to specific MVAP control IDs with evidence artifacts, owners, and phased implementation timelines aligned to the simulated May 2026 session sequence.

3.2 Topics Covered

- **Threat-to-control mapping:** CVE weaponization chain, LiteLLM RCE, shadow AI, poisoned RAG, LotL movement, classified spill, federal intel gap, transitive dependency chains
- **Pillar mitigations (P1P7):** Registry, OWASP LLM tests, SBOM/cosign, IR-AI-02 playbook, phishing/deepfake training, GPU attestation, KEV sweep, gateway hardening checklist
- **Government mitigations:** GOV-02 contractor audit, GOV-13 RAG spill prevention, P7-07 classified-adjacent checklist, P7-08 intel redundancy table
- **Implementation roadmap:** Phase 0 emergency (4 May), Phase 1 foundation (512 May), Phase 2 L2 operational (816 May), Phase 3 hardening (20 May+)
- **Metrics and ownership:** KPIs (KEV SLA $\geq 95\%$, MTTC $< 15m$, phishing $< 5\%$), pillar owner matrix, verification protocol

3.3 Principal Conclusions

1. **Priority P0 mitigations:** P7-11 gateway hardening + P7-10 SLA + P4-05 IR-AI-02 for AI gateway RCE; P7-07 + P1-03 for classified spill.
2. **Self-sufficient intel is non-negotiable** — CISA KEV + OSV + GitHub Advisory trinity; do not wait for NVD or depleted federal advisories.

3. **Pre-deploy compromise beats post-deploy hacking** — SBOM and cosign verification before model reaches inference cluster (Oliver Hansen thesis).
 4. **L2 maintenance requires semi-annual purple-team** and enforced CI/CD gates — aspirational documentation without tested playbooks fails Red/Blue standard.
 5. **Residual risk acceptance required** for GOV-02/GOV-13 until contractor program and spill controls mature — track quarterly with attestation evidence.
-

Chapter 4

3. Dialog Session Abstracts

Full transcripts: `output/Boardroom-Complete-Dialog-Transcript.md` (AICSR-DLG-2026-001). Below: topics and conclusions per session in simulated chronological order.

4.1 3.1 Introductions (Simulated 2026-05-01)

Source: `sessions/2026-06-22-cybersecurity-ai-diligence-introductions.md` Dialog section 2.1

4.1.1 Topics Covered

- Boardroom governance: Arthur Vance (moderator, non-voting) and Eleanor Vance (court reporter, verification)
- Twenty-seven voting participants introduced across nine expertise groups
- Debate protocol: one positive and one negative point per voter per round; 14/27 majority threshold
- Governing research question framed for all future sessions
- Participant perspectives: compliance, architecture, SOC, zero-day, red/blue rapid response, human factors

4.1.2 Conclusions

- Boardroom charter ratified; all thirty-one profiles seated and speaking order established
 - Core research question locked: minimum viable AI diligence for enterprise LLM deployment
 - Verification ledger protocol adopted — external sources required; speculation must be labeled
 - Session adjourned with unanimous readiness to begin twenty-round deliberation
-

4.2 3.2 Rounds 120 (Simulated 2026-05-02)

Source: `sessions/2026-06-22-cybersecurity-ai-diligence-rounds-1-20.md` Dialog section 2.2

4.2.1 Topics Covered

- Rounds 14: MVAP definition, NIST AI RMF mapping, OWASP LLM Top 10, compliance officer reconciliation
- Rounds 59: SBOM/supply chain, RAG sanitization, model lineage, red-team mandate, SOC telemetry on AI APIs
- Rounds 1014: Cryptography/embeddings, human deepfake risk, cloud shared responsibility, mobile scope, OT AI tier
- Rounds 1518: Firmware/GPU scope, kill-chain mapping, CCO budget, minimum viable vs aspirational controls
- Rounds 1920: Formal pillar votes and final MVAP adoption with dissent record

4.2.2 Conclusions

- **MVAP adopted 22/27** — five pillars: Governance (25/27), App/LLM (24/27), Supply Chain (23/27), Detection (24/27), Human (22/27)
- NIST AI RMF and OWASP LLM Top 10 confirmed as mandatory baselines
- Material dissent recorded: firmware mandatory all tiers (9/27), SLSA L3 + 24h KEV (8/27), OT funding (11/27), academic robustness metrics (6/27)
- Next steps mandated: publish MVAP spec, run Pillar 2+4 Red/Blue exercise, 90-day maturity review

4.3 3.3 Pillar 2+4 Red/Blue Exercise (Simulated 2026-05-04)

Source: sessions/2026-06-22-mvap-pillar-2-4-red-blue-exercise.md Dialog section 2.3

4.3.1 Topics Covered

- 72-hour adversarial exercise against staging RAG chatbot on EKS
- Red Rapid: shadow recon, prompt injection, poisoned RAG PDF, overprivileged export_tool, kerberoast, 220GB exfil
- Blue Rapid: token baseline alert T+6m, LotL detection T+14h, IR-AI-01 containment T+11h, immutable backup restore
- Pillar 2 tests: guardrails (94% block), RAG sanitization, tool authorization, SAST/DAST
- Pillar 4 tests: API logging, token baselines, ATT&CK mapping, SOAR playbook, backup integrity

4.3.2 Conclusions

- **Pillar 2 L1: NOT READY** — P2-03 failed (7/10 poisoned RAG); P2-04 tool scope overprivileged
- **Pillar 4 L1: CONDITIONAL** — detection fired but correlation/MTTC gaps; unified SOAR needed
- 16/27 agree **30-day remediation sprint** before L1 certification
- Remediation priorities: P2-04 ACL scoping, P2-03 blocklist, P4-05 unified AI abuse playbook

4.4 3.4 Remediation Pass 1 (Simulated 2026-05-05)

Source: `sessions/2026-07-22-mvap-p2-03-p4-05-remediation-validation-1.md` Dialog section 2.4

4.4.1 Topics Covered

- Staging re-validation of P2-03 poisoned RAG corpus after metadata scrubber deployment
- P4-05 SOAR playbook IR-AI-02 chaining P4-02 token alerts to P4-03 LotL detection
- Per-user tool ACL implementation for `export_tool` (P2-04)
- Multilingual injection bypass retest for P2-02

4.4.2 Conclusions

- P2-03 **9/10 conditional PASS** (16/27) — one metadata bypass remains
 - P4-05 **PASS** (19/27) — staging MTTC 22 minutes
 - P2-04 per-user ACL reconfirmed
 - Board orders Pass 2 production validation before L2 operational declaration
-

4.5 3.5 L2 Maturity Review (Simulated 2026-05-08)

Source: `sessions/2026-09-20-mvap-level-2-maturity-review.md` Dialog section 2.5

4.5.1 Topics Covered

- L1 checklist review per pillar: registry, SBOM, SOC telemetry, human simulations
- L2 requirements: purple-team validation + automated CI/CD pipeline gates
- Dissent reconciliation from MVAP v1.0 adoption
- MVAP v1.1 scope preview: firmware elevation, Pillar 7 zero-day, OT tier funding

4.5.2 Conclusions

- **L2 promoted 18/27** — conditional on Pass 2 production validation
 - L1 certified organization-wide; P2 conditional pending production retest
 - Pipeline gates defined: registry, SAST, SBOM cosign, KEV sweep, human sign-off on AI patches
 - Scheduled Pass 2 for P2-03 and P4-05 next session
-

4.6 3.6 Remediation Pass 2 (Simulated 2026-05-09)

Source: `sessions/2026-09-21-mvap-p2-03-p4-05-remediation-validation-2.md` Dialog section 2.6

4.6.1 Topics Covered

- Production-environment retest of P2-03 poisoned RAG with QR/DOCX scrubber

- Production SOAR IR-AI-02 MTTC measurement across three incident simulations
- P2-04 ACL persistence verification under load
- L2 operational declaration criteria

4.6.2 Conclusions

- P2-03 **11/12 production certified** (91.7%) — exceeds $\geq 9/10$ threshold
 - P4-05 **14-minute average MTTC** in production — meets L2 target ($< 15m$)
 - **Remediation sprint closed** — L2 operational v1.0
 - Verification ledger updated; backlog items queued for v1.1
-

4.7 3.7 MVAP v1.1 Adoption (Simulated 2026-05-12)

Source: `sessions/2026-12-20-mvap-v1-1-backlog-zero-day-government-risk.md` **Dialog section 2.7**

4.7.1 Topics Covered

- v1.0 backlog closure and remediation certification acknowledgment
- Pillar 7 proposal: open-source zero-day source evaluation (P7-01P7-08)
- Pillar 6 elevation: firmware/GPU mandatory tier-1
- Government classified infrastructure risk briefing: CISA vacuum, GAO 815 violations, Salt Typhoon, FOUO spill
- Living risk register GOV-01GOV-08 initial ratification

4.7.2 Conclusions

- **MVAP v1.1 adopted 20/27**
 - **Pillar 7 adopted 20/27** — KEV sweep, native SAST, source audit, intel redundancy
 - **P6 tier-1 mandatory 17/27**
 - **Government risk register ratified 22/27** at `mvap/ZERO-DAY-OPEN-SOURCE-RISK-ASSESSMENT.md`
-

4.8 3.8 P7-05 Tabletop (Simulated 2026-05-15)

Source: `sessions/2027-03-20-p7-05-zero-day-tabletop.md` **Dialog section 2.8**

4.8.1 Topics Covered

- LiteLLM KEV chain replay: CVE-2026-42208 SQLi, CVE-2026-42271 MCP injection, CVE-2026-48710 BadHost chain
- Classified spill-to-RAG subplot tied to GAO contractor violation patterns
- P7-11 gateway hardening tabletop walkthrough
- P7-10 KEV SLA timer exercise
- Cross-pillar coordination: P4-05 IR-AI-02 under zero-day pressure

4.8.2 Conclusions

- **Tabletop PASS 19/27**
 - Gaps identified: transitive SBOM (P7-09), 4h KEV SLA enforcement (P7-10), gateway config audit
 - Classified spill subplot exposed P7-07 drill inadequacy for tier-1 contractors
 - Drives MVAP v1.2 specification draft
-

4.9 3.9 Government Risk Review (Sim. Q1) (Simulated 2026-05-16)

Source: sessions/2027-03-20-quarterly-government-risk-review.md Dialog section 2.9

4.9.1 Topics Covered

- EO 14409 AI cybersecurity clearinghouse vs CISA capacity collapse (GOV-09)
- LiteLLM/Starlette KEV entries — AI gateway as new edge appliance (GOV-11, GOV-12)
- NVD backlog vs AI discovery speed (GOV-10) — validates P7-08 redundant feeds
- Salt Typhoon persistence update; GAO contractor program unchanged (GOV-02 Critical+)
- v1.2 preview votes: P7-09, P7-10, P7-11

4.9.2 Conclusions

- **Risk register v2.0 ratified 21/27** — GOV-09 through GOV-15 added
 - P7-10 4h KEV SLA passed 19/27; P7-11 gateway hardening 21/27 preview
 - Enterprises cannot rely on federal backstop — P7-08 redundancy mandatory
 - P6 tier-1+tier-2 firmware passed 18/27 in preview
-

4.10 3.10 MVAP v1.2 Adoption (Simulated 2026-05-20)

Source: sessions/2027-06-20-mvap-v1-2-adoption.md Dialog section 2.10

4.10.1 Topics Covered

- Final votes on P7-09 transitive SBOM, P7-10 KEV SLA, P7-11 gateway hardening
- P6 tier-1 + tier-2 firmware mandate
- P2-08 membership inference deferral to v1.2.1
- SLSA L3 tier-1 deferral to v1.3

4.10.2 Conclusions

- **MVAP v1.2 adopted 20/27**
- P7-10 and P7-11 **adopted**; P6 tier-2 **adopted**; P7-09 **conditional** (tier-1 mandatory follow-on)
- NullByte/Kira dissent on SLSA L3 recorded for v1.3
- Quarterly government re-attestation now includes P7-10 SLA and P7-11 audit evidence

4.11 3.11 Government Risk Review (Sim. Q2) (Simulated 2026-05-21)

Source: sessions/2027-06-20-quarterly-government-risk-review-q2.md Dialog section 2.11

4.11.1 Topics Covered

- GOV-01 through GOV-15 status reassessment post-v1.2 adoption
- EO 14409 clearinghouse partial operational status
- P7-11 compliance audit results (18/27 orgs PASS)
- GOV-11/GOV-12 mitigation via gateway hardening and transitive SBOM controls
- Residual Critical+ risks: GOV-02 contractor program, GOV-13 classified spill to RAG

4.11.2 Conclusions

- **Risk register reaffirmed 20/27**
 - GOV-11 and GOV-12 downgraded to MITIGATED for P7-11-compliant organizations
 - GOV-02 and GOV-13 remain highest residual enterprise risks
 - P7-08 redundant feeds caught advisories ahead of NVD this cycle
-

Chapter 5

4. Verification Ledger Abstract

Source: `sessions/verification-ledger.md`

5.1 Topics Covered

- Status taxonomy: Verified, Partial, Unverified, Projected Speculation
- Claim-by-claim audit trail from introductions through v1.2 adoption
- External source URLs for government risk, KEV entries, MVAP vote tallies
- Stripped hallucinations: zero prompt-injection risk, mandatory SLSA L3 Q3 2026, 99% deepfake accuracy

5.2 Conclusions

- Eleanor Vance court reporter protocol provides the authoritative boundary between simulation narrative and verifiable fact
 - All three reports (Study, Mitigation, Abstracts) should cite ledger status when asserting compliance or threat claims
 - Ledger is a living document — regenerate reports after ledger updates
-

Chapter 6

Footnotes and Reference Bibliography

Complete numbered bibliography of sources cited in this document. External URLs were verified during simulation; repository paths are inspectable locally.

1. **NIST AI Risk Management Framework 1.0** (`nist-airmf`)
 - <https://www.nist.gov/itl/ai-risk-management-framework>
2. **NIST Generative AI Profile (NIST.AI.600-1)** (`nist-genai`)
 - <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
3. **OWASP Top 10 for LLM Applications 2025** (`owasp-llm`)
 - <https://genai.owasp.org/llm-top-10/>
4. **CISA Known Exploited Vulnerabilities Catalog** (`cisa-kev`)
 - <https://www.cisa.gov/known-exploited-vulnerabilities-catalog>
5. **SLSA Supply-chain Levels for Software Artifacts v1.0** (`slsa`)
 - <https://slsa.dev/spec/v1.0/>
6. **MITRE ATLAS — Adversarial ML** (`mitre-atlas`)
 - <https://atlas.mitre.org/>
7. **GAO-26-107861 — 815 Classified Contractor Security Violations** (`gao-classified`)
 - <https://www.gao.gov/products/gao-26-107861>
8. **CSA Research Note — CISA Leadership Governance Vacuum (2026-04-24)** (`csa-cisa`)
 - <https://labs.cloudsecurityalliance.org/research/csa-research-note-cisa-leadership-governance-vacuum-20260424/>
9. **GCA — Salt Typhoon Across the Internet** (`gca-salt`)
 - <https://globalcyberalliance.org/new-report-salt-typhoon-across-the-internet/>
10. **CISA Advisory AA25-239A — Salt Typhoon** (`cisa-salt`)
 - <https://www.cisa.gov/news-events/cybersecurity-advisories/aa25-239a>
11. **Trend Micro — U.S. Public Sector Under Siege Q1 2026** (`trend-q1`)
 - https://www.trendmicro.com/en_us/research/26/d/us-public-sector-under-siege.html
12. **TechCrunch — Acting CISA chief uploaded FOUO docs to ChatGPT** (`techcrunch-cisa`)
 - <https://techcrunch.com/2026/01/28/trumps-acting-cybersecurity-chief-uploaded-sensitive-government-docs-to-chatgpt/>

13. **GAO-26-109159 — Water Sector Cybersecurity** (gao-water)
 - <https://www.gao.gov/products/gao-26-109159>
14. **White House EO 14409 — AI Innovation and Security (2026-06-02)** (eo-14409)
 - <https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/>
15. **CSA Whitepaper — NVD Infrastructure Crisis & AI Vulnerability Discovery** (csa-nvd)
 - https://labs.cloudsecurityalliance.org/wp-content/uploads/2026/05/CSA_whitepaper_NVD_infrastructure_csa-styled.pdf
16. **CISA Alert — CVE-2026-42271 LiteLLM added to KEV** (litellm-kev)
 - <https://www.cisa.gov/news-events/alerts/2026/06/08/cisa-adds-two-known-exploited-vulnerabilities-catalog>
17. **Horizon3.ai — LiteLLM chained with Starlette BadHost RCE** (horizon3-chain)
 - <https://horizon3.ai/attack-research/vulnerabilities/cve-2026-42271-chained-with-cve-2026-48710/>
18. **The Hacker News — LiteLLM CVE-2026-42208 exploited within 36h** (litellm-sqli)
 - <https://thehackernews.com/2026/04/litellm-cve-2026-42208-sql-injection.html>
19. **OSTIF — BadHost vulnerability in Starlette** (ostif-badhost)
 - <https://ostif.org/disclosing-the-badhost-vulnerability-in-starlette/>
20. **NJCCIC — Salt Typhoon targets House Committee emails** (njccic-house)
 - <https://www.cyber.nj.gov/Home/Components/News/News/1935/214>
21. **ISA/IEC 62443 Industrial Cybersecurity Standards** (iec-62443)
 - <https://www.isa.org/standards-and-publications/isa-standards/isa-iec-62443-series-of-standards>
22. **CyberScoop — FBI confirms Salt Typhoon still ongoing (Feb 2026)** (fbi-salt)
 - <https://cyberscoop.com/fbi-salt-typhoon-ongoing-threat-cybertalks-2026/>
23. **StateScoop — CISA ending MS-ISAC support** (ms-isac)
 - <https://statescoop.com/cisa-confirms-its-ending-ms-isac-support/>
24. **Bloomberg Government — 815 classified data violations summary** (bgov-gao)
 - <https://news.bgov.com/bloomberg-government-news/us-companies-had-815-classified-data-violations-gao-finds>
25. **AI Cyber Security Research Study AICSR-STUDY-2026-001** (study-ref)
 - output/Cyber-Security-AI-Diligence-Research-Study.md
26. **MVAP Complete Mitigation Strategy AICSR-MIT-2026-001** (mit-ref)

- output/MVAP-Complete-Mitigation-Strategy.md
27. **Boardroom Complete Dialog Transcript AICSR-DLG-2026-001** (dlg-ref)
 - output/Boardroom-Complete-Dialog-Transcript.md
 28. **Boardroom Comprehensive Abstracts AICSR-ABS-2026-001** (abs-ref)
 - output/Boardroom-Comprehensive-Abstracts.md
 29. **How the Research Was Done AICSR-METHOD-2026-001** (method-ref)
 - output/How-The-Research-Was-Done.md
 30. **Research Materials Index AICSR-MATINDEX-2026-001** (matindex-ref)
 - output/RESEARCH-MATERIALS-INDEX.md

6.1 Markdown Footnote Anchors