

MVAP Complete Mitigation Strategy

Enterprise Implementation Guide Derived from the AI Cyber Security Research Boardroom Study

AI Cyber Security Research Boardroom

2026-05-31

Contents

1	MVAP Complete Mitigation Strategy AICSR-MIT-2026-001	2
1.1	Boardroom Citation Context	2
1.2	Source Location	3
1.3	Captured Content	3
2	MVAP Complete Mitigation Strategy	3
2.1	Enterprise Implementation Guide	3
2.2	Table of Contents	3
3	1. Executive Summary	4
3.1	Strategic Thesis	4
3.2	Mitigation Priorities (Ordered)	4
3.3	Success Criteria	5
4	2. Threat-to-Control Mapping	5
5	3. Pillar 1 — Governance Mitigations	6
5.1	3.1 Control Implementation	6
5.2	3.2 Governance Mitigation Playbook	6
6	4. Pillar 2 — Application & LLM Mitigations	6
6.1	4.1 Critical Mitigations (Post-Remediation)	7
6.2	4.2 Application Mitigation Sequence	7
7	5. Pillar 3 — Supply Chain Mitigations	7
8	6. Pillar 4 — Detection & Response Mitigations	8
8.1	6.1 IR-AI-02 Unified Playbook	8
8.2	6.2 Detection Engineering	8
9	7. Pillar 5 — Human Layer Mitigations	8
10	8. Pillar 6 — Firmware Mitigations	9

11 9. Pillar 7 — Zero-Day & Open-Source Mitigations	9
11.1 9.1 Daily Operations (P7-01, P7-08)	9
11.2 9.2 v1.2 Adopted Mitigations	10
11.2.1 P7-10 KEV Patch SLA	10
11.2.2 P7-11 Gateway Hardening Checklist	10
11.2.3 P7-09 Transitive SBOM (Conditional)	10
11.3 9.3 Source Audit Program	10
12 10. Government & Classified-Adjacent Mitigations	10
12.1 10.1 Residual High Risks	11
12.2 10.2 Classified-Adjacent Checklist (P7-07)	11
12.3 10.3 Intel Redundancy (P7-08)	11
13 11. Implementation Roadmap	12
13.1 Phase 0 — Emergency (Sim. 2026-05-04)	12
13.2 Phase 1 — Foundation (Sim. 2026-05-05 through 2026-05-12)	12
13.3 Phase 2 — Operational L2 (Sim. 2026-05-08 through 2026-05-16)	12
13.4 Phase 3 — Hardening (Sim. 2026-05-20 through follow-on)	13
14 12. Metrics, Verification & Ownership	13
14.1 12.1 Key Performance Indicators	13
14.2 12.2 Pillar Ownership Matrix	14
14.3 12.3 Verification Protocol	14
15 Index	14
16 Footnotes and Reference Bibliography	15
16.1 Markdown Footnote Anchors	17

1 MVAP Complete Mitigation Strategy AICSR-MIT-2026-001

Field	Value
Reference key	mit-ref
Availability	AVAILABLE
Capture method	primary
Source type	local_repository
URL / path	output/MVAP-Complete-Mitigation-Strategy.md
Captured (UTC)	2026-06-23T03:31:40Z
Content type	text/markdown
HTTP status	n/a

1.1 Boardroom Citation Context

Mitigation strategy derived from study.

1.2 Source Location

output/MVAP-Complete-Mitigation-Strategy.md

1.3 Captured Content

2 MVAP Complete Mitigation Strategy

2.1 Enterprise Implementation Guide

Derived from AICSR-STUDY-2026-001 — Cyber-Security and AI Diligence Research Study

Document ID	AICSR-MIT-2026-001
Version	1.1
Publication Date	2026-05-31
Simulation Window	May 2026 (compressed chronology)
Parent Study	AICSR-STUDY-2026-001 v1.2
Framework	MVAP v1.2 (adopted simulated 2026-05-20)
Maturity Target	L2 operational L3 aspirational

*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`.*

All threat assessments reference AICSR-STUDY-2026-001 Section 1-8 (output/Cyber-Security-AI-Diligence-Research-Study.md). Dialog evidence: output/Boardroom-Complete-Dialog-Transcript.md.**

ewpage

2.2 Table of Contents

1. Executive Summary
2. Threat-to-Control Mapping
3. Pillar 1 — Governance Mitigations
4. Pillar 2 — Application & LLM Mitigations
5. Pillar 3 — Supply Chain Mitigations
6. Pillar 4 — Detection & Response Mitigations
7. Pillar 5 — Human Layer Mitigations
8. Pillar 6 — Firmware Mitigations

- 9. Pillar 7 — Zero-Day & Open-Source Mitigations
- 10. Government & Classified-Adjacent Mitigations
- 11. Implementation Roadmap
- 12. Metrics, Verification & Ownership
- 13. Index

ewpage

3 1. Executive Summary

This document provides a **complete mitigation strategy** for enterprises deploying production LLM and AI inference systems. It is derived directly from *AICSR-STUDY-2026-001* Section 1 (`output/Cyber-Security-AI-Diligence-Research-Study.md`) and implements controls adopted through MVAP v1.2 per *AICSR-STUDY-2026-001* Section 3.4 (`output/Cyber-Security-AI-Diligence-Research-Study.md`).

3.1 Strategic Thesis

Per *AICSR-STUDY-2026-001* Section 6.1 (`output/Cyber-Security-AI-Diligence-Research-Study.md`), enterprises **cannot rely on federal cybersecurity backstop** for AI/zero-day intelligence. Mitigation must be **self-sufficient**, **pipeline-embedded**, and **verifiable** through the Eleanor Vance verification ledger (*AICSR-STUDY-2026-001* Section Appendix A (`output/Cyber-Security-AI-Diligence-Research-Study.md`)).

3.2 Mitigation Priorities (Ordered)

Priority	Threat (Study ref)	Primary controls	Owner tier
P0	AI gateway RCE (Ch. 5.4)	P7-11, P7-10, P4-05	Blue Rapid + Kira/Oliver
P0	Classified spill to RAG (GOV-13)	P7-07, P1-03	Marcus Thorne, Tariq
P1	KEV weaponization <36h (Ch. 5.5)	P7-01, P7-10, P7-08	Sarah Jenkins, Kira
P1	Shadow AI bypass (Ch. 8.1)	P1-05, P2-04, P4-02	Victor Vance, Maya
P2	Poisoned RAG (Ch. 8.2)	P2-03, P2-06	Maya Patel, Synapse
P2	Contractor program gaps (GOV-02)	P7-07, P3-04	Marcus Thorne
P3	Human deepfake/phishing (Ch. 4.5)	P5-01P5-04	Susan Albright, Mateo
P3	Firmware/GPU KEV (Ch. 5.2)	P6-01P6-04	Aether, Hex

3.3 Success Criteria

- **L2 maintained:** Purple-team semi-annual PASS; CI/CD gates enforced (*AICSR-STUDY-2026-001* Section 3.1 (`output/Cyber-Security-AI-Diligence-Research-Study.md`))
- **P7-10 SLA:** $\geq 95\%$ tier-1 KEV patches within 4 hours (*AICSR-STUDY-2026-001* Section 3.4 (`output/Cyber-Security-AI-Diligence-Research-Study.md`))
- **IR-AI-02 MTTC:** < 15 minutes production average (*AICSR-STUDY-2026-001* Section 4.4 (`output/Cyber-Security-AI-Diligence-Research-Study.md`), *AICSR-STUDY-2026-001* Section 8.2 (`output/Cyber-Security-AI-Diligence-Research-Study.md`))
- **GOV residual risk:** GOV-02 and GOV-13 tracked in simulated reviews (*AICSR-STUDY-2026-001* Section 6.8 (`output/Cyber-Security-AI-Diligence-Research-Study.md`), *AICSR-STUDY-2026-001* Section 6.9 (`output/Cyber-Security-AI-Diligence-Research-Study.md`))

4 2. Threat-to-Control Mapping

Cross-reference: *AICSR-STUDY-2026-001* Section 5.5 (`output/Cyber-Security-AI-Diligence-Research-Study.md`) offensive chain and *AICSR-STUDY-2026-001* Section 6.7 (`output/Cyber-Security-AI-Diligence-Research-Study.md`) government risk register.

Attack Step	Study Evidence	MVAP Control	Mitigation Action
CVE feed scrape	Ch. 5.5 step 1	P7-01, P7-08	Daily KEV + OSV + GitHub Advisory poll
Attack surface map	Ch. 5.5 step 2	P1-01, P3-01	AI registry + SBOM-to-CMDB correlation
PoC selection / exploit	Ch. 5.4 LiteLLM	P7-11, P7-10	Gateway hardening + 4h SLA
LotL movement	Ch. 8.1 kerberoast	P4-03, P4-05	IR-AI-02 SOAR playbook
Exfiltration	Ch. 8.1 220GB	P4-02, P4-04	Token baselines + DLP on embedding stores
Classified spill ingest	GOV-13	P7-07	Classification marker scan on RAG ingest
Insider AI tool misuse	GOV-06	P1-05, P5-03	Shadow-AI prohibition + training
Federal intel gap	GOV-01, GOV-09	P7-08	Redundant ISAC + commercial feeds
NVD publication lag	GOV-10	P7-08, P7-09	OSV trinity + transitive SBOM

Attack Step	Study Evidence	MVAP Control	Mitigation Action
Transitive dep chain	GOV-12	P7-09, P7-11b	Full tree SBOM + Starlette pin

5 3. Pillar 1 — Governance Mitigations

Study basis: AICSR-STUDY-2026-001* Section 4.1 (output/Cyber-Security-AI-Diligence-Research-Study.md), AICSR-STUDY-2026-001 Section 3.2 (output/Cyber-Security-AI-Diligence-Research-Study.md) P1 (25/27)*

5.1 3.1 Control Implementation

Control	Mitigation	Evidence artifact
P1-01	Maintain AI system registry with risk tier	Registry export quarterly
P1-02	NIST AI RMF Govern/Map alignment ¹	RMF mapping worksheet
P1-03	Classified-adjacent tier tagging	IL4/IL5/FedRAMP High flags
P1-04	NIST GenAI Profile for tier-1 ²	Profile gap analysis
P1-05	Shadow-AI prohibition with DNS/proxy enforcement	Blocklist + exception workflow

5.2 3.2 Governance Mitigation Playbook

- Week 1:** Inventory all LLM endpoints, gateways, RAG pipelines, fine-tuning jobs
- Week 2:** Assign risk tier (1/2/3) per data sensitivity and internet exposure
- Week 3:** Map each system to MVAP pillar owners (see Section 12)
- Ongoing:** Re-attestation against GOV register per simulated review cadence (AICSR-STUDY-2026-001 Section 6.9 (output/Cyber-Security-AI-Diligence-Research-Study.md))

6 4. Pillar 2 — Application & LLM Mitigations

Study basis: AICSR-STUDY-2026-001* Section 4.2 (output/Cyber-Security-AI-Diligence-Research-Study.md), Red/Blue failure modes AICSR-STUDY-2026-001 Section 8.1 (output/Cyber-Security-AI-Diligence-Research-Study.md)

¹NIST AI Risk Management Framework 1.0. <https://www.nist.gov/itl/ai-risk-management-framework>

²NIST Generative AI Profile (NIST.AI.600-1). <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>

6.1 4.1 Critical Mitigations (Post-Remediation)

Control	Mitigation	Target metric
P2-03	RAG corpus sanitization + poison detection	>=9/10 poisoned corpus block (prod: 11/12)
P2-04	Tool-calling per-user ACLs	Zero overprivileged export_tool
P2-01	OWASP LLM Top 10 test suite ³	100% tier-1 pre-deploy
P2-06	Output encoding / prompt boundary isolation	SAST gate in CI
P2-07	MCP endpoint authZ	PROXY_ADMIN pattern (align P7-11c)

6.2 4.2 Application Mitigation Sequence

1. Deploy guardrails on all user-facing LLM interfaces
2. Sanitize RAG ingest (classification markers per P7-07 if gov-adjacent)
3. Restrict tool permissions to least-privilege per identity
4. Run OWASP LLM Top 10 automated tests in CI/CD (*AICSR-STUDY-2026-001* Section 3.3 (output/Cyber-Security-AI-Diligence-Research-Study.md) gate 2)
5. Block internet-exposed Gradio/Streamlit staging endpoints (shadow-AI pattern from *AICSR-STUDY-2026-001* Section 8.1 (output/Cyber-Security-AI-Diligence-Research-Study.md))

7 5. Pillar 3 — Supply Chain Mitigations

Study basis: AICSR-STUDY-2026-001* Section 4.3 (output/Cyber-Security-AI-Diligence-Research-Study.md), AICSR-STUDY-2026-001 Section 5.7 (output/Cyber-Security-AI-Diligence-Research-Study.md)*

Control	Mitigation	Tooling
P3-01	SPDX/CycloneDX SBOM per release	syft, trivy
P3-02	Sigstore cosign on model artifacts	cosign verify in deploy gate
P3-03	Dependency pinning (pip-lock, package-lock)	Renovate with KEV priority
P3-04	Contractor/vendor NISP attestation	Annual for classified-adjacent
P3-05	SLSA L2 builds (L3 deferred v1.3) ⁴	GitHub Actions provenance

³OWASP Top 10 for LLM Applications 2025. <https://genai.owasp.org/llm-top-10/>

⁴SLSA Supply-chain Levels for Software Artifacts v1.0. <https://slsa.dev/spec/v1.0/>

Mitigation priority: Pre-deploy compromise (Oliver Hansen thesis, *AICSR-STUDY-2026-001* Section 4.3 (`output/Cyber-Security-AI-Diligence-Research-Study.md`)) — verify SBOM and signatures **before** model reaches inference cluster.

8 6. Pillar 4 — Detection & Response Mitigations

Study basis: AICSR-STUDY-2026-001* Section 4.4 (`output/Cyber-Security-AI-Diligence-Research-Study.md`), IR-AI-02 *AICSR-STUDY-2026-001* Section 8.2 (`output/Cyber-Security-AI-Diligence-Research-Study.md`)*

8.1 6.1 IR-AI-02 Unified Playbook

Phase	Mitigation	MTTC target
Detect	P4-02 token usage baselines; P4-03 LotL on svc-llm accounts	T+6m alert
Contain	VLAN isolation; API key rotation; Kerberos ticket revocation	T+14m avg
Eradicate	SBOM diff; Starlette/LiteLLM version verify	T+30m
Recover	Immutable backup restore (Phoenix protocol)	T+4h

8.2 6.2 Detection Engineering

- Log all AI API calls with user identity, model, token count, tool invocations
- MITRE ATT&CK mapping for AI service accounts
- SOAR integration: P4-02 alert P4-05 auto-containment human escalation
- Hunt Host header anomalies (BadHost pattern, *AICSR-STUDY-2026-001* Section 5.4 (`output/Cyber-Security-AI-Diligence-Research-Study.md`))

9 7. Pillar 5 — Human Layer Mitigations

Study basis: AICSR-STUDY-2026-001* Section 4.5 (`output/Cyber-Security-AI-Diligence-Research-Study.md`), Salt Typhoon social engineering *AICSR-STUDY-2026-001* Section 6.4 (`output/Cyber-Security-AI-Diligence-Research-Study.md`)*

Control	Mitigation	Frequency
P5-01	AI-phishing simulations	Quarterly
P5-02	Deepfake executive voice/video training	Semi-annual

Control	Mitigation	Frequency
P5-03	Public AI tool policy (no FOUO/classified)	Annual attestation
P5-04	Helpdesk shadow-AI reporting channel	Continuous

Target: Maintain <5% phishing click rate (study baseline: 4.2% vs 8% industry, *AICSR-STUDY-2026-001* Section 4.5 (output/Cyber-Security-AI-Diligence-Research-Study.md)).

10 8. Pillar 6 — Firmware Mitigations

Study basis: AICSR-STUDY-2026-001* Section 4.6 (output/Cyber-Security-AI-Diligence-Research-Study.md), v1.2 P6 tier-1+tier-2 *AICSR-STUDY-2026-001* Section 3.4 (output/Cyber-Security-AI-Diligence-Research-Study.md)

Control	Mitigation	Tier
P6-01	GPU driver KEV sweep	Tier-1 + Tier-2 (v1.2)
P6-02	TPM attestation on inference nodes	Tier-1 mandatory
P6-03	CUDA/NVIDIA bulletin monitoring	Tier-1 + Tier-2
P6-04	syzkaller kernel path (isolated lab)	Tier-1 only

Dissent note: Aether/Hex continue advocacy for all-tier firmware (*AICSR-STUDY-2026-001* Section 7.3 (output/Cyber-Security-AI-Diligence-Research-Study.md)) — enterprise should plan v1.3 expansion.

11 9. Pillar 7 — Zero-Day & Open-Source Mitigations

Study basis: AICSR-STUDY-2026-001* Section 5 (output/Cyber-Security-AI-Diligence-Research-Study.md), P7 playbook, v1.2 controls *AICSR-STUDY-2026-001* Section 3.4 (output/Cyber-Security-AI-Diligence-Research-Study.md)*

11.1 9.1 Daily Operations (P7-01, P7-08)

steps:

- `fetch`: CISA KEV JSON feed
- `fetch`: OSV API
- `fetch`: GitHub Advisory Database
- `match`: sbom.spdx + pip-lock.json + apt-manifest.txt
- `fail_on`: KEV_CVE in AI gateway watchlist

AI gateway watchlist: LiteLLM, Langflow, vLLM, Ollama, TGI, Ray Serve, BentoML, Starlette, FastAPI.

11.2 9.2 v1.2 Adopted Mitigations

11.2.1 P7-10 KEV Patch SLA

Tier	SLA	Escalation
Tier-1	4 hours	CISO page if breached
Tier-2	24 hours	SOC director
Tier-3	7 days	Risk acceptance memo

11.2.2 P7-11 Gateway Hardening Checklist

- P7-11a: Admin/MCP test endpoints internal-only
- P7-11b: Starlette $\geq 1.0.1$ ⁵
- P7-11c: PROXY_ADMIN on destructive routes
- P7-11d: Reverse proxy deny POST /mcp-rest/test/*
- P7-11e: Weekly KEV AI package scan

11.2.3 P7-09 Transitive SBOM (Conditional)

- Tier-1: full Python/npm tree by end of simulation follow-on cycle
- Diff on every release; flag Starlette/FastAPI/uvicorn chains
- Human sign-off on AI-generated patch proposals (P7-04)

11.3 9.3 Source Audit Program

Tier	Scope	Frequency	Sign-off
Tier-1	Top 50 deps + custom parsers	Every release	Hex + Maya
Tier-2	Direct ML deps	Monthly	Maya
Tier-3	Transitive via SBOM	Quarterly	Automated + spot check

12 10. Government & Classified-Adjacent Mitigations

Study basis: AICSR-STUDY-2026-001* Section 6 (output/Cyber-Security-AI-Diligence-Research-Study.md), GOV-01GOV-15 AICSR-STUDY-2026-001 Section 6.7 (output/Cyber-Security-AI-Diligence-Research-Study.simulated reviews AICSR-STUDY-2026-001 Section 6.8 (output/Cyber-Security-AI-Diligence-Research-Study.AICSR-STUDY-2026-001 Section 6.9 (output/Cyber-Security-AI-Diligence-Research-Study.md)*

⁵OSTIF — BadHost vulnerability in Starlette. <https://ostif.org/disclosing-the-badhost-vulnerability-in-starlette/>

12.1 10.1 Residual High Risks

Risk	Mitigation strategy	MVAP controls
GOV-02 Contractor violations	Independent NISP audit; no classified in prod RAG	P7-07, P3-04
GOV-03 Salt Typhoon	Assume telecom compromise; encrypt sensitive comms	P5-02, P4-03
GOV-06 Insider AI spill	Block public LLM for gov data; DLP	P1-05, P5-03
GOV-09 EO 14409 gap	Do not wait for clearinghouse; P7-08 redundancy	P7-08
GOV-10 NVD lag	OSV + GitHub Advisory trinity	P7-08, P7-09
GOV-13 Spill to RAG	Classification marker scan; air-gap embeddings	P7-07

12.2 10.2 Classified-Adjacent Checklist (P7-07)

- Embeddings air-gapped from internet-facing RAG
- Spill detection on ingest (classification marker scan)
- Contractor NISP compliance attestation annual
- No FOUO/classified in public AI tools ⁶
- Tabletop including spill sub-plot (*AICSR-STUDY-2026-001* Section 8.3 (output/Cyber-Security-AI-Diligence-Research-Study.md))

12.3 10.3 Intel Redundancy (P7-08)

Source	Role
CISA KEV ⁷	Binding patch priority
OSV	Fast open-source advisory
GitHub Advisory	Package-specific
FS-ISAC / sector ISAC	Sector context
NCSC-UK / ACSC	Five Eyes parallel
CSA / Horizon3 research	AI chain analysis ⁸

⁶TechCrunch — Acting CISA chief uploaded FOUO docs to ChatGPT. <https://techcrunch.com/2026/01/28/trumps-acting-cybersecurity-chief-uploaded-sensitive-government-docs-to-chatgpt/>

⁷CISA Known Exploited Vulnerabilities Catalog. <https://www.cisa.gov/known-exploited-vulnerabilities-catalog>

⁸Horizon3.ai — LiteLLM chained with Starlette BadHost RCE. <https://horizon3.ai/attack-research/vulnerabilities/cve-2026-42271-chained-with-cve-2026-48710/>

13 11. Implementation Roadmap

Phased program aligned with *AICSR-STUDY-2026-001* Section 3.1
(`output/Cyber-Security-AI-Diligence-Research-Study.md`) maturity levels. Timelines map
to **simulated May 2026** session sequence.

13.1 Phase 0 — Emergency (Sim. 2026-05-04)

Action	Control	Study driver
Audit all AI gateways for P7-11 compliance	P7-11	Ch. 5.4
Enable KEV daily sweep on AI deps	P7-01	Ch. 5.5
Deploy IR-AI-02 playbook draft	P4-05	Ch. 8.1
Block shadow-AI endpoints	P1-05	Ch. 8.1

13.2 Phase 1 — Foundation (Sim. 2026-05-05 through 2026-05-12)

Action	Control	Owner
Complete AI system registry	P1-01	Victor Vance
SBOM + cosign on tier-1 pipelines	P3-01, P3-02	Oliver Hansen
RAG sanitization to $\geq 9/10$	P2-03	Maya Patel
P7-08 redundant intel feeds live	P7-08	Sarah Jenkins
AI-phishing simulation baseline	P5-01	Susan Albright

13.3 Phase 2 — Operational L2 (Sim. 2026-05-08 through 2026-05-16)

Action	Control	Owner
CI/CD gates per <i>AICSR-STUDY-2026-001</i> Section 3.3 (<code>output/Cyber-Security-AI-Diligence-Research-Study.md</code>)	All pillars	Elena Rostova
P7-10 SLA timers + ticketing	P7-10	Kira
Purple-team exercise (P7-05 pattern)	P7-05	Arthur Vance
GPU driver attestation tier-1+2	P6	Aether
Government risk re-attestation	GOV	Marcus Thorne

13.4 Phase 3 — Hardening (Sim. 2026-05-20 through follow-on)

Action	Control	Target
P7-09 transitive SBOM tier-1	P7-09	Conditional adoption
SLSA L3 evaluation (v1.3 prep)	P3-05	v1.3 vote
P2-08 membership inference (v1.2.1)	P2-08	If adopted
Semi-annual purple-team	L2	PASS $\geq 18/27$ criteria

14 12. Metrics, Verification & Ownership

14.1 12.1 Key Performance Indicators

KPI	Target	Study reference
KEV tier-1 patch SLA compliance	$\geq 95\%$ within 4h	<i>AICSR-STUDY-2026-001</i> Section 3.4 (<code>output/Cyber-Security-AI-Diligence-Research-Study.md</code>)
IR-AI-02 MTTC	<15 minutes	<i>AICSR-STUDY-2026-001</i> Section 8.2 (<code>output/Cyber-Security-AI-Diligence-Research-Study.md</code>)

KPI	Target	Study reference
Poisoned RAG block rate	$\geq 91\%$ (11/12)	<i>AICSR-STUDY-2026-001</i> Section 8.2 (output/Cyber-Security-AI-Diligence-Research-Study.md)
AI-phishing click rate	$< 5\%$	<i>AICSR-STUDY-2026-001</i> Section 4.5 (output/Cyber-Security-AI-Diligence-Research-Study.md)
P7-11 gateway audit pass	100% tier-1	<i>AICSR-STUDY-2026-001</i> Section 6.9 (output/Cyber-Security-AI-Diligence-Research-Study.md)
Shadow-AI incidents	Zero prod breaches	<i>AICSR-STUDY-2026-001</i> Section 8.1 (output/Cyber-Security-AI-Diligence-Research-Study.md)

14.2 12.2 Pillar Ownership Matrix

Pillar	Primary owners	Rapid response
P1 Governance	Victor Vance, Marcus Thorne	—
P2 Application	Maya Patel, Synapse	Jax Reed
P3 Supply Chain	Oliver Hansen, Hex	—
P4 Detection	Sarah Jenkins, Aegis	Shield, Phoenix
P5 Human	Susan Albright, Mateo Silva	—
P6 Firmware	Aether, Hex	—
P7 Zero-Day	Kira, NullByte, Oliver	Viper, Ghost
Verification	Eleanor Vance	—
Facilitation	Arthur Vance	—

14.3 12.3 Verification Protocol

All mitigation claims must be logged in `sessions/verification-ledger.md` per *AICSR-STUDY-2026-001* Section Appendix A (output/Cyber-Security-AI-Diligence-Research-Study.md). Dialog evidence: `output/Boardroom-Complete-Dialog-Transcript.md` (AICSR-DLG-2026-001).

15 Index

- **BadHost mitigation** — Sections 9.2, 2
- **Classified spill** — Sections 10, 9
- **Dialog transcript** — Section 12.3; `output/Boardroom-Complete-Dialog-Transcript.md`

- **GOV-02 contractor risk** — Section 10
 - **GOV-13 RAG spill** — Sections 10, 4
 - **IR-AI-02** — Section 6
 - **KEV patch SLA (P7-10)** — Sections 9, 11
 - **LiteLLM chain** — Sections 2, 9
 - **May 2026 simulation** — Title page
 - **MVAP L2** — Sections 1, 11
 - **P7-11 gateway hardening** — Sections 9, 11
 - **Parent study AICSR-STUDY-2026-001** — All sections
 - **Shadow AI** — Sections 3, 4, 11
 - **Simulation disclosure** — Title page
 - **Transitive SBOM (P7-09)** — Sections 9, 11
 - **Verification ledger** — Section 12.3
-

16 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

16.1 Markdown Footnote Anchors

Archived reference article for AICSR-STUDY-2026-001 footnote corpus.