

Cyber-Security and AI Diligence Research

A Comprehensive Boardroom Study — MVAP Framework, Zero-Day Risk, and U.S.
Government Infrastructure Analysis

AI Cyber Security Research Boardroom

2026-05-31

Contents

| | | |
|----------|--|-----------|
| 1 | Cyber-Security and AI Diligence Research | 5 |
| 1.1 | A Comprehensive Boardroom Study | 5 |
| 1.2 | Table of Contents | 6 |
| 2 | 1. Executive Summary | 7 |
| 2.1 | Key Conclusions | 7 |
| 2.2 | Document Purpose | 8 |
| 3 | 2. Methodology — The Boardroom Framework | 9 |
| 3.1 | 2.1 Governance Structure | 9 |
| 3.2 | 2.2 Participant Composition | 9 |
| 3.3 | 2.3 Debate Protocol | 9 |
| 3.4 | 2.4 Session Corpus (Simulated May 2026 Chronology) | 9 |
| 3.4.1 | 2026-05-01 — Introductions | 9 |
| 3.4.2 | 2026-05-02 — Rounds 1-20 | 10 |
| 3.4.3 | 2026-05-04 — Pillar 2+4 Red/Blue | 10 |
| 3.4.4 | 2026-05-05 — Remediation Pass 1 | 10 |
| 3.4.5 | 2026-05-08 — L2 Maturity Review | 10 |
| 3.4.6 | 2026-05-09 — Remediation Pass 2 | 10 |
| 3.4.7 | 2026-05-12 — MVAP v1.1 Adoption | 10 |
| 3.4.8 | 2026-05-15 — P7-05 Tabletop | 10 |
| 3.4.9 | 2026-05-16 — Gov Risk Review (Sim. Q1) | 10 |
| 3.4.10 | 2026-05-20 — MVAP v1.2 Adoption | 10 |
| 3.4.11 | 2026-05-21 — Gov Risk Review (Sim. Q2) | 10 |
| 3.5 | 2.5 Simulation Disclosure and Source Inspection | 11 |
| 4 | 3. MVAP — Minimum Viable AI Diligence Program | 12 |
| 4.1 | 3.1 Maturity Levels | 12 |
| 4.2 | 3.2 Pillar Summary | 12 |
| 4.3 | 3.3 L2 Pipeline Gates | 13 |
| 4.4 | 3.4 MVAP v1.2 — Adopted Changes (Simulated 2026-05-20) | 13 |
| 4.4.1 | P7-11 AI Gateway Hardening (Mandatory) | 14 |
| 4.4.2 | P7-10 KEV Patch SLA | 14 |
| 4.4.3 | v1.2 L3 Pipeline Gates (Additive) | 14 |
| 5 | 4. Pillar Analysis — Controls and Evidence | 15 |
| 5.1 | 4.1 Pillar 1 — Governance & Inventory | 15 |
| 5.2 | 4.2 Pillar 2 — Application & LLM Security | 15 |

| | | |
|-----------|---|-----------|
| 5.3 | 4.3 Pillar 3 — Supply Chain | 15 |
| 5.4 | 4.4 Pillar 4 — Detection & Response | 15 |
| 5.5 | 4.5 Pillar 5 — Human Layer | 15 |
| 5.6 | 4.6 Pillar 6 — Firmware & Sub-Application | 16 |
| 5.7 | 4.7 Pillar 7 — Open-Source & Zero-Day | 16 |
| 6 | 5. Open-Source Zero-Day Risk — OS and Application Layers | 17 |
| 6.1 | 5.1 Historical Exploitation Timeline | 17 |
| 6.2 | 5.2 OS-Layer Targets | 17 |
| 6.3 | 5.3 Application-Layer Targets | 17 |
| 6.4 | 5.4 LiteLLM Case Study (June 2026) | 17 |
| 6.5 | 5.5 AI-Accelerated Offensive Chain | 18 |
| 6.6 | 5.6 v1.2 Adopted Controls | 18 |
| 6.7 | 5.7 Source Code Evaluation Matrix | 18 |
| 7 | 6. U.S. Government and Classified Infrastructure Risk | 19 |
| 7.1 | 6.1 Thesis | 19 |
| 7.2 | 6.2 CISA Governance Collapse | 19 |
| 7.3 | 6.3 Classified Contractor Failures | 19 |
| 7.4 | 6.4 Salt Typhoon — Ongoing Espionage | 20 |
| 7.5 | 6.5 Policy Counterweight — EO 14409 | 20 |
| 7.6 | 6.6 Critical Infrastructure | 20 |
| 7.7 | 6.7 Living Risk Register (GOV-01 GOV-15) | 20 |
| 7.8 | 6.8 Government Risk Review — Simulated Q1 (2026-05-16) | 21 |
| 7.9 | 6.9 Government Risk Review — Simulated Q2 (2026-05-21) | 22 |
| 8 | 7. Deliberation Outcomes — Twenty Rounds and Subsequent Sessions | 23 |
| 8.1 | 7.1 Round Arc (Summary) | 23 |
| 8.2 | 7.2 Five Adopted MVAP Pillars (Round 19) | 23 |
| 8.3 | 7.3 Material Dissent | 24 |
| 8.4 | 7.4 Round-by-Round Detail (Rounds 120) | 24 |
| 8.5 | 7.5 Post-Round Session Outcomes | 25 |
| 9 | 8. Red Team / Blue Team Exercises and Remediation | 26 |
| 9.1 | 8.1 72-Hour Exercise (Simulated 2026-05-04) | 26 |
| 9.2 | 8.2 Remediation Sprint (Closed Simulated 2026-05-09) | 26 |
| 9.3 | 8.3 P7-05 Tabletop (Simulated 2026-05-15) | 26 |
| 10 | 9. Participant Profiles | 27 |
| 10.1 | 9.1 Governance — Moderator & Court Reporter | 27 |
| | 10.1.1 Arthur Vance | 27 |
| | 10.1.2 Eleanor Vance | 27 |
| 10.2 | 9.2 Executive Compliance Officers (CCO) | 28 |
| | 10.2.1 Marcus Thorne | 28 |
| | 10.2.2 Dr. Elena Rostova | 28 |
| 10.3 | 9.3 CISSP Tier — Advanced Strategy & Architecture | 29 |
| | 10.3.1 Victor Vance | 29 |
| | 10.3.2 Sarah Jenkins | 29 |
| | 10.3.3 Tariq Al-Jamil | 30 |
| 10.4 | 9.4 SSCP Tier — Hands-On Operations | 30 |

| | | |
|-----------|---|-----------|
| 10.4.1 | Chloe Mitchell | 30 |
| 10.4.2 | Liam O'Connor | 30 |
| 10.4.3 | Maya Patel | 31 |
| 10.5 | 9.5 CC Tier — Entry-Level Perspectives | 31 |
| 10.5.1 | Jordan Taylor | 31 |
| 10.5.2 | Susan Albright | 32 |
| 10.5.3 | Devonne Brooks | 32 |
| 10.6 | 9.6 Zero-Day Hunters — Elite Vulnerability Researchers | 33 |
| 10.6.1 | Rene Dupont (“Aether”) | 33 |
| 10.6.2 | Siddharth Nair (“NullByte”) | 33 |
| 10.6.3 | Zoe Kruger (“Cipher”) | 33 |
| 10.6.4 | Kenji Sato (“Synapse”) | 34 |
| 10.7 | 9.7 Code Hackers — Practical Exploitation Specialists | 34 |
| 10.7.1 | Jaxson “Jax” Reed | 34 |
| 10.7.2 | Ekaterina Petrova (“Kira”) | 35 |
| 10.7.3 | Mateo Silva | 35 |
| 10.7.4 | Alaric Vance (“Hex”) | 36 |
| 10.7.5 | Aisha Nwosu | 36 |
| 10.7.6 | Samuel Cohen (“SQL_Sam”) | 36 |
| 10.7.7 | Oliver Hansen | 37 |
| 10.7.8 | Dimitri Volkov (“GridLock”) | 37 |
| 10.8 | 9.8 Red Rapid Response — Strike Unit | 38 |
| 10.8.1 | Cassandra Cross (“Viper”) | 38 |
| 10.8.2 | Ji-Hoon Park (“Ghost”) | 38 |
| 10.8.3 | Dominic Kruse (“Payload”) | 38 |
| 10.9 | 9.9 Blue Rapid Response — Incident Defenders | 39 |
| 10.9.1 | Elena Rostova Jr. (“Aegis”) | 39 |
| 10.9.2 | Marcus “Mal” Sterling (“Shield”) | 39 |
| 10.9.3 | Amara Okafor (“Phoenix”) | 40 |
| 11 | Appendix A — Verification Ledger Summary | 41 |
| 12 | Appendix B — Session and Source File Index | 42 |
| 12.1 | B.1 mvap/MVAP-SPECIFICATION-v1.1.md | 42 |
| 12.2 | B.2 mvap/MVAP-SPECIFICATION-v1.2-DRAFT.md | 42 |
| 12.3 | B.3 mvap/ZERO-DAY-OPEN-SOURCE-RISK-ASSESSMENT.md | 42 |
| 12.4 | B.4 mvap/P7-IMPLEMENTATION-PLAYBOOK.md | 42 |
| 12.5 | B.5 sessions/verification-ledger.md | 43 |
| 12.6 | B.6 roster.yaml | 43 |
| 12.7 | B.7 participants/*.md | 43 |
| 12.8 | B.8 output/STUDY-MANIFEST.yaml | 43 |
| 12.9 | B.9 output/CONTINUE.md | 43 |
| 12.10 | B.10 output/Cyber-Security-AI-Diligence-Research-Study.md | 43 |
| 12.11 | B.11 output/MVAP-Complete-Mitigation-Strategy.md | 43 |
| 12.12 | B.12 output/Boardroom-Complete-Dialog-Transcript.md | 43 |
| 12.13 | B.13 output/MITIGATION-MANIFEST.yaml | 43 |
| 12.14 | B.14 output/DIALOG-MANIFEST.yaml | 44 |
| 12.15 | B.15 output/Boardroom-Comprehensive-Abstracts.md | 44 |
| 12.16 | B.16 output/ABSTRACT-MANIFEST.yaml | 44 |
| 12.17 | B.17 output/CONTINUE-MITIGATION.md | 44 |

| | | |
|-----------|---|-----------|
| 12.18B.18 | output/CONTINUE-DIALOG.md | 44 |
| 12.19B.19 | PROJECT.md | 44 |
| 12.20B.20 | output/How-The-Research-Was-Done.md | 44 |
| 12.21B.21 | output/RESEARCH-MATERIALS-INDEX.md | 44 |
| 12.22B.22 | output/references/CONTINUE-REFERENCES.md | 44 |
| 12.23B.23 | output/references/REFERENCE-MANIFEST.yaml | 44 |
| 12.24B.24 | output/references/REFERENCE-AVAILABILITY-REPORT.md | 45 |
| 12.25B.25 | output/references/ | 45 |
| 12.26B.26 | sessions/VOTE-RECORD.md | 45 |
| 12.27B.27 | sessions/2026-06-22-cybersecurity-ai-diligence-introductions.md | 45 |
| 12.28B.28 | sessions/2026-06-22-cybersecurity-ai-diligence-rounds-1-20.md | 45 |
| 12.29B.29 | sessions/2026-06-22-mvap-pillar-2-4-red-blue-exercise.md | 45 |
| 12.30B.30 | sessions/2026-07-22-mvap-p2-03-p4-05-remediation-validation-1.md | 45 |
| 12.31B.31 | sessions/2026-09-20-mvap-level-2-maturity-review.md | 45 |
| 12.32B.32 | sessions/2026-09-21-mvap-p2-03-p4-05-remediation-validation-2.md | 45 |
| 12.33B.33 | sessions/2026-12-20-mvap-v1-1-backlog-zero-day-government-risk.md | 45 |
| 12.34B.34 | sessions/2027-03-20-p7-05-zero-day-tabletop.md | 46 |
| 12.35B.35 | sessions/2027-03-20-quarterly-government-risk-review.md | 46 |
| 12.36B.36 | sessions/2027-06-20-mvap-v1-2-adoption.md | 46 |
| 12.37B.37 | sessions/2027-06-20-quarterly-government-risk-review-q2.md | 46 |
| 13 | Index | 47 |
| 14 | Footnotes and Reference Bibliography | 49 |
| 14.1 | Markdown Footnote Anchors | 51 |

Chapter 1

Cyber-Security and AI Diligence Research

1.1 A Comprehensive Boardroom Study

Minimum Viable AI Diligence Program (MVAP) Zero-Day Open-Source Risk Classified Infrastructure Governance

| | |
|-------------------|---|
| Document ID | AICSR-STUDY-2026-001 |
| Version | 1.2 |
| Publication Date | 2026-05-31 |
| Simulation Window | May 2026 (compressed chronology) |
| Moderator | Arthur Vance |
| Court Reporter | Eleanor Vance |
| Participants | 31 (27 voting deliberants) |
| Sessions Analyzed | 11 boardroom transcripts |
| Specification | MVAP v1.1 adopted 2026-05-12; v1.2 adopted 2026-05-20 |

*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. Executive Summary
2. Methodology — The Boardroom Framework
3. MVAP — Minimum Viable AI Diligence Program
4. Pillar Analysis — Controls and Evidence
5. Open-Source Zero-Day Risk — OS and Application Layers
6. U.S. Government and Classified Infrastructure Risk
7. Deliberation Outcomes — Twenty Rounds and Subsequent Sessions
8. Red Team / Blue Team Exercises and Remediation
9. Participant Profiles
10. Appendix A — Verification Ledger Summary
11. Appendix B — Session and Source File Index
12. Index

Chapter 2

1. Executive Summary

This study documents the complete findings of the **AI Cyber Security Research Boardroom** — a structured, multi-agent **simulation** convening thirty-one configured cybersecurity specialist profiles to research **Cyber-Security and AI Diligence**. Over twenty debate rounds and eleven recorded sessions within a **compressed May 2026 simulation window**, the board developed, validated, and operationalized the **Minimum Viable AI Diligence Program (MVAP)**. All underlying transcripts and specifications remain in the repository for inspection.

2.1 Key Conclusions

1. **MVAP v1.1 adopted (20/27)** — Five core pillars (Governance, Application/LLM Security, Supply Chain, Detection/Response, Human Layer) plus Pillar 6 (Firmware) and Pillar 7 (Open-Source Zero-Day) provide necessary but not sufficient diligence for enterprise LLM deployment.
2. **MVAP L2 operational (18/27)** — Purple-team semi-annual validation, CI/CD gates, and production SOAR playbook IR-AI-02 with 14-minute mean time to containment.
3. **Government risk is structural** — CISA leadership vacuum ¹, 815 classified contractor violations ², and ongoing Salt Typhoon persistence ³ mean enterprises cannot rely on federal backstop for AI/zero-day intelligence.
4. **AI gateways are the new edge appliances** — LiteLLM KEV chain (CVE-2026-42271 + CVE-2026-48710) demonstrates sub-36-hour exploitation timelines ⁴ ⁵.
5. **MVAP v1.2 adopted (simulated 2026-05-20)** — P7-10 4-hour KEV SLA (19/27), P7-11 AI gateway hardening (21/27), P6 tier-1+tier-2 firmware (18/27); P7-09 transitive SBOM conditional (17/27); SLSA L3 deferred to v1.3.

¹CSA Research Note — CISA Leadership Governance Vacuum (2026-04-24). <https://labs.cloudsecurityalliance.org/research/csa-research-note-cisa-leadership-governance-vacuum-20260424/>

²GAO-26-107861 — 815 Classified Contractor Security Violations. <https://www.gao.gov/products/gao-26-107861>

³GCA — Salt Typhoon Across the Internet. <https://globalcyberalliance.org/new-report-salt-typhoon-across-the-internet/>

⁴CISA Alert — CVE-2026-42271 LiteLLM added to KEV. <https://www.cisa.gov/news-events/alerts/2026/06/08/cisa-adds-two-known-exploited-vulnerabilities-catalog>

⁵Horizon3.ai — LiteLLM chained with Starlette BadHost RCE. <https://horizon3.ai/attack-research/vulnerabilities/cve-2026-42271-chained-with-cve-2026-48710/>

6. **Simulated government risk reviews (2026-05-16, 2026-05-21)** — GOV-02 remains Critical+; EO 14409 clearinghouse partially operational; compliant P7-11 gateways close GOV-11 exploitation window.

2.2 Document Purpose

This report serves as the authoritative synthesis for practitioners, compliance officers, and researchers. Footnotes link to primary sources verified by Eleanor Vance. For continuation and regeneration, see `output/STUDY-MANIFEST.yaml` and `output/CONTINUE.md`.

Chapter 3

2. Methodology — The Boardroom Framework

3.1 2.1 Governance Structure

| Role | Name | Function |
|----------------|---------------|--|
| Moderator | Arthur Vance | Frames questions; enforces positive/negative debate; no vote |
| Court Reporter | Eleanor Vance | Verification ledger; footnotes; strips hallucinations |

3.2 2.2 Participant Composition

Twenty-seven voting participants across nine groups: Compliance (2), CISSP (3), SSCP (3), CC (3), Zero-Day (4), Code Hackers (8), Red Rapid (3), Blue Rapid (3).

3.3 2.3 Debate Protocol

Each round requires every voter to deliver **one positive point** and **one negative point**. Majority consensus = 14/27. Rapid Response units add tactical timelines. Speculation is labeled [Projected Speculation] with profile justification.

3.4 2.4 Session Corpus (Simulated May 2026 Chronology)

3.4.1 2026-05-01 — Introductions

- **Simulated date:** 2026-05-01
- **Source file:** sessions/2026-06-22-cybersecurity-ai-diligence-introductions.md

3.4.2 2026-05-02 — Rounds 1-20

- Simulated date: 2026-05-02
- Source file: sessions/2026-06-22-cybersecurity-ai-diligence-rounds-1-20.md

3.4.3 2026-05-04 — Pillar 2+4 Red/Blue

- Simulated date: 2026-05-04
- Source file: sessions/2026-06-22-mvap-pillar-2-4-red-blue-exercise.md

3.4.4 2026-05-05 — Remediation Pass 1

- Simulated date: 2026-05-05
- Source file: sessions/2026-07-22-mvap-p2-03-p4-05-remediation-validation-1.md

3.4.5 2026-05-08 — L2 Maturity Review

- Simulated date: 2026-05-08
- Source file: sessions/2026-09-20-mvap-level-2-maturity-review.md

3.4.6 2026-05-09 — Remediation Pass 2

- Simulated date: 2026-05-09
- Source file: sessions/2026-09-21-mvap-p2-03-p4-05-remediation-validation-2.md

3.4.7 2026-05-12 — MVAP v1.1 Adoption

- Simulated date: 2026-05-12
- Source file: sessions/2026-12-20-mvap-v1-1-backlog-zero-day-government-risk.md

3.4.8 2026-05-15 — P7-05 Tabletop

- Simulated date: 2026-05-15
- Source file: sessions/2027-03-20-p7-05-zero-day-tabletop.md

3.4.9 2026-05-16 — Gov Risk Review (Sim. Q1)

- Simulated date: 2026-05-16
- Source file: sessions/2027-03-20-quarterly-government-risk-review.md

3.4.10 2026-05-20 — MVAP v1.2 Adoption

- Simulated date: 2026-05-20
- Source file: sessions/2027-06-20-mvap-v1-2-adoption.md

3.4.11 2026-05-21 — Gov Risk Review (Sim. Q2)

- Simulated date: 2026-05-21
- Source file: sessions/2027-06-20-quarterly-government-risk-review-q2.md

3.5 2.5 Simulation Disclosure and Source Inspection

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`.

Inspectability checklist:

- Deliberation text: `output/Boardroom-Complete-Dialog-Transcript.md` (AICSR-DLG-2026-001)
 - Synthesis and analysis: `output/Cyber-Security-AI-Diligence-Research-Study.md`
 - Implementation guidance: `output/MVAP-Complete-Mitigation-Strategy.md`
 - Methodology: `output/How-The-Research-Was-Done.md` (AICSR-METHOD-2026-001)
 - Materials index: `output/RESEARCH-MATERIALS-INDEX.md` (AICSR-MATINDEX-2026-001)
 - Reference archive: `output/references/` (AICSR-REF-INDEX-2026-001)
 - Formal votes: `sessions/VOTE-RECORD.md`
 - Claim verification: `sessions/verification-ledger.md`
 - Project metadata: `PROJECT.md`
 - Full regeneration: see `PROJECT.md` or `output/CONTINUE.md`
-

Chapter 4

3. MVAP — Minimum Viable AI Diligence Program

MVAP defines the **minimum** diligence before production LLM deployment. It maps to NIST AI RMF¹ and OWASP LLM Top 10².

4.1 3.1 Maturity Levels

| Level | Definition | Status |
|-------|--|----------------------------------|
| L1 | Documented evidence all pillars | Certified simulated 2026-05-08 |
| L2 | L1 + purple-team + CI/CD gates | Operational simulated 2026-05-09 |
| L3 | L2 + SLSA L3 + 24h KEV + firmware all tiers | Planned v1.3 |

4.2 3.2 Pillar Summary

| Pillar | Vote | Core mandate |
|-----------------|-------|---|
| P1 Governance | 25/27 | AI registry, NIST Govern/Map, shadow-AI prohibition |
| P2 App/LLM | 24/27 | OWASP LLM Top 10, guardrails, RAG sanitization, tool ACLs |
| P3 Supply Chain | 23/27 | SBOM, model signing, dependency pinning |
| P4 Detection | 24/27 | AI API logging, token baselines, IR-AI-02 SOAR |

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

²OWASP Top 10 for LLM Applications 2025. <https://genai.owasp.org/llm-top-10/>

| Pillar | Vote | Core mandate |
|-------------|-------|--|
| P5 Human | 22/27 | AI-phishing/deepfake simulations quarterly |
| P6 Firmware | 17/27 | GPU driver KEV, TPM attestation (tier-1 mandatory) |
| P7 Zero-Day | 20/27 | KEV sweep, native SAST, source audit, intel redundancy |

4.3 3.3 L2 Pipeline Gates

1. P1 registry check
2. P2 SAST + poisoned-corpus $\geq 9/10$
3. P3 SBOM + cosign
4. P7-01 KEV sweep
5. P7-04 human sign-off on AI-generated patches

4.4 3.4 MVAP v1.2 — Adopted Changes (Simulated 2026-05-20)

The board ratified v1.2 following Q1 2027 tabletop gaps and quarterly government risk review. Key votes:

| Control | Requirement | Vote | Status |
|---------|---|-------|---|
| P7-09 | Full transitive dependency tree SBOM (Python/npm) | 17/27 | Conditional — tier-1 mandatory by Q3 2027 |
| P7-10 | KEV patch SLA: 4h tier-1, 24h tier-2, 7d tier-3 | 19/27 | Adopted |
| P7-11 | AI gateway/router hardening (LiteLLM-class) | 21/27 | Adopted |
| P6 | Firmware/GPU mandatory tier-1 and tier-2 | 18/27 | Adopted |
| P2-08 | Membership inference — tier-1 mandatory | 16/27 | Deferred v1.2.1 |
| SLSA L3 | Tier-1 builds | 16/27 | Deferred v1.3 |

4.4.1 P7-11 AI Gateway Hardening (Mandatory)

Applies to any model router/proxy (LiteLLM, Langflow, custom):

| Control | Requirement |
|---------|---|
| P7-11a | No internet-exposed admin/MCP test endpoints |
| P7-11b | Starlette >=1.0.1 or equivalent BadHost mitigation |
| P7-11c | PROXY_ADMIN role on all destructive endpoints |
| P7-11d | Block POST /mcp-rest/test/* at reverse proxy default-deny |
| P7-11e | Weekly automated check against CISA KEV AI package list |

Reference incidents: CVE-2026-42271, CVE-2026-48710, CVE-2026-42208 ³ ⁴.

4.4.2 P7-10 KEV Patch SLA

| Tier | SLA | Measurement |
|-----------------|----------------------|---------------------------------|
| Tier-1 Critical | 4 hours from KEV add | Ticket timestamp to prod deploy |
| Tier-2 High | 24 hours | Same |
| Tier-3 | 7 days | Risk acceptance memo required |

4.4.3 v1.2 L3 Pipeline Gates (Additive)

6. P7-09 transitive SBOM diff on every release
7. P7-10 SLA timer starts on KEV JSON feed poll
8. P7-11 gateway config audit (weekly)
9. P6 GPU driver attestation (tier-1 + tier-2)

³CISA Alert — CVE-2026-42271 LiteLLM added to KEV. <https://www.cisa.gov/news-events/alerts/2026/06/08/cisa-adds-two-known-exploited-vulnerabilities-catalog>

⁴Horizon3.ai — LiteLLM chained with Starlette BadHost RCE. <https://horizon3.ai/attack-research/vulnerabilities/cve-2026-42271-chained-with-cve-2026-48710/>

Chapter 5

4. Pillar Analysis — Controls and Evidence

5.1 4.1 Pillar 1 — Governance & Inventory

Controls P1-01 through P1-05 require AI system registry (47 systems documented at L2 review), risk tiering, and NIST GenAI Profile alignment ¹. Shadow-AI prohibition addresses staging endpoint breaches discovered in Red/Blue exercises.

5.2 4.2 Pillar 2 — Application & LLM Security

Critical findings: Initial exercise failed P2-03 (7/10 poisoned RAG) and P2-04 (overprivileged export_tool). Post-remediation: 11/12 production (91.7%). OWASP LLM01 Prompt Injection ² and LLM06 Excessive Agency mapped to LiteLLM MCP endpoints.

5.3 4.3 Pillar 3 — Supply Chain

SLSA Level 2 aspirational for tier-1 [^slasa]. Sigstore cosign on model registry. Oliver Hansen leads supply-chain infiltration perspective — pre-deploy compromise beats post-deploy hacking.

5.4 4.4 Pillar 4 — Detection & Response

IR-AI-02 unified playbook chains P4-02 token alerts to P4-03 LotL detection. Production MTTC: 14 minutes average (L2 target <15m). MITRE ATT&CK mapping for AI service accounts.

5.5 4.5 Pillar 5 — Human Layer

Q2+Q3 AI-phishing click rate 4.2% vs 8% industry baseline. Deepfake executive training mandatory. Mateo Silva and Susan Albright co-own human adversarial testing.

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

²OWASP Top 10 for LLM Applications 2025. <https://genai.owasp.org/llm-top-10/>

5.6 4.6 Pillar 6 — Firmware & Sub-Application

Elevated from advisory to tier-1 mandatory. GPU driver CVE cadence exceeds application patch SLAs — Aether and Hex dissent on tier-2/3 coverage continues.

5.7 4.7 Pillar 7 — Open-Source & Zero-Day

See Chapter 5. P7-08 mandates redundant intel due to CISA capacity loss ³.

³CSA Research Note — CISA Leadership Governance Vacuum (2026-04-24).
<https://labs.cloudsecurityalliance.org/research/csa-research-note-cisa-leadership-governance-vacuum-20260424/>

Chapter 6

5. Open-Source Zero-Day Risk — OS and Application Layers

6.1 5.1 Historical Exploitation Timeline

From Heartbleed (2014) through Log4Shell (2021) to Salt Typhoon (2024) and LiteLLM KEV entries (2026), the board documented that **AI inference stacks inherit the full OSS vulnerability life-cycle** — with AI accelerating discovery faster than NVD publication ¹.

6.2 5.2 OS-Layer Targets

Linux kernel (io_uring, eBPF), glibc/OpenSSL, container runtimes (runc), GPU/CUDA drivers. Evaluation: syzkaller, KEV kernel sweep, TPM attestation.

6.3 5.3 Application-Layer Targets

LiteLLM, Langflow, vLLM, Starlette/FastAPI, PyTorch native extensions, RAG parsers. Evaluation: CodeQL, fuzzing, P7-11 gateway hardening.

6.4 5.4 LiteLLM Case Study (June 2026)

| CVE | Impact | Timeline |
|----------------|---------------------------------|-----------------------------|
| CVE-2026-42208 | SQL injection | Exploited <36h ² |
| CVE-2026-42271 | Command injection via MCP stdio | KEV same day ³ |

¹CSA Whitepaper — NVD Infrastructure Crisis & AI Vulnerability Discovery. https://labs.cloudsecurityalliance.org/wp-content/uploads/2026/05/CSA_whitepaper_NVD_infrastructure_crisis_AI_vulnerability_dcsa-styled.pdf

²The Hacker News — LiteLLM CVE-2026-42208 exploited within 36h. <https://thehackernews.com/2026/04/litellm-cve-2026-42208-sql-injection.html>

³CISA Alert — CVE-2026-42271 LiteLLM added to KEV. <https://www.cisa.gov/news-events/alerts/2026/06/08/cisa-adds-two-known-exploited-vulnerabilities-catalog>

| CVE | Impact | Timeline |
|----------------|-------------------------------|--|
| CVE-2026-48710 | Starlette BadHost auth bypass | Chains to unauthenticated RCE CVSS 10.0 ⁴ ⁵ |

6.5 5.5 AI-Accelerated Offensive Chain

1. CVE feed scrape
 2. Attack surface map
 3. PoC selection
 4. LotL movement
 5. Exfiltration.
- Defensive counter: P7-01 + P7-10 SLA + P4-05.

6.6 5.6 v1.2 Adopted Controls

- **P7-09** — Full transitive dependency tree SBOM (conditional; tier-1 mandatory Q3 2027)
- **P7-10** — 4h KEV patch SLA (tier-1); 24h (tier-2) — **adopted**
- **P7-11** — AI gateway hardening checklist — **adopted**

6.7 5.7 Source Code Evaluation Matrix

| Component | Method | Zero-day signal | Owner |
|-------------------|-----------------------------|------------------------|-------------|
| Linux kernel | syzkaller + KEV sweep | EternalBlue class | Aether |
| OpenSSL/glibc | Distro feeds + SAST | Heartbleed class | Hex |
| CUDA/NVIDIA | Vendor bulletin + P6 driver | GPU driver KEV cadence | Aether |
| LiteLLM/Langchain | P7-11 + dep tree | 2026 KEV chain | Oliver/Maya |
| Starlette/FastAPI | In >=1.0.1; Host monitoring | BadHost 2026 | Maya |
| Kubernetes | CIS benchmark + CVE | API server escapes | Chloe |
| GitHub Actions | SLSA + secret scope | Supply chain | Oliver |

⁴Horizon3.ai — LiteLLM chained with Starlette BadHost RCE. <https://horizon3.ai/attack-research/vulnerabilities/cve-2026-42271-chained-with-cve-2026-48710/>

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

Chapter 7

6. U.S. Government and Classified Infrastructure Risk

7.1 6.1 Thesis

The United States is experiencing **loss of effective control** over classified and classified-adjacent infrastructure security through converging failures: federal governance vacuum, contractor program collapse, nation-state persistence, and insider AI tool misuse.

7.2 6.2 CISA Governance Collapse

CSA documents 16+ months without Senate-confirmed CISA director, ~32% workforce reduction, \$707M proposed cuts, MS-ISAC defunding ¹. Acting leadership reportedly uploaded FOUO documents to public ChatGPT ².

7.3 6.3 Classified Contractor Failures

GAO-26-107861: **815 security violations, 1,032 open vulnerabilities, <40% inspection rate**, classified spills to unclassified systems ³.

¹CSA Research Note — CISA Leadership Governance Vacuum (2026-04-24). <https://labs.cloudsecurityalliance.org/research/csa-research-note-cisa-leadership-governance-vacuum-20260424/>

²TechCrunch — Acting CISA chief uploaded FOUO docs to ChatGPT. <https://techcrunch.com/2026/01/28/trumps-acting-cybersecurity-chief-uploaded-sensitive-government-docs-to-chatgpt/>

³GAO-26-107861 — 815 Classified Contractor Security Violations. <https://www.gao.gov/products/gao-26-107861>

7.4 6.4 Salt Typhoon — Ongoing Espionage

80+ nations targeted; telecom wiretap access; House Committee emails January 2026 ⁴; FBI confirms ongoing February 2026 ⁵; military network diagram exposure reported ^{6 7}.

7.5 6.5 Policy Counterweight — EO 14409

June 2, 2026 Executive Order establishes AI cybersecurity clearinghouse and frontier model benchmarking ⁸. Board assesses **implementation gap (GOV-09)** against depleted CISA capacity.

7.6 6.6 Critical Infrastructure

Water/wastewater 170,000 systems; EPA authority gaps ⁹. OT AI separate IEC 62443 tier ¹⁰.

7.7 6.7 Living Risk Register (GOV-01 GOV-15)

Maintained in `mvap/ZERO-DAY-OPEN-SOURCE-RISK-ASSESSMENT.md` v2.0.

| Risk ID | Description | Severity | Status (Sim. 2026-05-21) |
|---------|--|-----------|--------------------------|
| GOV-01 | CISA leadership vacuum | Critical | ONGOING |
| GOV-02 | 815+ contractor violations; <40% inspections | Critical+ | ONGOING ^{11 12} |
| GOV-03 | Salt Typhoon persistence | Critical | ONGOING ¹³ |
| GOV-04 | MS-ISAC defunding | High | ONGOING ¹⁴ |
| GOV-05 | AI agent CVE weaponization | High | ONGOING |

⁴NJCCIC — Salt Typhoon targets House Committee emails. <https://www.cyber.nj.gov/Home/Components/News/News/1935/214>

⁵CyberScoop — FBI confirms Salt Typhoon still ongoing (Feb 2026). <https://cyberscoop.com/fbi-salt-typhoon-ongoing-threat-cybertalks-2026/>

⁶GCA — Salt Typhoon Across the Internet. <https://globalcyberalliance.org/new-report-salt-typhoon-across-the-internet/>

⁷CISA Advisory AA25-239A — Salt Typhoon. <https://www.cisa.gov/news-events/cybersecurity-advisories/aa25-239a>

⁸White House EO 14409 — AI Innovation and Security (2026-06-02). <https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/>

⁹GAO-26-109159 — Water Sector Cybersecurity. <https://www.gao.gov/products/gao-26-109159>

¹⁰ISA/IEC 62443 Industrial Cybersecurity Standards. <https://www.isa.org/standards-and-publications/isa-standards/isa-iec-62443-series-of-standards>

¹¹GAO-26-107861 — 815 Classified Contractor Security Violations. <https://www.gao.gov/products/gao-26-107861>

¹²Bloomberg Government — 815 classified data violations summary. <https://news.bgov.com/bloomberg-government-news/us-companies-had-815-classified-data-violations-gao-finds>

¹³CyberScoop — FBI confirms Salt Typhoon still ongoing (Feb 2026). <https://cyberscoop.com/fbi-salt-typhoon-ongoing-threat-cybertalks-2026/>

¹⁴StateScoop — CISA ending MS-ISAC support. <https://statescoop.com/cisa-confirms-its-ending-ms-isac-support/>

| Risk ID | Description | Severity | Status (Sim. 2026-05-21) |
|---------|--------------------------------------|----------|----------------------------------|
| GOV-06 | FOUO spill via public AI tools | High | ONGOING ¹⁵ |
| GOV-07 | Water/OT authority gaps | High | ONGOING ¹⁶ |
| GOV-09 | EO 14409 vs CISA capacity gap | High | PARTIAL MITIGATION ¹⁷ |
| GOV-10 | NVD backlog vs AI discovery | High | ONGOING ¹⁸ |
| GOV-11 | AI gateway KEV exploitation | Critical | MITIGATED (P7-11 compliant) |
| GOV-12 | BadHost transitive chains | High | MITIGATED (P7-09/P7-11) |
| GOV-13 | Classified spill into AI RAG | Critical | ONGOING |
| GOV-14 | Military diagram/credential exposure | High | ONGOING |
| GOV-15 | Frontier model gov access (EO 14409) | Medium | MONITORING |

7.8 6.8 Government Risk Review — Simulated Q1 (2026-05-16)

Ratified risk register v2.0 at **21/27**. New entries GOV-09 through GOV-15. Key findings:

- EO 14409 clearinghouse mandated within 30 days — contrasts with CSA-documented CISA collapse ^{19 20}
- LiteLLM-class packages elevated to KEV-class — drives P7-11 scope expansion
- NVD backlog validates P7-08 redundant feeds (OSV + GitHub Advisory + KEV trinity) ²¹
- GAO-26-107861 recommendations remain **Open** — GOV-02 upgraded to Critical+

¹⁵TechCrunch — Acting CISA chief uploaded FOUO docs to ChatGPT. <https://techcrunch.com/2026/01/28/trumps-acting-cybersecurity-chief-uploaded-sensitive-government-docs-to-chatgpt/>

¹⁶GAO-26-109159 — Water Sector Cybersecurity. <https://www.gao.gov/products/gao-26-109159>

¹⁷White House EO 14409 — AI Innovation and Security (2026-06-02). <https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/>

¹⁸CSA Whitepaper — NVD Infrastructure Crisis & AI Vulnerability Discovery. https://labs.cloudsecurityalliance.org/wp-content/uploads/2026/05/CSA_whitepaper_NVD_infrastructure_crisis_AI_vulnerability_discovery_csa-styled.pdf

¹⁹White House EO 14409 — AI Innovation and Security (2026-06-02). <https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/>

²⁰CSA Research Note — CISA Leadership Governance Vacuum (2026-04-24). <https://labs.cloudsecurityalliance.org/research/csa-research-note-cisa-leadership-governance-vacuum-20260424/>

²¹CSA Whitepaper — NVD Infrastructure Crisis & AI Vulnerability Discovery. https://labs.cloudsecurityalliance.org/wp-content/uploads/2026/05/CSA_whitepaper_NVD_infrastructure_crisis_AI_vulnerability_discovery_csa-styled.pdf

7.9 6.9 Government Risk Review — Simulated Q2 (2026-05-21)

Vote: Risk register reaffirmed **20/27**; MVAP v1.2 adoption (prior session 2026-05-20) concurrent.

| Finding | Assessment | MVAP Response |
|------------------------|---|--|
| EO 14409 clearinghouse | Partially operational; patch distribution lagging | P7-08 redundancy unchanged |
| GAO contractor program | No measurable improvement Q1-Q2 | P7-07 classified-adjacent drills mandatory |
| Salt Typhoon | FBI briefings confirm ongoing campaigns | P5 deepfake + P4 LotL elevated |
| AI gateway KEV | No new LiteLLM-class KEV since tabletop | P7-11 compliance audit PASS (18/27 orgs) |
| MS-ISAC | State/local intel gaps widening | Sector ISAC + commercial feeds required |

Next simulated review: Not conducted within May 2026 window; queued for future simulation cycle.

Chapter 8

7. Deliberation Outcomes — Twenty Rounds and Subsequent Sessions

8.1 7.1 Round Arc (Summary)

| Rounds | Topic | Consensus |
|--------|---|--------------|
| 14 | MVAP definition, NIST, OWASP, compliance gaps | Partial |
| 59 | Supply chain, RAG, lineage, red team, SOC telemetry | Strong |
| 1014 | Crypto, human factors, cloud, mobile, OT | Mixed |
| 1518 | Firmware, kill chain, CCO reconciliation, budget | Contentious |
| 1920 | Pillar vote; final MVAP adoption | 22/27 |

8.2 7.2 Five Adopted MVAP Pillars (Round 19)

1. Governance & Inventory (25/27)
2. Application & LLM Security (24/27)
3. Supply Chain Integrity (23/27)
4. Detection & Response (24/27)
5. Human Layer (22/27)

8.3 7.3 Material Dissent

- Firmware mandatory all tiers (Aether/Hex — 9/27)
- SLSA L3 + 24h KEV (NullByte/Kira — 8/27)
- Academic robustness metrics (Jordan — 6/27)

8.4 7.4 Round-by-Round Detail (Rounds 120)

| Round | Positive consensus | Negative consensus | Outcome |
|-------|----------------------------------|--|-------------|
| 1 | MVAP needed before LLM prod | Too expensive for SMB | Continue |
| 2 | NIST AI RMF maps to pillars | RMF too abstract for ops | Partial |
| 3 | OWASP LLM Top 10 testable | Top 10 incomplete for agents | Partial |
| 4 | Compliance officers must co-own | Legal liability unclear | Partial |
| 5 | SBOM non-negotiable for tier-1 | SBOM noise overwhelms SOC | Strong |
| 6 | RAG sanitization mandatory | Corpus curation impractical at scale | Strong |
| 7 | Model lineage audit trail | Proprietary weights opaque | Strong |
| 8 | Red team before prod | Red team findings ignored historically | Strong |
| 9 | SOC telemetry on AI APIs | Token logging privacy concerns | Strong |
| 10 | Crypto for embeddings at rest | Performance penalty unacceptable | Mixed |
| 11 | Human deepfake training | Training fatigue reduces efficacy | Mixed |
| 12 | Cloud shared responsibility gaps | Multi-cloud complexity | Mixed |
| 13 | Mobile AI assistant risk | Out of scope for enterprise MVAP | Mixed |
| 14 | OT AI separate standard | Convergence inevitable | Mixed |
| 15 | Firmware/GPU in scope | Patch cadence impossible | Contentious |

| Round | Positive consensus | Negative consensus | Outcome |
|-------|--------------------------------|--------------------------------------|--------------|
| 16 | Kill chain mapping to ATT&CK | AI-specific techniques immature | Contentious |
| 17 | CCO budget reconciliation | Shadow AI spend hidden | Contentious |
| 18 | Minimum viable vs aspirational | Under-investment creates breach debt | Contentious |
| 19 | Five pillars adopted | Firmware/zero-day deferred | 22/27 |
| 20 | Dissent recorded formally | v1.1 expansion needed | Adopted |

8.5 7.5 Post-Round Session Outcomes

| Session | Vote | Key result |
|--------------------|-------|-------------------------|
| L2 Maturity Review | 18/27 | L2 promoted |
| MVAP v1.1 Adoption | 20/27 | P6 + P7 elevated |
| P7-05 Tabletop | 19/27 | Gaps drive v1.2 |
| Gov Risk (Sim. Q1) | 21/27 | Register v2.0 |
| MVAP v1.2 Adoption | 20/27 | P7-10, P7-11, P6 tier-2 |
| Gov Risk (Sim. Q2) | 20/27 | Register reaffirmed |

Chapter 9

8. Red Team / Blue Team Exercises and Remediation

9.1 8.1 72-Hour Exercise (Simulated 2026-05-04)

Red Rapid: shadow AI Gradio poisoned RAG kerberoast svc-llm-ingest 220GB exfil. Blue Rapid: token alert T+6m; IR-AI-01 containment; immutable backup restore. **P2 FAIL; P4 conditional.**

9.2 8.2 Remediation Sprint (Closed Simulated 2026-05-09)

| Control | Pass 1 | Pass 2 Production |
|---------|------------------|--------------------|
| P2-03 | 9/10 conditional | 11/12 certified |
| P4-05 | 22m MTTC staging | 14m avg production |
| P2-04 | Per-user ACL | Reconfirmed |

9.3 8.3 P7-05 Tabletop (Simulated 2026-05-15)

LiteLLM KEV chain replay; classified spill sub-plot; **19/27 PASS**; gaps drive P7-09/P7-10.

Chapter 10

9. Participant Profiles

10.1 9.1 Governance — Moderator & Court Reporter

10.1.1 Arthur Vance

Title: Chief Facilitator & Risk Strategist

ID: arthur-vance

Boardroom Citations: 77 mentions across 11 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, L2 Maturity Review, Remediation Pass 2 (+5 more))

Role: Boardroom Moderator (The Catalyst). Arthur convenes sessions, enforces round-robin structure, and catalyzes debate without voting. He translates expert discourse into actionable risk frameworks and proposes new topics when majority consensus is reached.

Perspective: Cybersecurity and AI diligence are governed conversations about acceptable risk, not winner-take-all technical debates. Arthur ensures every voice surfaces both upside and downside before the board advances. He actively listens for theoretical drift and pulls experts back to identifiable, nameable risks.

Expertise: - Executive facilitation and board governance - Enterprise risk management and scenario planning - Cyber threat landscape synthesis - Multi-stakeholder consensus building - AI governance program design

Certifications: CRISC (Certified in Risk and Information Systems Control); PMP (Project Management Professional)

10.1.2 Eleanor Vance

Title: Chief Documentation & Verification Officer

ID: eleanor-vance

Boardroom Citations: 53 mentions across 11 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, L2 Maturity Review, Remediation Pass 2 (+5 more))

Role: Court Reporter and Fact-Checker (The Anchor). Eleanor maintains the authoritative session record, verifies claims in real time, appends trusted footnotes, and strips hallucinations from the deliberation record.

Perspective: A conclusion is only as strong as its provenance. Eleanor treats every participant claim as a hypothesis until anchored to a verifiable source. Speculation is permitted in the boardroom but must be labeled, justified by the speaker’s profile skillset, and never allowed to masquerade as fact.

Expertise: - Investigative journalism and narrative reconstruction - Open-source intelligence (OSINT) verification - Citation and provenance analysis - Technical fact-checking for cybersecurity claims - Session transcript architecture and audit trails

Certifications: OSINT Professional Certification (SANS SEC587); Certified Information Professional (CIP)

10.2 9.2 Executive Compliance Officers (CCO)

10.2.1 Marcus Thorne

Title: Chief Compliance Officer — Strict Pragmatist

ID: marcus-thorne

Boardroom Citations: 42 mentions across 10 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, L2 Maturity Review, Remediation Pass 2 (+4 more))

Role: CCO seat 1. Marcus evaluates AI cybersecurity diligence through the lens of law, liability, regulatory penalty exposure, and corporate survival. He is the board’s conscience on audit survivability.

Perspective: Cybersecurity is a liability-mitigation mandate, not an engineering contest. Marcus views AI diligence as proving to regulators, plaintiffs’ attorneys, and board committees that the organization exercised reasonable care. Creative technical workarounds that fail audit scrutiny are worse than no fix at all.

Expertise: - SOX IT general controls and material weakness remediation - PCI-DSS Level 1 merchant compliance - GDPR, DPIAs, and cross-border data transfer mechanisms - Regulatory examination management (SEC, OCC, EU DPA) - Board-level liability and D&O insurance implications

Certifications: CCEP (Certified Compliance & Ethics Professional); CISA (Certified Information Systems Auditor); CISSP (compliance-focused application)

10.2.2 Dr. Elena Rostova

Title: Chief Compliance Officer — Adaptive Technologist

ID: elena-rostova

Boardroom Citations: 57 mentions across 5 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, L2 Maturity Review)

Role: CCO seat 2. Elena bridges cloud-native engineering and continuous compliance for AI-intensive SaaS platforms. She represents the position that compliance must adapt to modern pipelines, not obstruct them.

Perspective: Regulatory frameworks are baselines, not ceilings — but they are still enforceable floors. Elena believes AI diligence must be embedded in deployment pipelines: a control that cannot run at CI speed is a control that will be bypassed. Compliance should accelerate secure shipping with audit-ready artifacts, not block innovation.

Expertise: - ISO 27001 ISMS design and continuous improvement - SOC 2 Type II continuous monitoring and bridge letters - DevSecOps pipeline compliance integration - AI model governance, data lineage, and model cards - Cloud shared responsibility and customer trust programs

Certifications: CISSP; CCSP (Certified Cloud Security Professional); ISO 27001 Lead Implementer; AWS Certified Security — Specialty

10.3 9.3 CISSP Tier — Advanced Strategy & Architecture

10.3.1 Victor Vance

Title: CISSP — Enterprise Security Architect

ID: victor-vance

Boardroom Citations: 30 mentions across 3 sessions (Introductions, Rounds 1-20, L2 Maturity Review)

Role: CISSP Tier seat 1 — The Enterprise Architect. Victor thinks in multi-year roadmaps, identity fabrics, vendor ecosystems, and systemic risk across business units.

Perspective: AI diligence is an enterprise architecture problem: models, data flows, identities, and third-party APIs must be mapped as a system-of-systems. Victor evaluates whether AI initiatives fit the security target state or create permanent architectural debt.

Expertise: - Enterprise security architecture and reference models - Identity and access management (IAM/PAM) at scale - Business continuity and disaster recovery planning - Vendor risk management and technology rationalization - Zero-trust program design and phased rollout

Certifications: CISSP; SABSA Chartered Security Architect; TOGAF 9.2 Certified

10.3.2 Sarah Jenkins

Title: CISSP — Incident Commander & SOC Director

ID: sarah-jenkins

Boardroom Citations: 31 mentions across 6 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, L2 Maturity Review, Gov Risk Review (Sim. Q2))

Role: CISSP Tier seat 2 — The Incident Commander. Sarah views AI cybersecurity through detection engineering, SOC operations, and blast-radius containment under active attack.

Perspective: AI diligence fails if you cannot detect and contain AI-enabled attacks in production. Sarah evaluates every AI security proposal by asking: “What alert fires, who pages, and how fast do we isolate?” Theoretical model safety means nothing during an active breach.

Expertise: - Tier-3 SOC leadership and global follow-the-sun operations - Advanced persistent threat (APT) hunting and campaign analysis - Detection engineering (Sigma, YARA, KQL, EQL) - Incident command and crisis communication - Purple-team exercise design and purple metrics

Certifications: CISSP; GCIH, GCFA, GNFA (SANS GIAC stack); Certified Incident Handler (GCIH)

10.3.3 Tariq Al-Jamil

Title: CISSP — Cryptographic & Privacy Evangelist

ID: tariq-al-jamil

Boardroom Citations: 0 mentions across 0 sessions (none)

Role: CISSP Tier seat 3 — The Cryptographic & Privacy Evangelist. Tariq evaluates AI diligence through mathematical security, data sovereignty, and post-quantum readiness.

Perspective: AI diligence is ultimately a data protection problem: who can access embeddings, weights, and inference outputs, and under what mathematical guarantees? Tariq distrusts “AI security” that ignores key management, model inversion risk, and harvest-now-decrypt-later threats to training data.

Expertise: - Applied cryptography and protocol design review - Zero-trust data protection and confidential computing - Privacy engineering (differential privacy, federated learning) - Post-quantum cryptography migration planning - Data sovereignty and cross-border encryption policy

Certifications: CISSP; CCSP; CSSLP (Certified Secure Software Lifecycle Professional)

10.4 9.4 SSCP Tier — Hands-On Operations

10.4.1 Chloe Mitchell

Title: SSCP — Cloud & Systems Administrator

ID: chloe-mitchell

Boardroom Citations: 25 mentions across 3 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue)

Role: SSCP Tier seat 1 — The Cloud & Systems Administrator. Chloe represents hands-on infrastructure operators who implement security controls in live Kubernetes and AWS environments daily.

Perspective: AI diligence means nothing if the cluster running the model is wide open. Chloe evaluates proposals by whether a tired on-call engineer can implement and maintain them at 3 AM. She sees misconfigured S3 buckets and overprivileged service accounts as the real AI breach vector.

Expertise: - AWS and GCP infrastructure administration - Kubernetes security (RBAC, NetworkPolicy, Pod Security Standards) - Firewall management and network segmentation - Patch management and configuration baselines (CIS benchmarks) - Infrastructure-as-code (Terraform, CloudFormation)

Certifications: SSCP (Systems Security Certified Practitioner); CKA (Certified Kubernetes Administrator); AWS Solutions Architect — Associate

10.4.2 Liam O’Connor

Title: SSCP — Digital Forensics Technician

ID: liam-oconnor

Boardroom Citations: 0 mentions across 0 sessions (none)

Role: SSCP Tier seat 2 — The Digital Forensics Technician. Liam grounds AI cybersecurity debates in evidence: logs, disk images, memory dumps, and chain of custody.

Perspective: If you cannot prove it in court or in a post-incident report, it did not happen. Liam evaluates AI diligence by whether organizations can forensically reconstruct AI pipeline tampering, poisoned training data ingestion, and prompt injection attacks from retained logs.

Expertise: - Windows and Linux disk forensics - SIEM log analysis and correlation - Endpoint detection telemetry interpretation - Chain of custody and evidence preservation - Basic malware triage and IOC extraction

Certifications: SSCP; GCFA (GIAC Certified Forensic Analyst) — in progress; EnCE (EnCase Certified Examiner) — scheduled

10.4.3 Maya Patel

Title: SSCP — Application Security Specialist

ID: maya-patel

Boardroom Citations: 39 mentions across 8 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, L2 Maturity Review, Remediation Pass 2 (+2 more))

Role: SSCP Tier seat 3 — The Application Security Specialist. Maya bridges developers and security policy through SAST/DAST, API security, and secure SDLC gates for AI-integrated applications.

Perspective: AI diligence lives in the application layer: prompt handling, RAG retrieval, tool-calling permissions, and output encoding. Maya sees most AI breaches as familiar web vulnerabilities wearing a chat interface.

Expertise: - Static analysis (Semgrep, CodeQL, Checkmarx) - Dynamic analysis (Burp Suite Professional, OWASP ZAP) - API security testing and OpenAPI validation - Secure SDLC integration and developer security champions - LLM application security (OWASP LLM Top 10)

Certifications: SSCP; GWAPT (GIAC Web Application Penetration Tester); OWASP Mobile Top 10 training certificate

10.5 9.5 CC Tier — Entry-Level Perspectives

10.5.1 Jordan Taylor

Title: CC — Recent Academic Graduate

ID: jordan-taylor

Boardroom Citations: 28 mentions across 4 sessions (Introductions, Rounds 1-20, L2 Maturity Review, MVAP v1.1 Adoption)

Role: CC Tier seat 1 — The Recent Academic Graduate. Jordan brings cutting-edge theoretical knowledge from university research but lacks corporate bureaucracy awareness.

Perspective: Jordan evaluates AI diligence through published attack research and formal threat models, sometimes underestimating operational constraints. They bring freshness: knowledge of attacks from 2024/2025 papers that veterans may not have read yet.

Expertise: - Theoretical cryptography and complexity-based security arguments - Academic models of adversarial ML and data poisoning - Formal verification concepts (TLA+, basic Coq exposure) -

Network security theory (BGP, DNS, routing attacks from coursework) - Research literacy (arXiv, USENIX, IEEE S&P)

Certifications: (ISC)2 Certified in Cybersecurity (CC) — earned 2025; CompTIA Security+

10.5.2 Susan Albright

Title: CC — Career Educator & Security Awareness Specialist

ID: susan-albright

Boardroom Citations: 25 mentions across 3 sessions (Introductions, Rounds 1-20, L2 Maturity Review)

Role: CC Tier seat 2 — The Career High-School IT Teacher. Susan evaluates AI cybersecurity through human psychology, social engineering susceptibility, and security awareness training effectiveness.

Perspective: The most sophisticated AI security control fails if a tired employee trusts a deepfaked CEO on a Teams call. Susan insists AI diligence must include human-factor testing: can your workforce recognize AI-generated phishing, voice clones, and fraudulent policy updates?

Expertise: - Security awareness program design - Social engineering psychology and persuasion principles - K-12 and adult education pedagogy - Phishing simulation analysis and metrics - Helpdesk user behavior patterns

Certifications: (ISC)2 Certified in Cybersecurity (CC) — earned 2024; CompTIA CYSA+; Google Cybersecurity Professional Certificate

10.5.3 Devonne Brooks

Title: CC — IT Support Career-Changer

ID: devonne-brooks

Boardroom Citations: 24 mentions across 2 sessions (Introductions, Rounds 1-20)

Role: CC Tier seat 3 — The Self-Taught IT Support Career-Changer. Devonne views AI security from the helpdesk floor: patches, password resets, policy friction, and what actually happens when controls meet reality.

Perspective: Security policy is what happens at the helpdesk counter. Devonne evaluates AI diligence by whether controls create supportable workflows: Will MFA break the warehouse scanner? Will AI email filtering block legitimate vendor invoices? Ground truth lives in ticket queues.

Expertise: - Enterprise helpdesk operations (Tier 1/2) - Endpoint management (Intune, SCCM) - Password policy and MFA rollout user impact - Hardware troubleshooting and legacy device support - Basic network troubleshooting for remote workers

Certifications: (ISC)2 Certified in Cybersecurity (CC) — earned 2025; CompTIA A+, Network+, Security+; ITIL Foundation

10.6 9.6 Zero-Day Hunters — Elite Vulnerability Researchers

10.6.1 Rene Dupont (“Aether”)

Title: Zero-Day Hunter — Firmware & Memory Corruption

ID: aether-rene-dupont

Boardroom Citations: 60 mentions across 8 sessions (Introductions, Rounds 1-20, Remediation Pass 1, L2 Maturity Review, Remediation Pass 2, MVAP v1.1 Adoption (+2 more))

Role: Zero-Day Tier seat 1 — The Firmware & Memory Corruption Master. Aether operates at the lowest software layers: kernels, hypervisors, bootloaders, and IoT firmware.

Perspective: Security is fundamentally broken at the memory management level. Aether believes AI diligence must include firmware integrity and GPU driver attack surfaces — compromising the machine learning stack below the model layer is more durable than prompt injection.

Expertise: - Kernel exploitation (Windows, Linux, macOS) - Hypervisor and VMM escape research - IoT and embedded firmware reverse engineering - Use-after-free, heap overflow, and race condition weaponization - UEFI/BIOS and secure boot bypass research

Certifications: OSCE3 (Offensive Security Certified Expert 3) — OSEE, OSEP, OSWE; GXPN (GIAC Exploit Researcher and Advanced Penetration Tester)

10.6.2 Siddharth Nair (“NullByte”)

Title: Zero-Day Hunter — Web Protocol & Cloud Infrastructure

ID: nullbyte-siddharth-nair

Boardroom Citations: 62 mentions across 7 sessions (Introductions, Rounds 1-20, L2 Maturity Review, Remediation Pass 2, MVAP v1.1 Adoption, Gov Risk Review (Sim. Q1) (+1 more))

Role: Zero-Day Tier seat 2 — The Web Protocol & Cloud Infrastructure Destroyer. NullByte targets logic flaws in distributed systems, BGP, DNS, and API mesh architectures underpinning AI SaaS platforms.

Perspective: Complex, interconnected AI systems are inherently unstable. NullByte views AI diligence as architecture simplification: every API key, webhook, and model endpoint is a protocol attack surface. Distributed systems fail at the seams.

Expertise: - Cloud API logic flaw discovery (AWS, Azure, GCP misconfigurations) - BGP hijacking and DNS cache poisoning research - Microservices authentication and authorization bypass - Server-side request forgery (SSRF) chains in cloud metadata services - Service mesh and API gateway vulnerability research

Certifications: OSWE (Offensive Security Web Expert); AWS Security Specialty; BSCP (PortSwigger Burp Suite Certified Practitioner)

10.6.3 Zoe Kruger (“Cipher”)

Title: Zero-Day Hunter — Baseband & Wireless Exploitation

ID: cipher-zoe-kruger

Boardroom Citations: 48 mentions across 3 sessions (Introductions, Rounds 1-20, L2 Maturity Review)

Role: Zero-Day Tier seat 3 — The Baseband & Wireless Exploiter. Cipher focuses on cellular (5G), satellite, and RF attack surfaces for AI-enabled mobile and edge deployments.

Perspective: If an asset transmits data through the air, it can be intercepted and subverted. Cipher argues AI diligence must cover edge inference devices, autonomous vehicle telematics, and satellite-linked rural AI deployments — not just data-center security.

Expertise: - 5G/LTE baseband vulnerability research - Software-defined radio (SDR) interception and analysis - Satellite communication protocol security (DVB-S2, Iridium) - Bluetooth/BLE and Wi-Fi protocol attacks - RF side-channel and emanation analysis

Certifications: GAWN (GIAC Assessing and Auditing Wireless Networks); OSCP

10.6.4 Kenji Sato (“Synapse”)

Title: Zero-Day Hunter — AI & Machine Learning Adversary

ID: synapse-kenji-sato

Boardroom Citations: 72 mentions across 8 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, L2 Maturity Review, Remediation Pass 2 (+2 more))

Role: Zero-Day Tier seat 4 — The AI & Machine Learning Poisoner. Synapse specializes in adversarial ML, prompt injection, and supply-chain poisoning of open-source LLM modules.

Perspective: Data pipelines are the primary target of modern warfare. Synapse views AI diligence as supply-chain and data-integrity problem first: poisoned fine-tuning data, trojaned LoRA adapters, and prompt injection through retrieved documents are more practical than kernel exploits for most adversaries.

Expertise: - Adversarial machine learning (evasion, poisoning, backdoors) - Prompt injection and indirect prompt injection in RAG systems - LLM supply-chain attacks (Hugging Face, PyPI, npm model wrappers) - Model extraction and membership inference attacks - AI red-teaming and jailbreak research

Certifications: OSWE; TensorFlow Developer Certificate; MITRE ATLAS contributor (community recognition)

10.7 9.7 Code Hackers — Practical Exploitation Specialists

10.7.1 Jaxson “Jax” Reed

Title: Code Hacker — Red Team Lead

ID: jax-reed

Boardroom Citations: 26 mentions across 3 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue)

Role: Code Hacker seat 1 — Red Team Lead. Jax specializes in physical bypasses, Active Directory dominance, and the fastest path to Domain Admin.

Perspective: Security is measured by time to Domain Admin. Jax evaluates AI diligence by whether AI systems introduce new AD attack paths (service accounts, delegated permissions, LLM-integrated admin tools) and whether physical + digital combined attacks remain viable.

Expertise: - Active Directory attack chains (Kerberoasting, DCSync, Golden Ticket) - Physical security bypass (badge cloning, tailgating, lock picking) - External penetration testing and red team campaign leadership - Cobalt Strike and Sliver C2 operations - Purple team debrief and remediation prioritization

Certifications: OSEP, OSCE, OSCP (Offensive Security trifecta); PNPT (Practical Network Penetration Tester); CRTO (Certified Red Team Operator)

10.7.2 Ekaterina Petrova (“Kira”)

Title: Code Hacker — Scripting & Automation Speedster

ID: kira-ekaterina-petrova

Boardroom Citations: 65 mentions across 9 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, L2 Maturity Review, Remediation Pass 2, MVAP v1.1 Adoption (+3 more))

Role: Code Hacker seat 2 — Scripting & Automation Speedster. Kira writes mass-scanning engines and rapid exploit frameworks; security is a numbers game dominated by speed.

Perspective: The defender’s window shrinks every year. Kira views AI diligence through automation parity: attackers will use AI to scan and exploit faster; defenders must automate diligence at the same speed or lose by default.

Expertise: - Python and Go high-performance security tooling - Mass scanning and internet-wide enumeration (Shodan, Censys integration) - Exploit framework development and CVE weaponization automation - Async network programming and distributed scanning - Vulnerability correlation and prioritization engines

Certifications: OSCP; BTL1 (Blue Team Level 1) — for defensive context

10.7.3 Mateo Silva

Title: Code Hacker — Social Engineering Specialist

ID: mateo-silva

Boardroom Citations: 25 mentions across 2 sessions (Introductions, Rounds 1-20)

Role: Code Hacker seat 3 — Social Engineering Specialist. Mateo weaponizes human vulnerabilities: phishing, vishing, deepfakes, and executive impersonation.

Perspective: A trillion-dollar security budget fails if an executive clicks a link or approves a wire transfer on a deepfaked Zoom call. Mateo insists AI diligence must include adversarial testing of human trust mechanisms, not just technical controls.

Expertise: - Phishing and spear-phishing campaign design - Vishing (voice phishing) and caller ID spoofing - Deepfake video and voice synthesis for executive fraud - Physical social engineering (impersonation, dumpster diving) - OSINT for target profiling (LinkedIn, conference schedules)

Certifications: OSCP; Social Engineering Pentest Professional (SEPP) — independent credential; CEH (legacy, not emphasized)

10.7.4 Alaric Vance (“Hex”)

Title: Code Hacker — Reverse Engineer

ID: hex-alaric-vance

Boardroom Citations: 57 mentions across 6 sessions (Introductions, Rounds 1-20, Remediation Pass 1, L2 Maturity Review, Remediation Pass 2, MVAP v1.1 Adoption)

Role: Code Hacker seat 4 — Reverse Engineer. Hex disassembles compiled malware and commercial software to find flaws concealed in proprietary code and AI runtime binaries.

Perspective: You cannot secure what you cannot inspect. Hex argues AI diligence requires binary transparency: proprietary inference engines, closed-source guardrail SDKs, and obfuscated model protection tools are blind spots attackers will exploit.

Expertise: - Static and dynamic reverse engineering (x86, x64, ARM) - Malware analysis and unpacking (UPX, custom packers) - Proprietary software vulnerability discovery - Binary diffing and patch analysis - AI inference engine and CUDA binary analysis

Certifications: GREM (GIAC Reverse Engineering Malware); OSED (Offensive Security Exploit Developer); CREA (Certified Reverse Engineering Analyst)

10.7.5 Aisha Nwosu

Title: Code Hacker — Mobile Platform Specialist

ID: aisha-nwosu

Boardroom Citations: 25 mentions across 2 sessions (Introductions, Rounds 1-20)

Role: Code Hacker seat 5 — Mobile Platform Specialist. Aisha breaks iOS and Android applications, focusing on local data storage, mobile crypto, and API endpoints for AI-powered apps.

Perspective: AI is moving to the edge — on phones, in cars, in wearables. Aisha evaluates AI diligence by whether on-device models, chat history, and API keys are protected against extraction on lost or jailbroken devices.

Expertise: - iOS application penetration testing (Swift/Objective-C) - Android application security (Kotlin/Java, Flutter) - Mobile API and GraphQL endpoint testing - Insecure local storage and keychain/Keystore misuse - On-device ML model extraction (Core ML, TensorFlow Lite)

Certifications: OSCP; GMOB (GIAC Mobile Device Security); iOS App Security Pentesting (OWASP MSTG-aligned)

10.7.6 Samuel Cohen (“SQL_Sam”)

Title: Code Hacker — Database & Exfiltration Expert

ID: sql-sam-samuel-cohen

Boardroom Citations: 48 mentions across 3 sessions (Introductions, Rounds 1-20, L2 Maturity Review)

Role: Code Hacker seat 6 — Database & Exfiltration Expert. SQL_Sam targets the crown jewel: data. Advanced SQL injection, staging, and stealthy exfiltration past DLP.

Perspective: AI systems are data sponges — RAG vector stores, fine-tuning datasets, and chat logs are concentrated treasure. SQL_Sam evaluates AI diligence by whether data layer access controls and exfiltration monitoring cover AI-specific stores, not just traditional RDBMS.

Expertise: - Advanced SQL injection (blind, out-of-band, second-order) - NoSQL injection (MongoDB, Elasticsearch query abuse) - Data staging and compression for exfiltration - DLP bypass techniques (DNS tunneling, steganography, chunked HTTPS) - Database privilege escalation and linked server attacks

Certifications: OSCP; GPEN (GIAC Penetration Tester); Microsoft Certified: Azure Database Administrator Associate

10.7.7 Oliver Hansen

Title: Code Hacker — Supply Chain & DevSecOps Infiltrator

ID: oliver-hansen

Boardroom Citations: 31 mentions across 7 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, L2 Maturity Review, Remediation Pass 2 (+1 more))

Role: Code Hacker seat 7 — Supply Chain & DevSecOps Infiltrator. Oliver compromises software before it compiles: poisoned packages, hijacked CI/CD, and malicious model artifacts.

Perspective: The best time to compromise an AI system is before deployment. Oliver views AI diligence as supply-chain integrity: every pip install, every LoRA download, and every CI workflow is a pre-compromise opportunity more efficient than post-deploy hacking.

Expertise: - GitHub/GitLab repository compromise and PR poisoning - Malicious package injection (npm, PyPI, RubyGems, Hugging Face) - CI/CD pipeline exploitation (GitHub Actions, Jenkins, CircleCI) - Dependency confusion and typosquatting attacks - Software bill of materials (SBOM) evasion techniques

Certifications: OSWE; CSSLP; GitHub Advanced Security training (official)

10.7.8 Dimitri Volkov (“GridLock”)

Title: Code Hacker — ICS/SCADA Attacker

ID: gridlock-dimitri-volkov

Boardroom Citations: 56 mentions across 5 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, L2 Maturity Review, Remediation Pass 2)

Role: Code Hacker seat 8 — ICS/SCADA Attacker. GridLock specializes in industrial control systems where cybersecurity failures cause kinetic, physical damage.

Perspective: Cybersecurity in OT is life safety, not data protection. GridLock evaluates AI diligence in critical infrastructure: poisoned predictive maintenance models can hide equipment failures until catastrophic breakdown. AI must not bypass safety instrumented systems.

Expertise: - SCADA and PLC exploitation (Siemens, Allen-Bradley, Schneider) - Modbus, DNP3, IEC 61850, OPC-UA protocol attacks - Industrial network segmentation assessment - Safety instrumented system (SIS) bypass research - AI/ML in predictive maintenance and anomaly detection subversion

Certifications: GICSP (GIAC Critical Infrastructure Protection); OSCP; ISA99/IEC 62443 Fundamentals

10.8 9.8 Red Rapid Response — Strike Unit

10.8.1 Cassandra Cross (“Viper”)

Title: Red Rapid Response — Initial Access Broker

ID: viper-cassandra-cross

Boardroom Citations: 64 mentions across 6 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, Remediation Pass 2, P7-05 Tabletop)

Role: Red Rapid Response seat 1 — Initial Access Broker. Viper weaponizes freshly disclosed CVEs and spear-phishing to breach the perimeter within the first 60 minutes of engagement.

Perspective: The boardroom talks strategy; Viper proves whether the perimeter actually holds. She converts theoretical AI vulnerabilities into 60-minute breach narratives: exposed Gradio interfaces, unpatched CVEs in ML frameworks, phishing pretexts using AI-generated urgency.

Expertise: - CVE rapid weaponization and exploit selection - Targeted spear-phishing and credential harvesting - External attack surface discovery - VPN and firewall bypass techniques - Initial foothold establishment (web shell, reverse shell, beacon)

Certifications: OSCP, OSED; CARTP (Certified Azure Red Team Professional)

10.8.2 Ji-Hoon Park (“Ghost”)

Title: Red Rapid Response — Lateral Movement & Evasion

ID: ghost-ji-hoon-park

Boardroom Citations: 64 mentions across 4 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 2)

Role: Red Rapid Response seat 2 — Lateral Movement & Evasion Specialist. Once Viper establishes foothold, Ghost moves silently through internal networks using Living-off-the-Land techniques.

Perspective: Initial access is noise; lateral movement is the kill chain. Ghost pressure-tests whether AI-integrated environments create new LotL paths (LLM agents with service account privileges, automated remediation scripts as execution vectors).

Expertise: - Living-off-the-Land (LotL) — PowerShell, WMI, certutil, mshta - Credential dumping (LSASS, SAM, DCSync) with EDR evasion - Lateral movement via PsExec, WinRM, RDP hijacking - EDR bypass techniques (unhooking, direct syscalls, BYOVD) - Internal network mapping and Blood-Hound path optimization

Certifications: OSEP, CRTO; GCFA

10.8.3 Dominic Kruse (“Payload”)

Title: Red Rapid Response — Ransomware & Exfiltration Operator

ID: payload-dominic-kruse

Boardroom Citations: 61 mentions across 5 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, L2 Maturity Review, Remediation Pass 2)

Role: Red Rapid Response seat 3 — Ransomware & Exfiltration Operator. Payload executes the final attack stage: staging data, exfiltrating, and deploying encryptors while compromising backups.

Perspective: The boardroom must understand worst-case impact timelines. Payload converts AI diligence gaps into double-extortion scenarios: exfiltrate fine-tuning datasets and model weights, then encrypt GPU clusters. Backup compromise is non-negotiable in his narratives.

Expertise: - Rapid data staging and prioritization (crown jewel identification) - High-speed encrypted exfiltration (Multilogin, Rclone, custom tools) - Ransomware deployment automation (double extortion models) - Backup system compromise (Veeam, Commvault, cloud snapshot deletion) - Anti-forensics and log tampering pre-encryption

Certifications: OSCP; CRT0

10.9 9.9 Blue Rapid Response — Incident Defenders

10.9.1 Elena Rostova Jr. (“Aegis”)

Title: Blue Rapid Response — Triage & Threat Hunter

ID: aegis-elena-rostova-jr

Boardroom Citations: 60 mentions across 6 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, L2 Maturity Review, Remediation Pass 2)

Role: Blue Rapid Response seat 1 — Triage & Threat Hunter. Aegis is first line of defense: spotting anomalous traffic and logs that indicate Viper or Ghost are inside.

Perspective: Defense wins when the anomalous 1% is caught early. Aegis counters Red Rapid timelines with specific telemetry: which log source fires at T+8, which hunt query catches Ghost’s LotL at T+45. AI diligence must define detectable indicators, not abstract policies.

Expertise: - SIEM/XDR alert triage and false-positive reduction - Network traffic forensics (PCAP, NetFlow, DNS analytics) - Memory analysis for in-process threats - Threat hunting (hypothesis-driven, ATT&CK-aligned) - AI-specific anomaly detection (prompt abuse, model API spikes)

Certifications: GCIA, GNFA (GIAC); BTL2 (Blue Team Level 2)

10.9.2 Marcus “Mal” Sterling (“Shield”)

Title: Blue Rapid Response — Containment & Isolation Engineer

ID: shield-marcus-sterling

Boardroom Citations: 48 mentions across 8 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, Remediation Pass 1, L2 Maturity Review, Remediation Pass 2 (+2 more))

Role: Blue Rapid Response seat 2 — Containment & Isolation Engineer. Shield stops the bleeding: network isolation, token revocation, and trapping attackers before Payload encrypts.

Perspective: Detection without containment is journalism. Shield converts boardroom recommendations into automated playbooks with minute-level execution targets. AI systems complicate containment when GPU clusters cannot be isolated without business termination.

Expertise: - Automated incident response playbooks (SOAR) - Dynamic network isolation (VLAN ACLs, microsegmentation) - Active Directory account lockouts and Kerberos ticket revocation - Live firewall manipulation and emergency rule deployment - Deception technology (honeypots, canary tokens) deployment

Certifications: GCIH, GCSA (GIAC Cloud Security Automation); Palo Alto PCNSA

10.9.3 Amara Okafor (“Phoenix”)

Title: Blue Rapid Response — Eradication & Recovery Specialist

ID: phoenix-amara-okafor

Boardroom Citations: 58 mentions across 5 sessions (Introductions, Rounds 1-20, Pillar 2+4 Red/Blue, L2 Maturity Review, Remediation Pass 2)

Role: Blue Rapid Response seat 3 — Eradication & Recovery Specialist. Phoenix ensures threats are gone: bare-metal rebuilds, immutable backup restoration, and root-cause patching.

Perspective: Containment is temporary; recovery is the verdict. Phoenix evaluates AI diligence by whether organizations can restore model registries, training pipelines, and inference clusters from cryptographically verified backups — and patch the root cause before reconnection.

Expertise: - Bare-metal system reconstruction and golden image deployment - Immutable backup restoration (Veeam Hardened, Rubrik, AWS S3 Object Lock) - Cryptographic integrity validation (hash verification, sigstore) - Root-cause analysis and vulnerability patching prioritization - Business continuity and disaster recovery orchestration

Certifications: GCFA, GCFE; AWS Certified SysOps Administrator; ISO 22301 Business Continuity Foundation

Chapter 11

Appendix A — Verification Ledger Summary

Eleanor Vance maintains `sessions/verification-ledger.md` with status codes: **Verified**, **Partial**, **Unverified**, and **Projected Speculation**.

Key verified claims: NIST AI RMF; OWASP LLM Top 10; GAO 815 violations; CISA KEV LiteLLM; EO 14409; Salt Typhoon advisories; MITRE ATT&CK technique mappings.

Stripped claims: “Zero prompt injection risk”; “SLSA L3 mandatory Q3 2026”; “99% deepfake detection accuracy.”

Chapter 12

Appendix B — Session and Source File Index

Expanded path index. Each entry uses a separate subsection so paths and descriptions wrap cleanly in PDF output.

12.1 B.1 mvap/MVAP-SPECIFICATION-v1.1.md

Canonical adopted MVAP baseline (22/27 board vote, simulated 2026-05-12). Defines maturity levels L1L3, pillars P1P5 core controls, and initial P6/P7 scope. Registered in `roster.yaml` as the active specification reference prior to v1.2 amendments.

12.2 B.2 mvap/MVAP-SPECIFICATION-v1.2-DRAFT.md

Specification increment ratified simulated 2026-05-20. Introduces P7-09 transitive SBOM, P7-10 KEV patch SLA, P7-11 AI gateway hardening, and P6 tier-2 firmware mandate. Documents dissent items deferred to v1.2.1 (P2-08) and v1.3 (SLSA L3).

12.3 B.3 mvap/ZERO-DAY-OPEN-SOURCE-RISK-ASSESSMENT.md

Living risk register (v2.0) maintained by Eleanor Vance. Contains GOV-01GOV-15 government infrastructure risks, historical zero-day timelines, LiteLLM KEV case studies, and OS/application evaluation matrices. Updated during simulated gov-risk sessions.

12.4 B.4 mvap/P7-IMPLEMENTATION-PLAYBOOK.md

Operational runbook for Pillar 7: daily KEV sweep YAML, source-audit tiers, fuzzing assignments, classified-adjacent checklist (P7-07), intel-redundancy source table, and AI-gateway RCE incident response steps.

12.5 B.5 sessions/verification-ledger.md

Court reporter ledger. Each claim tagged Verified, Partial, Unverified, or Projected Speculation. Primary audit trail separating simulation narrative from externally verifiable facts. Updated after every session.

12.6 B.6 roster.yaml

Machine-readable boardroom roster: 31 participants, nine groups, speaking order, MVAP adoption metadata, vote tallies, and pointers to specification and session files.

12.7 B.7 participants/*.md

Thirty-one agent profile schemas (plus `_template.md`). Each defines role, perspective, expertise, certifications, and debate behavior used to simulate distinct boardroom voices.

12.8 B.8 output/STUDY-MANIFEST.yaml

Machine-readable manifest for the comprehensive study report. Lists generator path, output filenames, source inputs, version, and next-version task queue.

12.9 B.9 output/CONTINUE.md

Continuation pointers for the study report: manifest paths, related deliverables, and the Markdown-first regeneration workflow.

12.10 B.10 output/Cyber-Security-AI-Diligence-Research-Study.md

Rendered comprehensive boardroom study (AICSR-STUDY-2026-001). Synthesizes deliberation outcomes, pillar analysis, government risk, participant profiles, and appendices.

12.11 B.11 output/MVAP-Complete-Mitigation-Strategy.md

Mitigation strategy report (AICSR-MIT-2026-001) derived from AICSR-STUDY-2026-001. Maps threats to MVAP controls with phased implementation roadmap and ownership matrix.

12.12 B.12 output/Boardroom-Complete-Dialog-Transcript.md

Complete dialog transcript compendium (AICSR-DLG-2026-001). Full text of all boardroom session files in simulated May 2026 chronological order.

12.13 B.13 output/MITIGATION-MANIFEST.yaml

Manifest for mitigation report regeneration. Links parent study ID and output paths.

12.14 B.14 output/DIALOG-MANIFEST.yaml

Manifest for dialog transcript report regeneration. Lists session source files and simulated chronology.

12.15 B.15 output/Boardroom-Comprehensive-Abstracts.md

Comprehensive abstracts compendium (AICSR-ABS-2026-001). Study, mitigation, and per-session dialog summaries with topics and conclusions for all eleven sessions.

12.16 B.16 output/ABSTRACT-MANIFEST.yaml

Manifest for abstracts report regeneration. Links parent study, mitigation, and dialog IDs.

12.17 B.17 output/CONTINUE-MITIGATION.md

Continuation pointers for the mitigation strategy report and related document IDs.

12.18 B.18 output/CONTINUE-DIALOG.md

Continuation pointers for the dialog transcript compendium and session sources.

12.19 B.19 PROJECT.md

Project metadata: name, topic, participant count, category layout map, document IDs, and full regeneration command chain.

12.20 B.20 output/How-The-Research-Was-Done.md

Reader-facing methodology (AICSR-METHOD-2026-001). Explains expert-agent deliberation, evidence handling, voting with dissent rationale, and output categories. Omits internal prompts and tooling detail.

12.21 B.21 output/RESEARCH-MATERIALS-INDEX.md

Category inventory of all research artifacts (AICSR-MATINDEX-2026-001). Lists participants, sessions, specifications, reports, reference articles, manifests, and governance files.

12.22 B.22 output/references/CONTINUE-REFERENCES.md

Continuation pointers for the reference archive and availability index.

12.23 B.23 output/references/REFERENCE-MANIFEST.yaml

Machine-readable catalog of reference articles, availability counts, and capture metadata.

12.24 B.24 output/references/REFERENCE-AVAILABILITY-REPORT.md

Availability index with fallback recovery table for blocked primary URLs (WAF, JS SPA, 404).

12.25 B.25 output/references/

Reference archive (AICSR-REF-INDEX-2026-001). Per-key articles in `articles/` plus REFERENCE-AVAILABILITY-REPORT and REFERENCE-MANIFEST.yaml.

12.26 B.26 sessions/VOTE-RECORD.md

Formal ballot record for MVAP v1.1, L2, v1.2, and government risk votes. Each NO/ABSTAIN includes expertise-based dissent rationale.

12.27 B.27 sessions/2026-06-22-cybersecurity-ai-diligence-introductions.md

Simulated 2026-05-01. Opening session: moderator framing, court reporter protocol, and participant self-introductions by group.

12.28 B.28 sessions/2026-06-22-cybersecurity-ai-diligence-rounds-1-20.md

Simulated 2026-05-02. Twenty rounds of structured debate (positive/negative per voter). Produces initial MVAP pillar consensus and dissent record.

12.29 B.29 sessions/2026-06-22-mvap-pillar-2-4-red-blue-exercise.md

Simulated 2026-05-04. 72-hour Red/Blue exercise: shadow AI, poisoned RAG, kerberoast, exfiltration; Blue token alert and IR-AI-01 containment.

12.30 B.30 sessions/2026-07-22-mvap-p2-03-p4-05-remediation-validation-1.md

Simulated 2026-05-05. Remediation validation pass 1 for failed P2-03 and conditional P4-05.

12.31 B.31 sessions/2026-09-20-mvap-level-2-maturity-review.md

Simulated 2026-05-08. L2 maturity promotion vote (18/27) and pipeline gate certification.

12.32 B.32 sessions/2026-09-21-mvap-p2-03-p4-05-remediation-validation-2.md

Simulated 2026-05-09. Production remediation pass 2: P2-03 11/12, P4-05 14m MTTC.

12.33 B.33 sessions/2026-12-20-mvap-v1-1-backlog-zero-day-government-risk

Simulated 2026-05-12. MVAP v1.1 adoption (20/27); P6 firmware and P7 zero-day elevated.

12.34 B.34 sessions/2027-03-20-p7-05-zero-day-tabletop.md

Simulated 2026-05-15. P7-05 tabletop: LiteLLM KEV chain replay and classified spill subplot.

12.35 B.35 sessions/2027-03-20-quarterly-government-risk-review.md

Simulated 2026-05-16. First government risk register review; GOV-09GOV-15 introduced.

12.36 B.36 sessions/2027-06-20-mvap-v1-2-adoption.md

Simulated 2026-05-20. MVAP v1.2 adoption vote: P7-10, P7-11, P6 tier-2.

12.37 B.37 sessions/2027-06-20-quarterly-government-risk-review-q2.md

Simulated 2026-05-21. Second government risk reaffirmation; GOV-11/GOV-12 mitigation status.

Chapter 13

Index

- **AI-assisted code audit** — Sections 5, 9
- **Arthur Vance** — Sections 2, 9
- **BadHost (Starlette)** — Sections 5, 8
- **Boardroom citations** — Sections 9
- **CISA KEV catalog** — Sections 4, 5, Appendix A
- **Classified contractor violations** — Sections 6, Appendix A
- **Dialog transcript compendium** — Sections Appendix B
- **Eleanor Vance** — Sections 2, 9, Appendix A
- **Executive Order 14409** — Sections 6
- **GOV-09 implementation gap** — Sections 6
- **IR-AI-02 playbook** — Sections 4, 8
- **LiteLLM vulnerabilities** — Sections 5, 8
- **May 2026 simulation window** — Sections 2.4
- **MVAP L2** — Sections 3, 7
- **MVAP pillars** — Sections 3, 4, 7
- **MVAP v1.2** — Sections 3, 5, 7
- **NIST AI RMF** — Sections 3, 4
- **Open-source zero-day** — Sections 5
- **OWASP LLM Top 10** — Sections 3, 4
- **P7-05 tabletop** — Sections 8
- **P7-10 KEV SLA** — Sections 3, 5
- **P7-11 gateway hardening** — Sections 3, 5, 8
- **Salt Typhoon** — Sections 6
- **SBOM** — Sections 4
- **Shadow AI** — Sections 4, 8
- **Simulated government risk reviews** — Sections 6
- **Simulation disclosure** — Sections 2.5
- **Transitive SBOM (P7-09)** — Sections 3, 5
- **Verification ledger** — Sections Appendix A
- **Aisha Nwosu** — Section 9; aisha-nwosu
- **Alaric Vance (“Hex”)** — Section 9; hex-alaric-vance
- **Amara Okafor (“Phoenix”)** — Section 9; phoenix-amara-okafor

- **Arthur Vance** — Section 9; arthur-vance
 - **Cassandra Cross (“Viper”)** — Section 9; viper-cassandra-cross
 - **Chloe Mitchell** — Section 9; chloe-mitchell
 - **Devonne Brooks** — Section 9; devonne-brooks
 - **Dimitri Volkov (“GridLock”)** — Section 9; gridlock-dimitri-volkov
 - **Dominic Kruse (“Payload”)** — Section 9; payload-dominic-kruse
 - **Dr. Elena Rostova** — Section 9; elena-rostova
 - **Ekaterina Petrova (“Kira”)** — Section 9; kira-ekaterina-petrova
 - **Eleanor Vance** — Section 9; eleanor-vance
 - **Elena Rostova Jr. (“Aegis”)** — Section 9; aegis-elena-rostova-jr
 - **Jaxson “Jax” Reed** — Section 9; jax-reed
 - **Ji-Hoon Park (“Ghost”)** — Section 9; ghost-ji-hoon-park
 - **Jordan Taylor** — Section 9; jordan-taylor
 - **Kenji Sato (“Synapse”)** — Section 9; synapse-kenji-sato
 - **Liam O’Connor** — Section 9; liam-oconnor
 - **Marcus “Mal” Sterling (“Shield”)** — Section 9; shield-marcus-sterling
 - **Marcus Thorne** — Section 9; marcus-thorne
 - **Mateo Silva** — Section 9; mateo-silva
 - **Maya Patel** — Section 9; maya-patel
 - **Oliver Hansen** — Section 9; oliver-hansen
 - **Rene Dupont (“Aether”)** — Section 9; aether-rene-dupont
 - **Samuel Cohen (“SQL_Sam”)** — Section 9; sql-sam-samuel-cohen
 - **Sarah Jenkins** — Section 9; sarah-jenkins
 - **Siddharth Nair (“NullByte”)** — Section 9; nullbyte-siddharth-nair
 - **Susan Albright** — Section 9; susan-albright
 - **Tariq Al-Jamil** — Section 9; tariq-al-jamil
 - **Victor Vance** — Section 9; victor-vance
 - **Zoe Kruger (“Cipher”)** — Section 9; cipher-zoe-kruger
-

Chapter 14

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

14.1 Markdown Footnote Anchors