

The MYTHOS Playbook, Volume II

Joseph P. Conroy

2026

Contents

The MYTHOS Playbook, Volume II	1
The Governance Model	1
Copyright	2
Two-Volume Set — Multi-Format Edition	2
A Note on the EPUB Edition	2
Legal and Editorial Disclosures	2
Publisher	3
Preface	4
Where Volume I Ended	4
What Volume II Does	4
Continuity with Volume I	5
What This Volume Is Not	5
Acknowledgments	6
Institutional Acknowledgments	6
VectorCertain Engineering	6
Family	6
Part IV — Frameworks	7
Chapter 20: Mapping MYTHOS to OWASP, MITRE ATLAS, and NIST	8
Chapter 20 Purpose	8
20.1 Reading the Cross-Walk Matrices	8
20.2 OWASP LLM Top 10 v2025 MYTHOS T1–T7	10
Table 20.2.1 — OWASP LLM Top 10 v2025 × MYTHOS T1–T7	10
Worked example: an OWASP LLM finding becomes a MYTHOS evidence package	13
20.3 OWASP Top 10 for Agentic Applications (Dec 2025) MYTHOS	14
Table 20.3.1 — OWASP Top 10 for Agentic Applications (Dec 2025) × MYTHOS T1–T7	14
Worked example: an A1 / A7 evidence-of-control response	18
20.4 MITRE ATLAS v5.1.0 MYTHOS	19
Table 20.4.1 — MITRE ATLAS v5.1.0 (selected techniques) × MYTHOS T1–T7	19
Worked example: an ATLAS-driven adversarial-ML threat-modeling exercise	22
20.5 MITRE ATT&CK Enterprise MYTHOS (for SOC Integration)	22
Table 20.5.1 — MITRE ATT&CK Enterprise (selected techniques) × MYTHOS T1–T7	23
Worked example: integrating MYTHOS coverage into the SOC’s existing ATT&CK-based detection content	25
20.6 Using the Matrices in Tabletop Exercises and Vendor RFPs	26
Tabletop exercises	26
Vendor RFPs	27
Where the matrices stop	27
20.7 CISO’s Five Next Steps	28
Cross-References	28
References	29
Chapter 21: NIST AI RMF Implementation for CISOs	30

21.1 Framework Introduction	30
21.1.1 Framework history and provenance	30
21.1.2 Scope and audience	31
21.1.3 Regulatory and contractual force	31
21.1.4 Why this framework matters for an agentic-AI deployment	32
21.2 Framework Architecture	32
21.2.1 The four functions: Govern, Map, Measure, Manage	32
21.2.2 The 72 subcategories and the GenAI Profile	33
21.2.3 Versioning and revision pacing	33
21.3 MYTHOS Coverage Mapping	34
21.3.1 Cross-walk methodology	34
21.3.2 The cross-walk matrix	34
21.3.3 By-MYTHOS-architecture-surface cross-walk	38
21.3.4 Sprint 4 vector inbound forward-ref resolution	38
21.4 Per-Subcategory Narrative	39
21.4.1 GOVERN 1.1 — AI risk management policies	39
21.4.2 GOVERN 1.4 / GOVERN 2.1 — Responsibility and role allocation	40
21.4.3 GOVERN 5.1 / GOVERN 6.1 — Value-chain engagement and third-party AI risk	40
21.4.4 MAP 1.1 / MAP 3.1 — Intended purpose and AI capabilities documentation	41
21.4.5 MAP 5.1 / MAP 5.2 — Operational-environment risk identification and quantification	41
21.4.6 MEASURE 1.1 / MEASURE 2.1 — Measurement methods and validity	41
21.4.7 MEASURE 2.6 / MEASURE 2.7 — Safety, security, and resilience	42
21.4.8 MEASURE 2.10 / MEASURE 3.1 / MEASURE 4.1 — Privacy, impact assessment, and emergent risk	42
21.4.9 MANAGE 1.1 / MANAGE 2.1 — Risk treatment options and strategic alignment	43
21.4.10 MANAGE 3.1 / MANAGE 4.1 / MANAGE 4.2 — Risk response, continuous monitor- ing, pre / post deployment	43
21.5 SecureAgent Evidence Surface	43
21.5.1 Artifact catalog — by NIST AI RMF function	43
21.5.2 Block-time and reproducibility evidence	44
21.5.3 Certification implications	45
21.6 MYTHOS Conformance Block — NIST AI RMF (NIST AI 100-1) + Generative AI Profile (NIST AI 600-1)	45
21.6.A Coverage — controls fully met	45
21.6.B Evidence — artifacts and logs available	46
21.6.C Gaps — controls not fully met	46
21.6.D Forward-Fits — Sprint 4+ architectural extensions to close gaps	47
21.6.E Certification Path — NIST AI RMF certification implications	47
21.6.F The AI Security Maturity Model (Tier 1 → Tier 3)	47
21.7 CISO’s Five Next Steps	48
Without SecureAgent: What You Can Do Today	48
21.8 Summary	49
References	49
Cross-References	49

Chapter 22: CRI FS AI RMF: Operationalizing the 230 Control Objectives **51**

Chapter Purpose Statement	51
§22.1 Framework Introduction	51
§22.1.1 Framework history and provenance	51
§22.1.2 Scope and audience	52
§22.1.3 Regulatory and contractual force	52
§22.1.4 Why this framework matters now — the “1999 problem”	53
§22.2 Framework Architecture	54
§22.2.1 The framework’s organizing principle — Govern / Map / Measure / Manage	54
§22.2.2 The 230 control objectives — three-level DS notation	54
§22.2.3 Companion documents and cross-references	55
§22.2.4 Versioning and the next revision	56

§22.3 MYTHOS Coverage Mapping	56
§22.3.1 Cross-walk methodology	56
§22.3.2 The cross-walk matrix — load-bearing control objectives	57
§22.3.3 Vector-by-vector cross-walk — resolving the seven Sprint 4 inbound forward-refs	60
§22.3.4 Load-bearing-control flagging	63
§22.4 Per-Control Narrative — The Twelve Load-Bearing Objectives	63
§22.4.1 GV.PO-01.01 — AI risk management policy is established	63
§22.4.2 GV.OV-01.01 — AI system inventory is maintained	64
§22.4.3 MP.CN-01.01 — AI system context and intended use are documented	64
§22.4.4 MP.RM-01.01 — AI system threat model is maintained	65
§22.4.5 PR.AC-01.01 — Identities and credentials are issued, managed, verified, revoked, and audited	65
§22.4.6 PR.AC-04.01 — Access permissions and authorizations are managed at the agentic boundary	66
§22.4.7 PR.PT-01.01 — Audit/log records are determined, documented, implemented, and reviewed	66
§22.4.8 PR.PT-04.01 — Communication and control networks are protected at the agentic egress boundary	67
§22.4.9 MS.TS-01.01 — AI system testing is conducted before deployment	67
§22.4.10 MS.TS-03.01 — AI system testing produces immutable evidence	68
§22.4.11 DE.AE-02.01 — Anomalies and events are analyzed for AI-specific signatures	68
§22.4.12 DE.CM-04.02 — AI-system-specific malicious-context detection is implemented	69
§22.5 SecureAgent Evidence Surface	69
§22.5.1 Artifact catalog	69
§22.5.2 Per-artifact evidence mapping	71
§22.5.3 Evidence-export discipline	71
§22.5.4 Block-time and reproducibility evidence	71
§22.5.5 Examination implications — no formal certification pathway	72
§22.6 MYTHOS Conformance Block — CRI FS AI RMF Profile v1.2.1	72
§22.6.A Coverage — controls fully met	72
§22.6.B Evidence — artifacts and logs available	73
§22.6.C Gaps — controls not fully met	73
§22.6.D Forward-Fits — Sprint 4+ architectural extensions	74
§22.6.E Certification Path — the SecureAgent conformance-suite case	75
§22.7 CISO's Five Next Steps	75
Without SecureAgent: What You Can Do Today	76
References	76
Cross-References	76
Chapter 23: ISO/IEC 42001 Implementation	78
Chapter Purpose	78
23.1 Framework Introduction	78
23.1.1 Framework history and provenance	78
23.1.2 Scope and audience	79
23.1.3 Regulatory or contractual force	79
23.1.4 Why this framework matters for an agentic-AI deployment	80
23.2 Framework Architecture	80
23.2.1 The framework's organizing principle — Clauses 4–10 plus Annex A	80
23.2.2 The control-class enumeration — Annex A	81
23.2.3 Cross-references to companion documents	82
23.2.4 Versioning and revision pacing	83
23.3 MYTHOS Coverage Mapping	83
23.3.1 Cross-walk methodology	83
23.3.2 The cross-walk matrix	84
23.3.3 Per-MYTHOS-architecture-surface cross-walk	88
23.3.4 Load-bearing controls flagged for §23.4	88
23.4 Per-Control Narrative	88

23.4.1 A.4.2 — AI system policy	89
23.4.2 A.5.2 — AI risk management	89
23.4.3 A.6.2.4 — AI system verification and validation	90
23.4.4 A.6.2.6 — AI system event logging	91
23.4.5 A.6.2.7 — AI system decision trace	91
23.4.6 A.6.2.8 — AI system data quality	92
23.4.7 A.7.3 — AI system performance monitoring	92
23.4.8 A.7.5 — AI system human oversight	92
23.4.9 A.10.2 — Allocation of responsibilities — third party	93
23.5 SecureAgent Evidence Surface	93
23.5.1 Artifact catalog	93
23.5.2 Per-artifact evidence mapping	94
23.5.3 Evidence-export discipline	95
23.5.4 Block-time and reproducibility evidence	95
23.5.5 Certification implications	95
23.6 MYTHOS Conformance Block — ISO/IEC 42001:2023	96
23.6.A Coverage — controls fully met	96
23.6.B Evidence — artifacts and logs available	97
23.6.C Gaps — controls not fully met	97
23.6.D Forward-Fits — Sprint 4+ architectural extensions to close gaps	97
23.6.E Certification Path — registered Certification Bodies and accreditation chain	98
23.7 CISO’s Five Next Steps	99
Without SecureAgent — What You Can Do Today	99
Cross-References	100
References	100
Chapter 24: OWASP LLM Top 10 v2025 + Agentic Top 10: Defensive Playbook	101
Chapter 24 Purpose	101
24.1 Framework Introduction	102
24.1.1 OWASP Foundation history and provenance	102
24.1.2 Scope and audience	102
24.1.3 Regulatory and contractual force	102
24.1.4 Why these frameworks matter for an agentic-AI deployment	103
24.2 Framework Architecture	103
24.2.1 The OWASP LLM Top 10 v2025 — verbatim category enumeration	103
24.2.2 The OWASP Agentic Top 10 (December 2025) — verbatim category enumeration	104
24.2.3 The OWASP FinBot CTF — third-party self-assessment exercise	105
24.2.4 How the two registers interact	105
24.2.5 Versioning and revision pacing	105
24.3 MYTHOS Coverage Mapping	105
24.3.1 Cross-walk methodology	105
24.3.2 The cross-walk matrix — OWASP categories to MYTHOS recipes	105
24.3.3 Vector-by-vector cross-walk — OWASP Agentic Top 10 to MYTHOS T1 through T7	109
24.3.4 Load-bearing categories for §24.4 deep walk	110
24.4 Per-Category Narrative — Twelve Load-Bearing OWASP Categories	110
24.4.1 LLM01 Prompt Injection	110
24.4.2 LLM02 Sensitive Information Disclosure	111
24.4.3 LLM03 Supply Chain	111
24.4.4 LLM05 Improper Output Handling	111
24.4.5 LLM06 Excessive Agency	112
24.4.6 LLM07 System Prompt Leakage	112
24.4.7 LLM08 Vector and Embedding Weaknesses	112
24.4.8 LLM09 Misinformation	113
24.4.9 A1 Memory Poisoning	113
24.4.10 A2 Tool Misuse	113
24.4.11 A3 Privilege Compromise	114
24.4.12 A6 Intent Breaking and Goal Manipulation	114

24.5 SecureAgent Evidence Surface	115
24.5.1 Artifact catalog — what SecureAgent produces per OWASP category	115
24.5.2 Per-artifact evidence mapping	115
24.5.3 Evidence-export discipline	116
24.5.4 Block-time and reproducibility evidence	116
24.5.5 The OWASP FinBot CTF — third-party self-assessment exercise	116
24.5.6 Crumpton/MITRE positioning at OWASP-coverage evidence	116
24.6 MYTHOS Conformance Block — OWASP LLM Top 10 v2025 + Agentic Top 10 (December 2025)	117
§24.6.A Coverage — controls fully met	117
§24.6.B Evidence — artifacts and logs available	118
§24.6.C Gaps — controls not fully met	118
§24.6.D Forward-Fits — Sprint 4+ architectural extensions	119
§24.6.E Certification Path — framework-specific certification implications	120
24.7 CISO’s Five Next Steps	120
Without SecureAgent: What You Can Do Today	120
References	121
Cross-References	121

Chapter 25: NVIDIA AI Red Team Controls and Integration 122

Chapter Purpose Statement	122
§25.1 Framework Introduction	122
25.1.1 Framework history and provenance	122
25.1.2 Scope and audience	122
25.1.3 Regulatory or contractual force	123
25.1.4 Why this framework matters — the GRC × ML development cross-product	123
§25.2 Framework Architecture	124
25.2.1 The framework’s organizing principle — mandatory trio plus recommended six	124
25.2.2 The mandatory trio	124
25.2.3 The recommended six	125
25.2.4 Cross-references to companion documents	127
25.2.5 Versioning and revision pacing	127
§25.3 MYTHOS Coverage Mapping	128
25.3.1 Cross-walk methodology	128
25.3.2 The cross-walk matrix — by NVIDIA control	129
25.3.3 Vector-by-vector cross-walk	130
25.3.4 Load-bearing controls — the controls examiners and auditors spend the most time on	132
§25.4 Per-Control Narrative	133
25.4.1 M1 — Network egress lockdown	133
25.4.2 M2 — Workspace-only writes	134
25.4.3 M3 — Configuration protection	134
25.4.4 R1 — Full virtualization	135
25.4.5 R2 — Secret injection	135
25.4.6 R3 — Lifecycle management	135
25.4.7 R4 — Read restrictions	136
25.4.8 R5 — Approval architectures	136
25.4.9 R6 — IDE sandboxing	136
§25.5 SecureAgent Evidence Surface	137
25.5.1 Artifact catalog	137
25.5.2 Per-artifact evidence mapping	137
25.5.3 Evidence-export discipline	138
25.5.4 Block-time and reproducibility evidence	138
25.5.5 Certification implications — and the GRC × ML development validation pattern	138
§25.6 MYTHOS Conformance Block — NVIDIA AI Red Team Framework	139
§25.6.A Coverage — controls fully met	139
§25.6.B Evidence — artifacts/logs available	140
§25.6.C Gaps — controls not fully met	140
§25.6.D Forward-Fits — Sprint 4+ architectural extensions to close gaps	141

§25.6.E Certification Path — framework-specific certification implications	141
§25.7 CISO's Five Next Steps	142
Without SecureAgent — What You Can Do Today	142
References	143
Cross-References	143
Part V — SOC	144
Chapter 26: SIEM Integration for Agentic AI	145
§26.1 Operational Context	145
§26.2 Substantive Architecture	146
§26.2.1 The GTID Telemetry Schema	147
§26.2.2 AGL-SG ComplianceAuditEvent Stream	148
§26.2.3 Pipeline-Stage Telemetry Coverage	148
§26.2.4 Ingestion Patterns: Push vs Pull, Multi-Tenant Considerations	149
§26.2.5 Cardinality Control: The Load-Bearing Discipline	149
§26.2.6 Retention Discipline: Three Tiers Aligned to Framework Evidence	150
§26.3 MYTHOS / SecureAgent Mapping — Inbound Forward-Reference Resolution	150
§26.3.1 Inbound Forward-Reference Resolution Table — Sprint 4 (18 forward-refs)	150
§26.3.2 Inbound Forward-Reference Resolution Table — Sprint 5 (27 forward-refs)	153
§26.3.3 Per-Vector Detection-Capability Mapping (T1–T7)	156
§26.3.4 Per-Pipeline-Stage Telemetry Coverage Mapping	160
§26.3.5 Multi-Framework Cross-Walk	161
§26.4 Per-SIEM-Stack Narrative	161
§26.4.1 Splunk Enterprise Security	161
§26.4.2 Microsoft Sentinel	165
§26.4.3 Google Chronicle	168
§26.4.4 Elastic Security	171
§26.4.5 Optional Fifth Stack	173
§26.5 SecureAgent Production Surface — Cross-Correlation Analytics, Dashboards, Schema Documen- tation	173
§26.5.1 The Cross-Correlation Analytic Library	173
§26.5.2 Per-Platform Reference Dashboards	174
§26.5.3 Telemetry Schema Documentation Reference	175
§26.5.4 Forward-References to Sprint 7+ Chapters	175
§26.6 Conformance Block (Strict 5-Subsection)	175
§26.6.A Capability — what this chapter enables operationally	175
§26.6.B Evidence — artifacts, logs, dashboards, and retention attestations	176
§26.6.C Gaps — controls not fully met	177
§26.6.D Forward-Fits — Sprint 7+ extensions	177
§26.6.E Operational Path — SIEM ingestion deployment runbook	178
§26.7 CISO's Five Next Steps	178
Footnotes	179
Cross-References	179
Chapter 27: The AI-SOC: A Capability Overlay, Not a New Department	181
Chapter 27 Purpose	181
27.1 Operational Context	181
27.1.1 The reader's question and the IBM framing answer	181
27.1.2 Where you are in the deployment arc	182
27.1.3 Audience	182
27.1.4 Prior-chapter dependencies	183
27.2 Substantive Architecture — The Capability Overlay Model	183
27.2.1 What the capability overlay is and is not	183
27.2.2 The three new role types	184
27.2.3 Existing-tier amendments	184
27.2.4 Skill prerequisites and learning curves	185

27.3 MYTHOS / SecureAgent Mapping — Per-Role Organization	185
27.3.1 Per-role rather than per-vector — the divergence framing	185
27.3.2 Inbound forward-reference resolution	186
27.3.3 AI-SOC analyst per-role mapping	186
27.3.4 AI-SOC architect per-role mapping	187
27.3.5 AI-SOC IR specialist per-role mapping	187
27.4 Per-Role Narrative	188
27.4.1 The AI-SOC analyst — daily activities, escalation triggers, runbook access	188
27.4.2 The AI-SOC architect — design responsibilities, vendor-evaluation tie-in	188
27.4.3 The AI-SOC IR specialist — incident-response runbooks, forensic integration	189
27.4.4 Capacity curves — staffing rule of thumb and the K1 cross-cut	190
27.4.5 Hiring, training, ramp-up timelines	191
27.5 SecureAgent Production Surface	191
27.5.1 Three reference job descriptions	191
27.5.2 RACI matrix — AI-SOC × existing-tier intersection	192
27.5.3 Runbook library framework	195
27.5.4 Tabletop exercise integration	196
27.5.5 Production-surface forward references	196
27.6 MYTHOS Conformance Block — AI-SOC Capability Overlay	196
§27.6.A Capability — AI-SOC operational capability via overlay	196
§27.6.B Evidence — org chart amendments, role descriptions, runbook libraries, training records, capacity reports	197
§27.6.C Gaps — role-coverage gaps, capacity ceilings, skill-pipeline assumptions	197
§27.6.D Forward-Fits — Sprint 7+ extensions	198
§27.6.E Operational Path — staffing rollout sequence	198
27.7 CISO’s Five Next Steps	199
Without SecureAgent: What You Can Do Today	199
References	200

Chapter 28: Adversary Emulation for Agentic Environments	201
§28.1 Operational Context	201
§28.2 Substantive Architecture — MITRE ATT&CK Adversary Emulation, Agentic Extensions, and Cyber-Range Design	202
§28.2.1 MITRE ATT&CK adversary emulation as canonical inheritance	202
§28.2.2 Agentic-specific extensions to the MITRE paradigm	204
§28.2.3 Agentic cyber-range design	204
§28.3 MYTHOS / SecureAgent Mapping — Per-Vector Emulation, Per-Pipeline-Stage Shadow-Mode Coverage	205
§28.3.1 Per-vector emulation scenario mapping (T1–T7)	206
§28.3.2 Per-pipeline-stage shadow-mode operational coverage	207
§28.3.3 SecureAgent shadow-mode operationalization during exercises	208
§28.4 Per-Vector Scenario Narrative — Seven Canonical Scenarios at Operational Depth	210
§28.4.1 The Last Ones 32-step exploit chain (T1 Autonomous Multi-Step Exploitation)	210
§28.4.2 Cooling Tower OT scenario (T6 Sandbox Escape Exploitation)	211
§28.4.3 Clinejection — clinical-context prompt-injection scenario (T2 Unsanctioned Scope Creep)	213
§28.4.4 Evaluation-Aware Drift scenario (T3 Invisible / Deceptive Reasoning)	214
§28.4.5 Audit-Replay Tampering scenario (T4 Track-Covering Log Manipulation)	215
§28.4.6 Non-Human Identity Hijack scenario (T5 Credential Theft)	215
§28.4.7 Tool-Registry Inflation scenario (T7 Capability Proliferation)	216
§28.4.8 Purple-team integration with HOTS — continuous-evaluation overlay	217
§28.5 SecureAgent Production Surface — Scenario Library, Cyber-Range Reference, Purple-Team Template	217
§28.5.1 Scenario library structure	217
§28.5.2 Cyber-range architecture reference	218
§28.5.3 Purple-team session template	219
§28.5.4 Per-vector emulation-result reporting surface	219
§28.6 MYTHOS Conformance Block — Adversary Emulation for Agentic Environments	220

§28.6.A Capability — adversary-emulation against agentic threats	220
§28.6.B Evidence — scenario library inventories, exercise results, purple-team reports, cyber-range configurations	221
§28.6.C Gaps — scenario-library coverage gaps and purple-team frequency limits	222
§28.6.D Forward-Fits — Sprint 7+ extensions	222
§28.6.E Operational Path — red-team adoption sequence	223
§28.6.F Composition with NVIDIA Red-Team Validation Pattern (chapter-specific extension)	224
§28.7 The CISO's Five Next Steps	224
References	225
Cross-References	225

Chapter 29: Vendor Evaluation: A Statistical Field Test **227**

Chapter Purpose Statement	227
§29.1 Operational Context	227
29.1.1 The vendor-claim translation problem	227
29.1.2 Where you are in the deployment arc	228
29.1.3 Audience	228
29.1.4 The vendor-claim skeptic's checklist	229
29.1.5 What the chapter operationalizes	229
§29.2 Substantive Architecture	230
29.2.1 What the statistical field-test protocol is	230
29.2.2 The five protocol stages	230
29.2.3 Clopper-Pearson 3-sigma scenario count tables (Stage 1 canonical reference)	231
29.2.4 Confusion-matrix disclosure requirements	231
29.2.5 Per-vector detection-rate non-blending discipline	232
§29.3 MYTHOS / SecureAgent Mapping	233
29.3.1 The mapping surface	233
29.3.2 Per-vector evaluation methodology mapping (T1–T7)	234
29.3.3 Sub-technique distribution (registry citation)	235
29.3.4 Per-claim-type evaluation methodology mapping	236
29.3.5 Pipeline-stage mapping	236
29.3.6 SecureAgent applies the methodology to itself — the worked example	237
§29.4 Per-Evaluation-Stage Narrative	239
29.4.1 Stage 1 — Scenario-count derivation (planning the evaluation)	239
29.4.2 Stage 2 — Scenario execution (running the evaluation)	239
29.4.3 Stage 3 — Confusion-matrix construction (the disclosure standard)	240
29.4.4 Stage 4 — Per-vector report generation (the procurement deliverable)	241
29.4.5 Stage 5 — Ongoing-validation cadence (post-procurement detection of drift)	242
§29.5 SecureAgent Production Surface	243
29.5.1 What this section ships	243
29.5.2 RFP language library — 12 canonical clauses	243
29.5.3 Evaluation-report template	245
29.5.4 Vendor-claim audit checklist (operational deployment of §29.1.4)	246
29.5.5 Forward-references to Sprint 7+ chapters	247
§29.6 Conformance Block	247
§29.6.A Capability — vendor-evaluation via statistical field-test methodology	247
§29.6.B Evidence — what the institution produces and retains	248
§29.6.C Gaps — what the methodology does not cover	249
§29.6.D Forward-Fits — Sprint 7+ extensions	249
§29.6.E Operational Path — procurement adoption sequence	250
§29.6.F RFP Language Library Index — topic cross-reference	251
§29.7 CISO's Five Next Steps	251
Without SecureAgent — What You Can Do Today	252
References	253
Cross-References	253

Chapter 30: Deployment Topologies and Capacity Planning	255
§30.1 Operational Context	255
§30.2 Substantive Architecture	257
§30.2.1 The Three-Topology Taxonomy	257
§30.2.2 The Latency-Budget Grammar	258
§30.2.3 The Capacity-Sizing Model	259
§30.2.4 The Legacy-Hardware Deployment Narrative	259
§30.2.5 The HA/DR Posture	260
§30.2.6 The Deployment-Cost Reduction Thesis (Pointer to §30.4)	260
§30.3 MYTHOS / SecureAgent Mapping — Inbound Forward-Reference Resolution	261
§30.3.1 Inbound Forward-Reference Resolution Table — Sprint 1 + Sprint 4 (21 forward-refs)	261
§30.3.2 Inbound Forward-Reference Resolution Table — Sprint 5 (28 forward-refs)	263
§30.3.3 Inbound Forward-Reference Resolution Table — Sprint 6 (21 forward-refs)	266
§30.3.4 Per-Vector Deployment-Coverage Mapping (T1–T7)	271
§30.3.5 Per-Pipeline-Stage Deployment Coverage Mapping (AMRS / Four-Gate / AGL-SG)	273
§30.4 Per-Topology Narrative	274
§30.4.1 The Edge Deployment Narrative	274
§30.4.2 The Gateway Deployment Narrative	275
§30.4.3 The Orchestrator-Integrated Deployment Narrative	276
§30.4.4 The Legacy-Hardware Deployment Narrative	278
§30.4.5 The HA/DR Deployment Narrative	279
§30.4.6 The Very-Large-Scale Staffing-Posture Narrative	280
§30.5 SecureAgent Production Surface — Deployment Templates, Cost Disclosure, and Cross-Vendor Coverage	280
§30.5.1 Deployment Topology Decision Matrix	280
§30.5.2 Deployment Specification and Deployment Plan Templates	280
§30.5.3 AI System Card and RACI Matrix Templates	281
§30.5.4 Deployment-Cost Disclosure Template	281
§30.5.5 Cross-Vendor Surface Coverage Disclosure	282
§30.5.6 Hub-and-Spoke Architecture Artifact (Very-Large-Scale)	282
§30.5.7 Forward-References to Later Chapters	282
§30.6 Conformance Block (Strict 5-Subsection)	283
§30.6.A Capability — what this chapter enables operationally	283
§30.6.B Evidence — artifacts, logs, and attestations	283
§30.6.C Gaps — known limitations	284
§30.6.D Forward-Fits — Sprint 8+ extensions	284
§30.6.E Operational Path — deployment-rollout sequence	285
§30.6.F Per-Topology Reference Annex (chapter-specific extension)	285
§30.7 CISO's Five Next Steps	286
Footnotes	286
Cross-References	286
 Chapter 31: Governance for Non-Human Identities	 288
Chapter 31 Purpose	288
31.1 Operational Context	289
31.2 Substantive Architecture — The CSA NHI Framework and the Four Operational Axes	291
31.3 MYTHOS / SecureAgent Mapping	293
§31.3.1 Inbound Forward-Reference Resolution	293
§31.3.2 Per-Vector NHI Mapping (T1–T7)	296
§31.3.3 Per-Pipeline-Stage NHI Mapping	297
§31.3.4 GTID / ZGTID Architectural Distinction at NHI Audit-Chain Integration	298
31.4 Per-Topic Operational Narrative	298
§31.4.1 NHI Inventory and Classification at Execution Depth	298
§31.4.2 Runtime-Token Governance	299
§31.4.3 Agent Passports and Ed25519 Cryptographic Identity	300
§31.4.4 Behavioral-Baseline Maintenance — The AMRS V4 Deployment Surface	301
§31.4.5 Cross-Tenant Isolation — The ZGTID Deployment Surface	302

§31.4.6 Lifecycle Management at Full Per-Stage Depth	303
31.5 SecureAgent Production Surface	304
§31.5.1 NHI Governance Reference Architecture	304
§31.5.2 Per-NHI Exposure Card Library	304
§31.5.3 Credential-Lifecycle Runbook Library	305
§31.5.4 NHI Behavioral-Baseline Maintenance Reference	305
§31.5.5 Cross-Tenant Evidence-Package Template	305
§31.5.6 Production-Surface Forward References	306
31.6 MYTHOS Conformance Block — NHI Governance	306
§31.6.A Capability — Comprehensive NHI Governance	306
§31.6.B Evidence — NHI Inventory Reports, Credential Audit Logs, Behavioral-Baseline At- testations	306
§31.6.C Gaps — Uncovered NHI Classes, Isolation Limits, Baseline Drift Detection Limits . . .	307
§31.6.D Forward-Fits — Sprint 8+ Extensions	308
§31.6.E Operational Path — NHI Governance Rollout Sequence	308
31.7 CISO's Five Next Steps	308
References	309

Chapter 32: Working With the Board and the Examiner 310

Chapter Purpose Statement	310
§32.1 Operational Context	310
§32.1.1 The four-conversation problem	310
§32.1.2 Where you are in the deployment arc	311
§32.1.3 Audience	312
§32.1.4 The chapter's YMYL self-application	312
§32.2 Substantive Architecture	313
§32.2.1 The four-conversation framework	313
§32.2.2 The audit-chain projection principle	314
§32.2.3 Per-stakeholder evidence-class inheritance	315
§32.2.4 The disclosure-spectrum framework	315
§32.3 MYTHOS / SecureAgent Mapping	316
§32.3.1 Inbound forward-reference resolution	316
§32.3.2 Per-stakeholder × per-evidence-class mapping	317
§32.3.3 Per-vector mapping inheritance	318
§32.3.4 EU AI Act Article 12 record-keeping inheritance	318
§32.4 Per-Stakeholder Narrative	319
§32.4.1 Board narrative — translating MYTHOS into board-level language	319
§32.4.2 Examiner narrative — operational evidence under published control taxonomy	321
§32.4.3 Regulator narrative — materiality-thresholded public disclosure	322
§32.4.4 Insurer narrative — per-control attestation under carrier questionnaire	324
§32.5 SecureAgent Production Surface	325
§32.5.1 Board reporting templates — quarterly and annual	325
§32.5.2 Examiner evidence templates — CRI-aligned + ISO 42001-aligned	326
§32.5.3 SEC AI disclosure templates — Form 10-K Item 105 + Form 8-K Item 1.05 scaffold . . .	326
§32.5.4 Insurance attestation templates — annual carrier questionnaire + claim-time evidence package	327
§32.6 Conformance Block	328
§32.6.A Capability — board / examiner / regulator / insurer evidence production	328
§32.6.B Evidence — what the institution produces and retains	329
§32.6.C Gaps — what the chapter's discipline does not cover	329
§32.6.D Forward-Fits — Sprint 7+ extensions	330
§32.6.E Operational Path — board / examiner reporting cadence rollout	330
§32.7 CISO's Five Next Steps	331
Without SecureAgent — What You Can Do Today	332
References	333
Cross-References	333

Chapter 33: The Next 18 Months: What to Watch	334
Chapter Purpose Statement	334
§33.1 Operational Context	334
33.1.1 The forecast-readiness problem	334
33.1.2 Where you are in the deployment arc	335
33.1.3 Audience	335
33.1.4 The four trajectory axes — chapter framework in compressed form	336
33.1.5 What the chapter operationalizes	336
§33.2 Substantive Architecture	337
33.2.1 The trajectory-monitoring framework	337
33.2.2 The four sub-disciplines	337
33.2.3 Citation freshness as evidentiary discipline	338
33.2.4 Forecasting framing — language discipline	338
33.2.5 The decision-trigger framework	338
33.2.6 Integration with Chapter 34’s calendar-anchored roadmap	339
§33.3 MYTHOS / SecureAgent Mapping	339
33.3.1 Inbound forward-reference resolution	339
33.3.2 Per-vector trajectory mapping	340
33.3.3 Per-pipeline-stage trajectory mapping	341
33.3.4 Per-framework trajectory mapping	342
§33.4 Per-Trajectory Narrative	343
33.4.1 Frontier-model capability trajectory	343
33.4.2 Regulatory trajectory	344
33.4.3 Research-frontier trajectory	345
33.4.4 IETF / standards-body trajectory	346
33.4.5 SecureAgent capability roadmap (forecast posture)	347
§33.5 What to Watch + Decision Triggers (Forecast-Readiness Anchors) [ADAPTED]	348
33.5.1 Anchor design discipline	348
33.5.2 Frontier-model capability anchors	349
33.5.3 Regulatory-trajectory anchors	349
33.5.4 Research-frontier anchors	350
33.5.5 Standards-body anchors	350
33.5.6 Forward-references to Chapter 34 calendar-anchored roadmap	351
§33.6 Conformance Block	351
§33.6.A Capability — forecasting framework + monitoring discipline	351
§33.6.B Evidence — what the institution produces and retains	351
§33.6.C Gaps — what the discipline does not cover	352
§33.6.D Forward-Fits — Sprint 8+ extensions and future-edition cycles	352
§33.6.E Operational Path — strategic-planning cadence integration	353
§33.7 CISO’s Five Next Steps	353
References	354
 Chapter 34: The CISO’s 12-Month MYTHOS Implementation Roadmap	 355
§34.1 Months 1–3 — Foundation Phase	355
§34.2 Months 4–6 — Operational Phase	358
§34.3 Months 7–9 — Maturation Phase	361
§34.4 Months 10–12 — Optimization Phase	364
§34.5 Success Metrics + 18-Month Outlook	368
References	370
Cross-References	370
 Appendix A: The MYTHOS Technique-Page Template	 371
How to use this template	371
§A.1 Identification	371
§A.2 Description	372
§A.3 Sub-techniques	372
§A.3.1 Sub-technique 1	372

§A.3.2 Sub-technique 2	372
§A.3.3 Sub-technique 3	372
§A.3.N Additional sub-techniques	373
§A.4 Procedure Examples	373
§A.4.1 Procedure example 1	373
§A.4.2 Procedure example 2	373
§A.4.N Additional procedure examples	373
§A.5 Mitigations	373
§A.5.1 Mitigation 1	373
§A.5.2 Mitigation 2	374
§A.5.3 Mitigation 3	374
§A.5.N Additional mitigations	374
§A.6 Detection	374
§A.7 MYTHOS Certification Data	375
§A.7.1 Per-technique corpus	375
§A.7.2 Confusion matrix	375
§A.7.3 Statistical confidence	375
§A.7.4 Sub-technique-level breakdown (if data supports)	375
§A.7.5 Annotation discipline	375
§A.7.6 Architectural-separation discipline reminder	375
§A.8 CISO's Five Next Steps	376
§A.9 Cross-references	376
End of form	377

Appendix B: Confusion-Matrix Worksheet with Clopper-Pearson Calculator 378

B.1 How to Use This Appendix	378
B.2 The Worksheet Schema — TP / FP / FN / TN, Per-Vector and Corpus-Aggregate Rows, Clopper-Pearson Lower-Bound Output Cell	379
B.2.1 The four cells	379
B.2.2 The schema rows	379
B.2.3 The Clopper-Pearson lower-bound output cell	380
B.2.4 The pitfalls the schema prevents	380
B.3 Excel Computational Reference	380
B.3.1 Cell layout	380
B.3.2 The Clopper-Pearson lower-bound formula	381
B.3.3 Worked Excel reference — the corpus-aggregate row	381
B.3.4 Excel sanity-check column	382
B.4 Python Computational Reference	382
B.4.1 The script	382
B.4.2 What the script produces	384
B.4.3 Determinism	385
B.5 Per-Claim Worked-Example Library	385
B.5.1 Worked example one — the MR-1 corpus aggregate (5,857 / 0 / 100% / 99.65% three-sigma)	385
B.5.2 Worked example two — the per-vector confusion matrices (T1 through T7)	385
B.5.3 Worked example three — the 14,208-trial / 1.9636-TES / 98.2% internal evaluation	386
B.5.4 Worked example four — the MITRE ER7 cohort identity-protection rate (0% across 9 vendors)	387
B.5.5 Worked example five — the EDR industry-average false-positive rate (~1 in 3)	387
B.6 The [VectorCertain Internal] vs. [Third-Party-Reproduced] Annotation Discipline (Canonical Reference)	388
B.6.1 The two categories	388
B.6.2 Where the annotation appears	388
B.6.3 Self-application — the worksheet catches its own publisher	388
B.6.4 Worked-example cross-reference	388
B.7 Crumpton/MITRE Governance Overlay (Mandatory at Every ER8 / TES Reference)	389
B.8 What to Read Next	389

Appendix C: Full Cross-Reference Matrix	391
C.1 What This Appendix Is	391
C.2 How to Use the Matrix	392
C.3 Primary Cross-Reference Matrix — MYTHOS T1–T7 × Seven Frameworks	393
C.3.1 T1 — Autonomous Multi-Step Exploitation	393
C.3.2 T2 — Unsanctioned Scope Expansion	395
C.3.3 T3 — Invisible Deceptive Reasoning	396
C.3.4 T4 — Track-Covering Log Manipulation	397
C.3.5 T5 — Credential Theft and System Access	398
C.3.6 T6 — Sandbox Escape Exploitation	400
C.3.7 T7 — Capability Proliferation	401
C.4 Expanded Framework Axis — CSF 2.0, EU AI Act, FFIEC, PCI DSS 4.0, SOC 2	403
C.5 Per-IdP Non-Human-Identity Governance Cross-Walk	405
C.6 Crumpton/MITRE Positioning at Framework Cross-Walk Cells	410
Appendix D: The 828-Model MRM-CFS Reference Card	411
D.1 Purpose and How to Read This Card	411
D.2 Locked-Figures Registry Entries — The Six Architectural Specification Figures	411
D1 — Ensemble Cardinality (828 Micro-Recursive Models / 828 Segments)	412
D2 — Per-Segment Size Range (29–500 Bytes)	412
D3 — Per-Segment Processing Latency (<0.3 ms; PR-17 0.27 ms)	412
D4 — Accuracy Target (99.20%+)	413
D5 — Architecture (2-Tier Classify→Quantify)	413
D6 — Five Reachable Determinations (PERMIT / INHIBIT / DEFER / DEGRADE / ES-CALATE)	413
D.3 Functional Classifier Groups — The 828 Segments by Operational Pattern	414
Tier-One Functional Classifier Groups (Approximately 720 of 828 Segments)	414
Tier-Two Cascading-Fusion Functional Classifier Groups (Approximately 100 of 828 Segments)	415
Reserved Segments (Approximately 8 of 828)	415
D.4 Per-Group Profiles — Latency, Accuracy Contribution, Pipeline Integration	416
D.5 Ensemble Behavior — Voting, Abstention, Determination Composition	416
D.6 Verification Notes and Sealed Cross-References	417
Appendix E: MYTHOS Pipeline Rule Reference (44-Rule Taxonomy)	418
E.0 What This Appendix Is	418
E.1 AMRS V4 Admission Rules — Govern What Goes In	419
E-AMRS-1 — Monotonicity (BG-1)	419
E-AMRS-2 — Weight Normalization (BG-2)	419
E-AMRS-3 — Score Boundedness (BG-3)	419
E-AMRS-4 — P-Gate Supremacy (BG-4)	419
E-AMRS-5 — V-Gate Supremacy (BG-5)	420
E-AMRS-6 — Threshold Enforcement (BG-6)	420
E-AMRS-7 — Determinism (BG-7)	420
E-AMRS-8 — Audit Completeness (BG-8)	420
E-AMRS-9 — Tier 4 Hard Invalidation (BG-9)	420
E-AMRS-10 — Hash Integrity (BG-10)	420
E.2 HES1-SG Candidate-Diversity Rules — Gate 1	421
E-HES1-1 — Tier-1 Roster Architectural Diversity Floor	421
E-HES1-2 — Availability-Adjusted Roster Target (33 Models)	421
E-HES1-3 — Action-Class Registry, Default-Engaged	421
E.3 HCF2-SG Independence Cascade — Gate 2 (L1–L4)	421
E-HCF2-1 — L1: Architectural Independence Index Floor	422
E-HCF2-2 — L1: Adversarial Transferability Independence Index Floor	422
E-HCF2-3 — L1: Structural Tier Diversity (2 Distinct Architecture Tiers)	422
E-HCF2-4 — L2: Effective Voter Count Floor	422
E-HCF2-5 — L3: Cross-Coincident Failure Diversity Ceiling	422
E-HCF2-6 — L3: Lyapunov Vote-Trajectory Stability Ceiling	422

E-HCF2-7 — L3: Likelihood-Ratio Significance Floor	422
E-HCF2-8 — L4: CONSOL SPRT Acceptance / Rejection / Defer Boundaries	423
E-HCF2-9 — L4: Domain-Specific / Configuration	423
E.4 TEQ-SG Numerical Admissibility Rules — Gate 3	423
E-TEQ-1 — Representation-Invariance of the Discrete Determination	423
E-TEQ-2 — Identity-Trust Scoring is Behavioral, Not Credential-Based	423
E-TEQ-3 — Per-Action Re-Scoring (No Inheritance)	423
E.5 MRM-CFS-SG Execution-Governance Rules — Gate 4	424
E-MRMSG-1 — Architectural Separation: MRM-CFS Detection vs. MRM-CFS-SG Governance	424
E-MRMSG-2 — Supervisory Sufficiency Precondition	424
E-MRMSG-3 — 828-Segment Cardinality Lock (D1)	424
E-MRMSG-4 — Per-Segment Specifications (D2 / D3 / D4)	424
E-MRMSG-5 — Five-Determination Reachability (D6)	425
E-MRMSG-6 — Two-Tier Classify→Quantify Architecture (D5)	425
E.6 AGL-SG GTID-and-Audit-Chain Rules — The Court Reporter	425
E-AGL-1 — Genesis-Record Construction	425
E-AGL-2 — Previous-Hash Linkage	425
E-AGL-3 — RFC 8785 JSON Canonicalization (JCS)	426
E-AGL-4 — GovernancePolicyPin Binding	426
E-AGL-5 — LayerDetermination Comprehensiveness — Seven-Element Schema	426
E-AGL-6 — Hash Determinism (U-20)	426
E-AGL-7 — Append-Only Invariant: AppendOnlyViolation Behavior	426
E-AGL-8 — INHIBITED-Record Comprehensiveness (U-18)	427
E-AGL-9 — Authorization-Token Derivation, TTL, and GATED Default Posture	427
E-AGL-10 — Containment Depth Bounding (U-19)	427
E-AGL-11 — Inter-Agent Reference Content-Addressing — CGTID and ClearanceGTID First- Class Status	427
E-AGL-12 — ECDSA P-256 Optional Cryptographic Signature	427
E-AGL-13 — ComplianceAuditEvent at Every Governance Transition (U-40)	428
E.7 Forward-Reference Index	428
E.8 References	428

Appendix F: Sample GTID Audit Record and Verification Script **429**

F.1 What This Appendix Is, and What It Is Not	429
F.2 The Sample GTID Audit Chain	429
F.2.1 Genesis Record (gtid_id = gtid-000000000000000001)	430
F.2.2 Record 2 (gtid_id = gtid-000000000000000002) — INHIBITED Structuring	430
F.2.3 Record 3 (gtid_id = gtid-000000000000000003) — AUTHORIZED Routine Read	432
F.2.4 What the Sample Demonstrates	433
F.3 The Verification Script	433
F.3.1 Script Listing	433
F.3.2 What Each Check Establishes	437
F.4 Per-Step Procedural Narrative	437
Step 1 — Export the Audit Chain into the Sample-Record JSON Schema	437
Step 2 — Run the Verification Script Against the Export	437
Step 3 — Interpret the Output	438
Step 4 — Produce the Auditor-Defensible Attestation	438
F.5 What the Script Does Not Do	438
F.6 Summary	439

Appendix G: Vendor RFP Language Library **440**

G.1 How to Use This Library	440
G.2 Disclosure Clauses	440
Clause G-D-01 — Provenance disclosure	441
Clause G-D-02 — Methodology source disclosure	441
Clause G-D-03 — Confusion-matrix decomposition disclosure	441
Clause G-D-04 — Per-vector decomposition (non-blending discipline)	441

Clause G-D-05 — Latency-context disclosure	442
Clause G-D-06 — [Vendor-Internal Validation] versus [Third-Party-Reproduced] annotation discipline	442
G.3 Statistical Field-Test Clauses	442
Clause G-S-01 — Per-vector Clopper-Pearson three-sigma lower bound	442
Clause G-S-02 — Per-vector scenario-count sufficiency	443
Clause G-S-03 — Reproducibility-evidence requirements	443
Clause G-S-04 — Sub-technique cross-walk	443
Clause G-S-05 — TES-equivalent disclosure (where the vendor publishes a TES-style result)	444
G.4 Ongoing-Validation Clauses	444
Clause G-V-01 — Quarterly re-evaluation cadence	444
Clause G-V-02 — Drift-detection thresholds and contractual remedies	444
Clause G-V-03 — AMRS V4 baseline-refresh discipline (sub-hour cadence)	445
Clause G-V-04 — Roadmap disclosure for capability gates and pathway extensions	445
G.5 Service-Credit Clauses	445
Clause G-C-01 — Service-credit guarantee tied to per-vector lower bound	446
Clause G-C-02 — Deployment-cost denominator-and-numerator transparency	446
Clause G-C-03 — Per-quantization-tier performance disclosure	446
G.6 Audit-Trail Clauses	447
Clause G-A-01 — Governance-chain audit-trail export	447
Clause G-A-02 — Hash-determinism verification	447
Clause G-A-03 — Policy-pin disclosure	448
G.7 Cross-Walk Clauses	448
Clause G-X-01 — Framework-cross-walk publication	448
Clause G-X-02 — Crumpton/MITRE positioning preservation	448
G.8 Crumpton/MITRE Governance Overlay	449
G.9 Cross-Vendor Applicability Disclosure	449
G.10 Adoption Path	450
G.11 Cross-References	450

Appendix H: Glossary 451

How to use this glossary	451
A	451
Abstention	451
Action Approval Gate	452
Action Risk Classification	452
Adaptive Memory Relevance Scoring — see AMRS.	452
Adaptive Risk Scoring Classification — see ARSC.	452
Adversarial Threat Landscape for Artificial-Intelligence Systems — see ATLAS.	452
Adversarial Transferability Independence Index — also ATII	452
AgentCertain	452
Agent Governance Layer — Safety & Governance — see AGL-SG.	452
Agent Passport	452
AGL-SG — Agent Governance Layer — Safety & Governance	452
AII — Architectural Independence Index	452
AMRS — Adaptive Memory Relevance Scoring	452
AMRS V4	453
AppendOnlyViolation	453
Architectural Independence Index — see AII.	453
ARSC — Adaptive Risk Scoring Classification	453
ATII — see Adversarial Transferability Independence Index.	453
ATLAS — Adversarial Threat Landscape for Artificial-Intelligence Systems (MITRE)	453
ATT&CK — see MITRE ATT&CK.	453
AUTHORIZED	453
B	453
BG-1 through BG-10	453
Bidirectional Security Envelope	453

	Block-time discipline	453
C		454
	Candidate Diversity	454
	CFD — see Coincident Failure Diversity.	454
	Coincident Failure Diversity — also CFD	454
	Compliance Audit Event	454
	CONSOL SPRT	454
	Control Objective	454
	CRI FS AI RMF — Cyber Risk Institute Financial Services AI Risk Management Framework	454
	CRI Profile v1.2.1	454
	Crumpton — see Crumpton / MITRE Positioning.	455
	Crumpton / MITRE Positioning	455
	CSA NHI Framework — Cloud Security Alliance Non-Human Identity Management Framework	455
D		455
	D-1 Artifact	455
	DC — Detection Context (TES coefficient)	455
	DEFER / DEFERRED	455
	DEGRADE	455
	Detection Context — see DC.	455
	Detection Performance — see DP.	455
	Detection Sufficiency — see DS.	455
	Determination	455
	DP — Detection Performance (TES coefficient)	456
	DS — Detection Sufficiency (TES coefficient)	456
E		456
	Effective Voter Count — also n_{eff}	456
	Epistemic Trust Governance	456
	ESCALATE / ESCALATED	456
	Execution Governance — see MRM-CFS-SG.	456
F		456
	Forbidden Language Registry	456
	Functional Classifier Group	456
G		457
	Gate-decision time — see Block-time discipline.	457
	GovernanceChain	457
	GovernanceDetermination	457
	Governance Pipeline	457
	GovernancePolicyPin	457
	Governance Transaction Identifier — see GTID.	457
	GTID — GovernanceTransactionIdentifier	457
H		457
	HCF2-SG — Hierarchical Cascading Framework — Safety & Governance	457
	HES1-SG — Hybrid Ensemble System — Safety & Governance	457
	Hierarchical Cascading Framework — Safety & Governance — see HCF2-SG.	458
	H-Neuron Homology Hypothesis	458
	H-Neurons	458
	HOTS — H-Neuron Overcompliance Test Suite	458
	HOTSResult	458
	Hybrid Ensemble System — Safety & Governance — see HES1-SG.	458
I		458
	Independence Cascade	458
	INHIBIT / INHIBITED	458
	ISO/IEC 23894:2023	458
	ISO/IEC 42001:2023	459
	ISO/IEC 5338:2023	459
L		459
	Living Corpus	459

M	459
Micro-Recursive Model — Cascading Fusion System — see MRM-CFS.	459
Micro-Recursive Model — Cascading Fusion System — Safety & Governance — see MRM-CFS-SG.	459
Microsoft Entra Agent ID	459
MITRE ATLAS — see ATLAS.	459
MITRE ATT&CK	459
MITRE ATT&CK Evaluations — see Crumpton / MITRE Positioning, MITRE ER7, MITRE ER8.	459
MITRE ER7	459
MITRE ER8	459
MRM-CFS — Micro-Recursive Model — Cascading Fusion System	460
MRM-CFS-SG — Micro-Recursive Model — Cascading Fusion System — Safety & Governance	460
MYTHOS	460
N	460
n_eff — see Effective Voter Count.	460
NHI — Non-Human Identity	460
NIST AI RMF — National Institute of Standards and Technology Artificial Intelligence Risk Management Framework	460
Numerical Admissibility	460
O	460
Obviousness-Type Double Patenting — also ODP	460
OWASP Agentic Top 10	461
OWASP Foundation	461
OWASP LLM Top 10 v2025	461
P	461
PERMIT	461
P-gate — Provenance Gate	461
Post-Execution Detection	461
PP — Process Performance (TES coefficient)	461
Pre-Execution Governance	461
Priority Classification Routing — also PCR	462
Project Glasswing	462
Provenance	462
R	462
Retired Form	462
ROME	462
S	462
S1 Escalation	462
S1 / S2 / S3 — HOTS severity levels	462
Safety Envelope	462
Sandwich	463
SecureAgent	463
Sub-objective	463
Sub-technique ID	463
Supervisory Sufficiency	463
T	463
T1 — Autonomous Multi-Step Exploitation	463
T2 — Unsanctioned Scope Expansion	463
T3 — Invisible Deceptive Reasoning	463
T4 — Track-Covering Log Manipulation	463
T5 — Credential Theft	464
T6 — Sandbox Escape	464
T7 — Capability Proliferation	464
TEQ-SG — Trust & Execution Governance — Safety & Governance	464
TES — Total Evaluation Score	464
Test 18 (Use vs. Mention Discriminator)	464

Total Evaluation Score — see TES.	464
Trust & Execution Governance — Safety & Governance — see TEQ-SG.	464
Two-Layer Defense	464
V	464
Validation	464
VectorAgents	465
VectorCertain	465
VectorTransForm	465
V-gate — Validation Gate	465
Z	465
Zero-knowledge GTID — see ZGTID.	465
ZGTID — Zero-knowledge GTID	465
Numeric / Compound Terms	465
14,208 / 38 / 0 — TES validation triple	465
17-Part MYTHOS Threat Intelligence Series	465
44-Rule Pipeline — see Pipeline Rule Reference (FLAG-008 / Appendix E)	466
508-Point Unified Governance	466
81.4 Percent Cross-Model Correlation	466
828-Model MRM-CFS Ensemble	466
Retired-Form Quick Reference	466
Source-of-Truth Authority	467

Appendix I: Further Reading: Annotated Bibliography 468

How to use this appendix	468
I.1 Desk-Reference Companions	468
I.1.1 Bill Bonney, Gary Hayslip, and Matt Stamper. <i>CISO Desk Reference Guide: A Practical Guide for CISOs</i> . Volumes 1 and 2. CISO DRG Joint Venture, Volume 1 (1st edition 2016, 2nd edition 2019), Volume 2 (2018). Available at https://cisodrg.com/	469
I.1.2 Raj Badhwar. <i>The CISO's Next Frontier: AI, Post-Quantum Cryptography, and Advanced Security Paradigms</i> . Springer, 2021. Available at https://www.booksamillion.com/p/Cisos-Next-Frontier/Raj-Badhwar/9783030753535	469
I.1.3 Rick Howard. <i>Cybersecurity First Principles: A Reboot of Strategy and Tactics</i> . Wiley, 2023. Available at https://www.wiley.com/en-us/Cybersecurity+First+Principles:+A+Reboot+of+Strategy+and+Tactics-p-9781394173099	469
I.2 Threat-Practitioner Anchors	469
I.2.1 Marcus J. Carey and Jennifer Jin. <i>Tribe of Hackers</i> series — <i>Tribe of Hackers: Cybersecurity Advice from the Best Hackers in the World</i> (2019); <i>Tribe of Hackers Red Team</i> (2019); <i>Tribe of Hackers Blue Team</i> (2020); <i>Tribe of Hackers Security Leaders</i> (2020). Wiley.	470
I.2.2 Kevin D. Mitnick (with William L. Simon). <i>The Art of Intrusion: The Real Stories Behind the Exploits of Hackers, Intruders and Deceivers</i> . Wiley, 2005.	470
I.3 Operational-Practice References	470
I.3.1 Richard Bejtlich. <i>The Practice of Network Security Monitoring: Understanding Incident Detection and Response</i> . No Starch Press (in association with O'Reilly distribution channels), 2013. Available at https://www.oreilly.com/library/view/the-practice-of/9781457185175/	470
I.3.2 Cody Roberts and Erik Brown (Editors). <i>Practical Threat Intelligence and Data-Driven Threat Hunting: A Hands-On Guide to Threat Hunting with the ATT&CK Framework and Open-Source Tools</i> . Second Edition. Packt Publishing (distributed via O'Reilly), 2024. Available at https://www.oreilly.com/library/view/practical-threat-intelligence/9781803233758/	471
I.4 Statistical-Methodology References	471
I.4.1 Air Force Institute of Technology, Scientific Test and Analysis Techniques Center of Excellence (STATCOE). <i>Confidence Intervals for Proportions: Methods, Comparisons, and Recommendations</i> . Technical brief (current edition). Available at https://www.afit.edu/STAT/	471

I.4.2	Wikipedia. <i>Binomial proportion confidence interval</i> . Available at https://en.wikipedia.org/wiki/Binomial_proportion_confidence_interval	471
I.4.3	Wikipedia. <i>68–95–99.7 rule</i> . Available at https://en.wikipedia.org/wiki/68%E2%80%99395%E2%80%9999.7_rule	472
I.4.4	Yujie Gao, Xiaofeng Wang, et al. (Tsinghua University). <i>H-Neurons: A Mechanistic Account of Hallucination and Safety Bypass in Large Language Models</i> . arXiv:2512.01797v2, December 2024. Available at https://arxiv.org/abs/2512.01797	472
I.5	Standards-and-Frameworks References	472
I.5.1	National Institute of Standards and Technology. <i>Artificial Intelligence Risk Management Framework (NIST AI 100-1)</i> . January 2023. Available at https://www.nist.gov/itl/ai-risk-management-framework . Companion: NIST AI 600-1, <i>Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile</i> , July 2024.	472
I.5.2	Cyber Risk Institute (in collaboration with U.S. Department of the Treasury). <i>CRI Financial Services Artificial Intelligence Risk Management Framework</i> . February 2026. Available at https://cyberriskinstitute.org/ . Companion: <i>CRI Profile</i> (current version 1.2.1, February 2026).	473
I.5.3	International Organization for Standardization and International Electrotechnical Commission. <i>ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system</i> . December 2023. Available at https://www.iso.org/standard/81230.html . Companions: ISO/IEC 23894:2023 <i>Information technology — Artificial intelligence — Guidance on risk management</i> ; ISO/IEC 5338:2023 <i>Information technology — Artificial intelligence — AI system life-cycle processes</i>	473
I.5.4	OWASP Foundation. <i>OWASP Top 10 for Large Language Model Applications v2025</i> and <i>OWASP Top 10 for Agentic AI Applications</i> (December 2025). Available at https://genai.owasp.org/llm-top-10/ and https://genai.owasp.org/llm-top-10/agentic-ai/	473
I.6	Crumpton/MITRE Correspondence Entry	474
I.6.1	Crumpton, Lex. <i>MITRE ATT&CK Evaluations Technical Lead correspondence with VectorCertain LLC</i> . April 8, 2026. Sealed canonical reference at mythos-source-library/citations/mitre	474
	How this bibliography is maintained	474
References		475
Bibliography		476
About the Author		477
Colophon		478

The MYTHOS Playbook, Volume II

The Governance Model

by Joseph P. Conroy

Founder & CEO, VectorCertain, LLC

VectorCertain, LLC

Casco, Maine

2026

Copyright

The MYTHOS Playbook, Volume II: The Governance Model

© 2026 VectorCertain, LLC. All rights reserved.

First Edition (EPUB) published 2026 by VectorCertain, LLC, Casco, Maine.

ISBN 979-8-9960296-5-5 (EPUB).

Format: EPUB3 (reflowable digital edition).

No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from the publisher, except by a reviewer who may quote brief passages in a review.

Two-Volume Set — Multi-Format Edition

This is the EPUB edition of the second of a two-volume set. Each volume is published in three formats:

- **Volume I: The Adversary Landscape**
 - Paperback (ISBN 979-8-9960296-0-0)
 - Hardcover (ISBN 979-8-9960296-6-2)
 - EPUB (ISBN 979-8-9960296-3-1)
- **Volume II: The Governance Model**
 - Paperback (ISBN 979-8-9960296-4-8)
 - Hardcover (ISBN 979-8-9960296-7-9)
 - **EPUB (ISBN 979-8-9960296-5-5) — this edition**

Each format is a self-contained reference work with the full bibliography (297 entries) and Appendices A–I reproduced in both volumes.

A Note on the EPUB Edition

This EPUB is a reflowable digital edition designed for the working CISO who reads on a Kindle, Apple Books, or other EPUB3-compatible reader. Footnotes appear as on-tap popups on devices that support EPUB3 popup notes (Kindle, Apple Books) and as inline endnote markers on legacy readers. The print indexes (paperback + hardcover) are NOT reproduced in this digital edition; full-text search via the reader’s native search tools is the equivalent navigation tool. Cross-references to Volume I appear as plain text (e.g., “see Volume I, Chapter 8”); to access that material, locate the companion volume by ISBN at your retailer.

Legal and Editorial Disclosures

This book is for informational purposes for cybersecurity, governance, and risk leaders, and does not constitute legal, financial, or security advice for any specific environment. All implementation decisions should be made in consultation with qualified counsel, your institution’s risk and compliance teams, and current regulatory guidance.