MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / CYBERSECURITY AND NETWORK SECURITY

Vulnerability Management and Exposure Assessment Standard

Defines discovery, risk, prioritization, remediation, exceptions, proof, and metrics.

PERMANENT DOCUMENT ID
MYTHOS-SEC-0008

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
8 Atomic Criteria / 4 Families

PERMANENT PUBLICATION IDENTITY

Vulnerability Management and Exposure Assessment Standard

Defines discovery, risk, prioritization, remediation, exceptions, proof, and metrics.

vulnerability-management

DOCUMENT ID
MYTHOS-SEC-0008

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

8	4	4	1
ATOMIC CRITERIA	CAPABILITY FAMILIES	ASSESSMENT PROFILES	PERMANENT IDENTITY

Publication doctrine. This standard establishes durable requirements for its defined scope. Interpretation must preserve the permanent document identity, recorded evidence, and formal revision history.

DOMAIN CONTROL SYSTEM

MYTHOS-SEC-0008 inside the cybersecurity and network security standard constellation

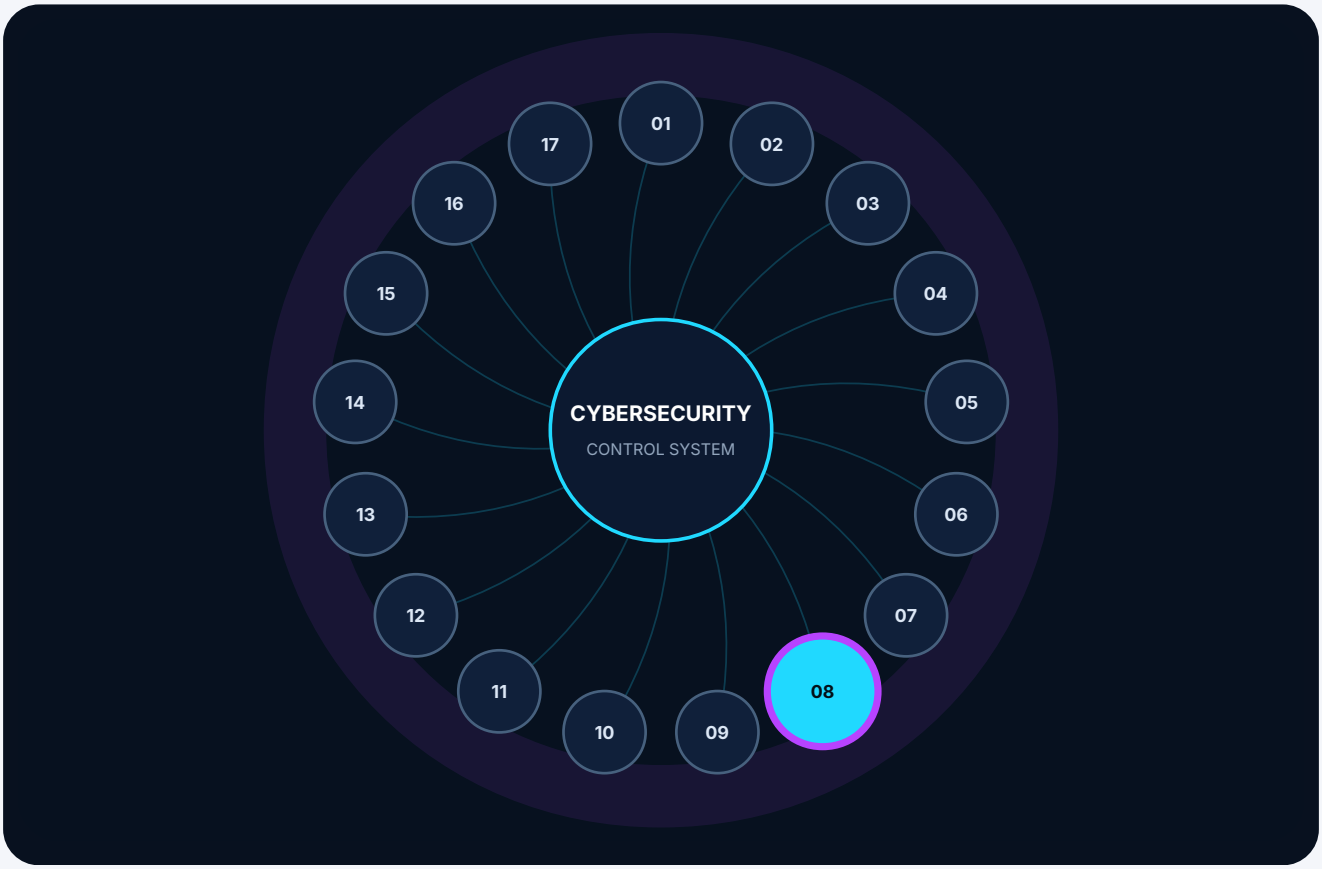

Connected assurance

Architecture, policy, implementation, operations, and evidence are evaluated as one controlled system.

Specialization rule

Companion standards may add precision, but they cannot silently weaken this publication's requirements.

CLOSED-LOOP ARCHITECTURE

Intent becomes trustworthy only when observed state returns as evidence

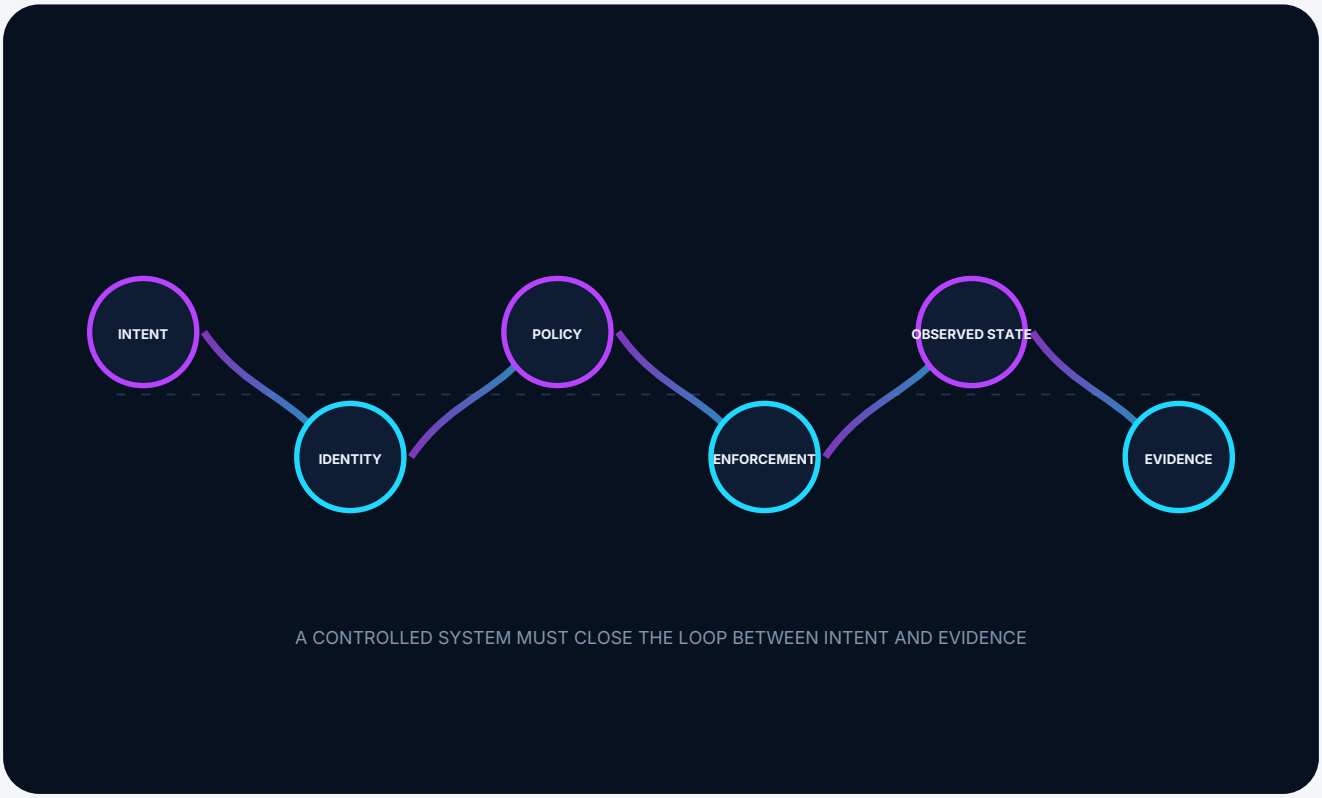

State separation

Desired, approved, deployed, observed, historical, and simulated states must remain distinguishable.

Recoverable change

Material change requires validation, authorization, monitoring, containment, and a credible recovery path.

EVALUATION TOPOLOGY

The capability families and atomic criteria of MYTHOS-SEC-0008

Capability families

Families expose the major engineering surfaces and preserve the source weighting model.

Atomic criteria

Each criterion carries a question, maturity spectrum, evidence expectations, assessment method, and operational signals.

NAVIGATION

Contents

ENGINEERING DOCTRINE

0. Executive Mandate

Traditional evaluation of vulnerability management is insufficient.

Organizations measure security by the number of Critical CVEs they can patch in 30 days. They generate lists of 100,000 vulnerabilities, hand them to an understaffed SOC, and wonder why they get breached. A flat list of CVEs without business context, network topology, and exploitability analysis is not a vulnerability management program; it is a liability inventory.

This standard exists because:

1. **CVE Counts are Meaningless:** A Critical CVE on an isolated, internal, air-gapped server is irrelevant. A Medium CVE on an internet-facing API gateway is a ticking time bomb. Vulnerability management MUST be contextual, factoring in network reachability, asset criticality, and threat intelligence.
2. **Exposure Management is the Evolution:** The industry is shifting from Vulnerability Management (scanning code) to Exposure Management (analyzing the composite attack surface). The system MUST map CVEs to active attack paths, not just static asset databases.
3. **AI Introduces New Vulnerability Classes:** Traditional VM tools scan code dependencies. They do not scan AI model weights for backdoors, prompt injection vulnerabilities, or poisoned training data. An AI Bill of Materials (AIBOM) is required to manage this new attack surface.
4. **Automated Remediation is Dangerous:** Giving an AI agent the authority to auto-patch production infrastructure based on a vulnerability report is a recipe for an outage. Patches MUST be simulated and verified in a digital twin before deployment.

This document establishes the **Mythos Vulnerability & Exposure Standard (MVES)**, a comprehensive, evidence-based methodology for evaluating the contextual prioritization, attack path analysis, and remediation governance of security vulnerabilities.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Vulnerability Management (VM) platforms and scanners (SAST, DAST, SCA)
- Exposure Management and Attack Surface Management (ASM)
- Contextual prioritization (EPSS, asset criticality, network reachability)
- Attack graph generation and blast radius analysis
- AIBOM generation and AI-specific vulnerability scanning (weight poisoning, prompt injection)
- Remediation governance and patch simulation
- Continuous compliance and reporting

1.2 Out of Scope

- General network security architecture (covered in MYTHOS-SEC-0007)
- Firewall policy enforcement (covered in MYTHOS-SEC-0006)

1.3 Audience

- Security engineers and architects designing VM/EM pipelines
- SOC analysts and incident responders
- Platform engineering teams managing patch deployments
- AI governance, risk, and compliance teams
- Standards bodies seeking a reference vulnerability methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Vulnerability Dimensions

Dimension	Definition
Exposure Management	The continuous process of identifying, prioritizing, and remediating vulnerabilities based on their actual exploitability and impact within a specific environment.
EPSS	Exploit Prediction Scoring System. A data-driven model for predicting the likelihood of exploitability of a vulnerability in the wild.
Attack Graph	A mathematical map of all paths an attacker can take through a network, mapping vulnerabilities to network topology.
AIBOM	AI Bill of Materials. An inventory of model weights, training data sources, and dependencies used to identify AI-specific vulnerabilities.
Contextual Prioritization	Ranking vulnerabilities based on asset criticality, network exposure, threat intel, and EPSS, rather than static CVSS scores.

2.2 The Vulnerability Doctrine

> A list of CVEs is not a security program. A CVSS score is not a prioritization. A patch without simulation is an outage. An AI model is a vulnerability surface.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; flat CVE lists, no context
1	Basic scanning but no contextual prioritization
2	Functional but lacks attack graph analysis
3	Solid production performance; contextual scoring, AIBOM support
4	Strong performance; attack graph integration, patch simulation
5	Best-in-class; sets the standard for exposure management

Weighting

Category	Weight	Rationale
Contextual Prioritization & EPSS	20%	Flat lists cause alert fatigue
Attack Graph & Blast Radius	20%	Reachability determines impact
AIBOM & AI Surface Scanning	15%	AI models are a new attack surface
Remediation & Patch Governance	15%	Auto-patching causes outages
Continuous Discovery & Compliance	10%	Point-in-time scans are obsolete
Evidence & Audit	10%	Unverified patches are a liability
Operational Lifecycle	10%	VM must be governed continuously

Total: 100%

A system MUST score at least 3.0 in Contextual Prioritization and Attack Graph to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

Capability claims are bounded by evidence, context, and reproducible assessment.

8 atomic criteria across 4 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Contextual Prioritization and EPSS

SOURCE WEIGHT 20%

4.1.1
Exploitability Scoring

4.1.2
Asset Criticality

MYTHOS-SEC-0008-EVAL-4.1.1

Exploitability Scoring

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Contextual Prioritization and EPSS	4.1.1	20%	Evidence-led

Does the system prioritize vulnerabilities based on real-world exploitability?

0
Pure CVSS-based prioritization

1
Basic manual threat intel integration

2
Automated threat intel feeds

3
EPSS integration and automated scoring

4
Contextual scoring (EPSS + asset criticality + network exposure)

5
Leading prioritization: formally verified scoring model, continuous EPSS updates, and AI-driven context correlation

ENGINEERING INTERPRETATION

This criterion evaluates whether exploitability scoring is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

MYTHOS-SEC-0008-EVAL-4.1.2

Asset Criticality

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Contextual Prioritization and EPSS	4.1.2	20%	Evidence-led

Does the system factor in the business criticality of the affected asset?

0
All assets treated equally

1
Basic asset tagging

2
Automated asset classification

3
Dynamic criticality based on blast radius

4
Real-time criticality updates based on network topology

5
Leading criticality: formally verified asset graphs and continuous exposure correlation

ENGINEERING INTERPRETATION

This criterion evaluates whether asset criticality is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

4.2 / CAPABILITY FAMILY

Attack Graph and Blast Radius

SOURCE WEIGHT 20%

4.2.1
Attack Path Mapping

4.2.2
Blast Radius Enforcement

MYTHOS-SEC-0008-EVAL-4.2.1

Attack Path Mapping

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Attack Graph and Blast Radius	4.2.1	20%	Evidence-led

Does the system map vulnerabilities to active network paths?

0
No network context

1
Basic subnet mapping

2
Static topology mapping

3
Dynamic topology mapping with reachability analysis

4
Real-time attack graph generation

5
Leading mapping: formally verified attack graphs, cryptographic topology sealing, and continuous blast radius calculation

ENGINEERING INTERPRETATION

This criterion evaluates whether attack path mapping is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

MYTHOS-SEC-0008-EVAL-4.2.2

Blast Radius Enforcement

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Attack Graph and Blast Radius	4.2.2	20%	Evidence-led

Does the system block changes that increase the blast radius of a vulnerability?

0
No enforcement

1
Basic advisory warnings

2
Policy gates on critical assets

3
Automated blocking of risky network changes

4
Digital twin simulation of exposure impact

5
Leading enforcement: formally verified exposure limits, cryptographic sealing, and automatic rollback

ENGINEERING INTERPRETATION

This criterion evaluates whether blast radius enforcement is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- asset inventory, scanner evidence, KEV/EPSS context, remediation tickets, and exception records
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. reproduce representative findings and verify remediation or containment
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- severity-only prioritization, duplicate assets, unowned findings, and expired exceptions
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- KEV remediation time
- exposure-weighted backlog
- reopen rate
- coverage of in-scope assets or flows

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

4.3 / CAPABILITY FAMILY

AIBOM and AI Surface Scanning

SOURCE WEIGHT 15%

4.3.1
AI Vulnerability Scanning

4.3.2
Supply Chain Integration

MYTHOS-SEC-0008-EVAL-4.3.1

AI Vulnerability Scanning

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
AIBOM and AI Surface Scanning	**4.3.1**	**15%**	**Evidence-led**

Does the system scan AI artifacts (weights, models, prompts) for vulnerabilities?

0
No AI scanning

1
Basic dependency scanning for AI libraries

2
AIBOM generation and tracking

3
Prompt injection and weight poisoning detection

4
Behavioral sandboxing of AI artifacts

5
Leading AI scanning: continuous behavioral analysis, formally verified safety boundaries, and adversarial red-team suites

ENGINEERING INTERPRETATION

This criterion evaluates whether ai vulnerability scanning is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- asset inventory, scanner evidence, KEV/EPSS context, remediation tickets, and exception records
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. reproduce representative findings and verify remediation or containment
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- severity-only prioritization, duplicate assets, unowned findings, and expired exceptions
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- KEV remediation time
- exposure-weighted backlog
- reopen rate
- coverage of in-scope assets or flows

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

MYTHOS-SEC-0008-EVAL-4.3.2

Supply Chain Integration

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
AIBOM and AI Surface Scanning	4.3.2	15%	Evidence-led

Does the system integrate vulnerability data with AIBOM/SBOM data?

0
No integration

1
Manual cross-referencing

2
Automated SBOM scanning

3
AIBOM and SBOM correlation

4
Real-time exposure analysis based on AIBOM

5
Leading integration: signed AIBOMs, transparency logs, and continuous monitoring

ENGINEERING INTERPRETATION

This criterion evaluates whether supply chain integration is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- asset inventory, scanner evidence, KEV/EPSS context, remediation tickets, and exception records
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. reproduce representative findings and verify remediation or containment
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- severity-only prioritization, duplicate assets, unowned findings, and expired exceptions
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- KEV remediation time
- exposure-weighted backlog
- reopen rate
- coverage of in-scope assets or flows

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

4.4 / CAPABILITY FAMILY

Remediation and Patch Governance

SOURCE WEIGHT 15%

4.4.1
Patch Simulation

4.4.2
AI Authority Isolation

MYTHOS-SEC-0008-EVAL-4.4.1

Patch Simulation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Remediation and Patch Governance	4.4.1	15%	Evidence-led

Are patches simulated before deployment?

0
No simulation; direct deployment

1
Basic dry-run

2
Text-based diff

3
Graph-based blast radius analysis

4
Digital twin simulation

5
Leading simulation: cryptographic plan sealing, rollback confidence, and formally verified topology

ENGINEERING INTERPRETATION

This criterion evaluates whether patch simulation is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

MYTHOS-SEC-0008-EVAL-4.4.2

AI Authority Isolation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Remediation and Patch Governance	4.4.2	15%	Evidence-led

> Can an AI agent auto-patch infrastructure without human verification?

0
Fully autonomous patching

1
AI proposes, auto-applies

2
AI proposes, human reviews text, applies

3
AI proposes, twin simulates, human reviews

4
AI proposes, twin simulates, policy gates, dual-control applies

5
Leading isolation: formally verified authority boundary, zero auto-patching

ENGINEERING INTERPRETATION

This criterion evaluates whether ai authority isolation is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

LIFECYCLE AND ASSURANCE

Part V. The Mythos Vulnerability Suite (MVS)

Traditional VM benchmarks measure scan speed. The MVS evaluates contextual prioritization, attack graph accuracy, and remediation safety.

5.1 Suite Composition

5.1.1 MVS-Context

- **Prioritization:** 100+ tests verifying vulnerabilities are ranked by EPSS and asset criticality, not CVSS.
- **Attack Graph:** 50+ tests verifying vulnerabilities are mapped to active network paths.

5.1.2 MVS-AI

- **AIBOM Scanning:** 50+ tests verifying the system detects poisoned weights and prompt injection vulnerabilities.
- **Patch Simulation:** 50+ tests verifying patches are simulated in a digital twin before deployment.

5.2 Execution Protocol

PUBLICATION BOUNDARY

Candidate standard. Permanent identity. Evidence before authority.

Approval requires designated technical, security, legal, accessibility, and governance review with recorded findings and dispositions.