MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / CYBERSECURITY AND NETWORK SECURITY

Security Architecture and Advanced Design Principles Standard

Defines trust boundaries, threat modeling, secure defaults, isolation, resilience, and assurance.

PERMANENT DOCUMENT ID
MYTHOS-SEC-0004

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
9 Atomic Criteria / 7 Families

PERMANENT PUBLICATION IDENTITY

Security Architecture and Advanced Design Principles Standard

Defines trust boundaries, threat modeling, secure defaults, isolation, resilience, and assurance.

DOCUMENT ID
MYTHOS-SEC-0004

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

9	7	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-0004 inside the cybersecurity and network security standard constellation

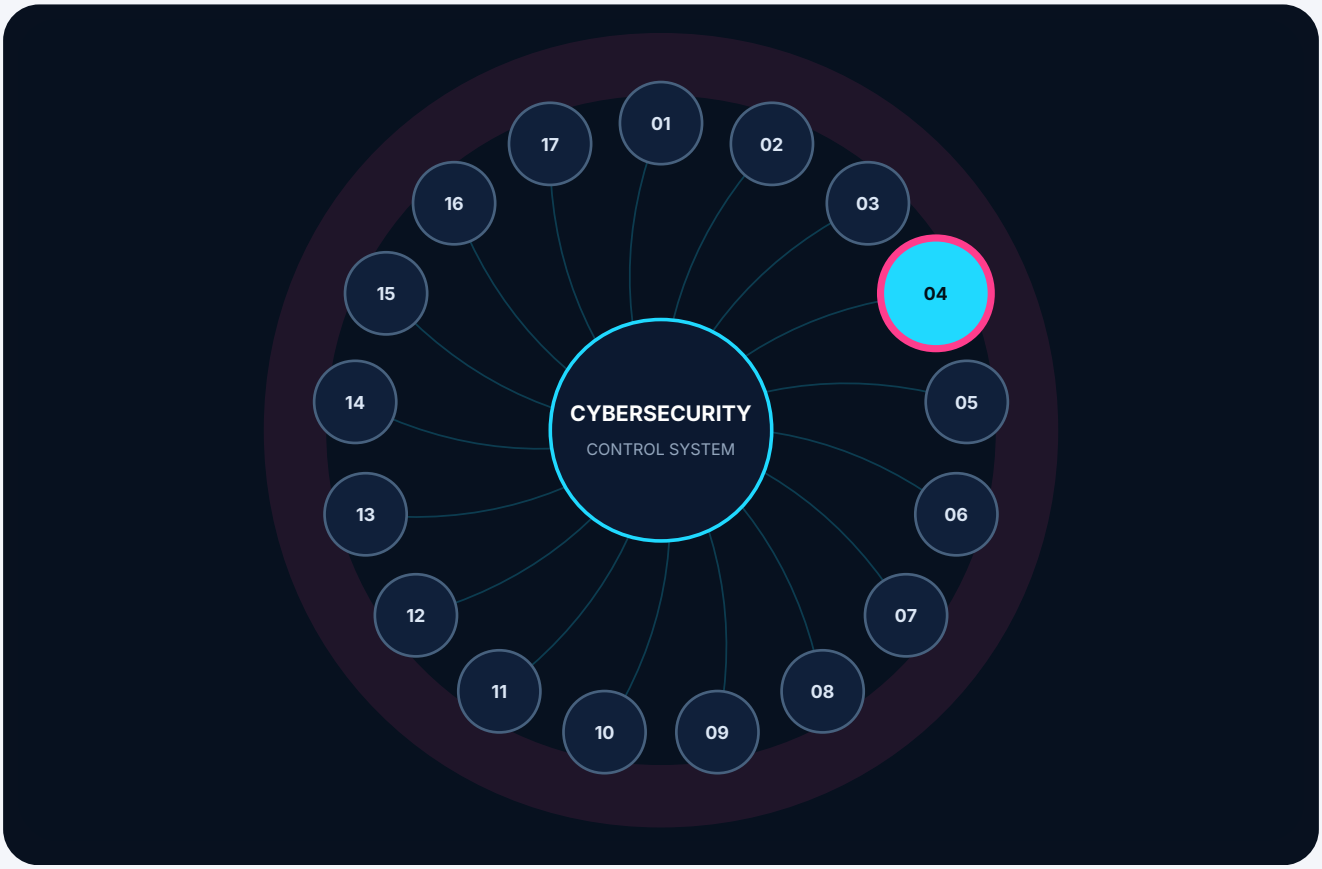

Connected assurance

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

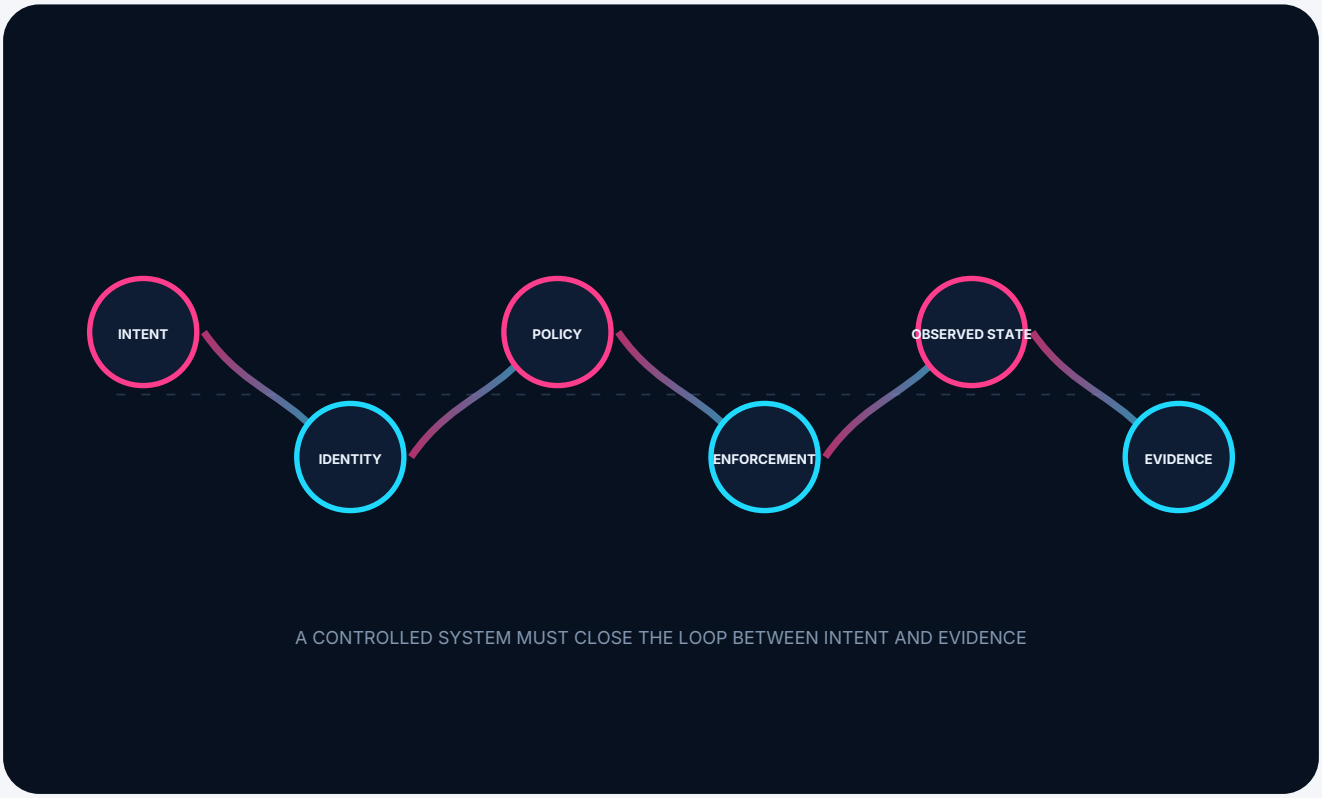

State separation

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-0004

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 software security architecture is insufficient.

Organizations measure security by compliance checklists (SOC2, ISO 27001) and vulnerability scan counts. They do not measure whether the system's architecture is fundamentally capable of surviving an intelligent, adaptive adversary or a catastrophic AI hallucination. A system can pass every compliance audit and still collapse the moment an LLM is granted tool access.

This standard exists because:

1. **Compliance is Not Architecture:** Checking boxes for encryption at rest and MFA does not guarantee that the system's internal trust boundaries are sound. Security MUST be an architectural property, not a post-deployment overlay.
2. **Threat Modeling is Mandatory, Not Optional:** If a system is designed without explicit abuse cases and threat models (STRIDE), it is designed to fail. The integration of unpredictable AI agents makes proactive threat modeling exponentially more critical.
3. **Formal Verification is the Only Proof:** Testing shows that bugs exist; formal verification (TLA+/Apalache) proves they do not. For critical state machines - like agent task graphs or authentication flows - testing is insufficient. Mathematical proof is required.
4. **Secure-by-Default is Non-Negotiable:** Systems must fail closed. Open by default is a liability. If a policy engine goes down, the system must halt, not grant access.

This document establishes the **Mythos Security Architecture Standard (MSAS)**, a comprehensive, evidence-based methodology for evaluating the foundational security principles of any complex technical platform.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Secure-by-design architectural principles (least privilege, defense in depth, fail-closed)
- Threat modeling and abuse case development (STRIDE, attack trees)
- Formal verification of critical state machines (TLA+, Apalache)
- Cryptographic architecture and crypto-agility
- Supply chain integrity and dependency isolation
- Incident response, break-glass procedures, and blast radius mitigation
- Zero-trust network and application architecture
- AI-specific security controls (prompt injection defense, authority separation)

1.2 Out of Scope

- Specific cryptographic algorithms (covered in MYTHOS-SEC-0001)
- Zero-trust policy enforcement (covered in MYTHOS-SEC-0003)
- Software supply chain scanning (covered in MYTHOS-SEC-0002)

1.3 Audience

- Principal engineers and architects designing platform infrastructure
- Security teams evaluating systemic risk
- AI governance, risk, and compliance teams
- Standards bodies seeking a reference security architecture methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Architecture Dimensions

Dimension	Definition
Secure-by-Design	Security is built into requirements and architecture, not patched at the end.
Defense in Depth	Multiple layers of security controls; if one fails, others remain.
Fail-Closed	If a security control fails, the system denies access rather than granting it.
Threat Modeling	The proactive process of identifying assets, adversaries, attack vectors, and mitigations.
Formal Verification	The use of mathematical logic to prove that a system meets its specification.
Blast Radius	The maximum potential damage if a component is compromised.

2.2 The Security Architecture Doctrine

> Compliance is not architecture. Testing is not proof. A system that fails open is a liability. A model that reasons without boundaries is a threat.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; no security architecture
1	Basic compliance but no architectural enforcement
2	Functional but lacks formal verification or threat modeling
3	Solid production performance; secure-by-design, threat modeled
4	Strong performance; formal verification, defense in depth
5	Best-in-class; sets the standard for security architecture

Weighting

Category	Weight	Rationale
Threat Modeling & Abuse Cases	20%	Unmodeled threats cause systemic failures
Secure-by-Design & Fail-Closed	20%	Open-by-default is a critical liability
Formal Verification	15%	Testing cannot prove the absence of bugs
Defense in Depth & Blast Radius	15%	Single points of failure cause outages
Cryptographic Architecture	10%	Hardcoded crypto is un-migratable
Incident Response & Break-Glass	10%	Unrecoverable systems are not production-ready
Zero-Trust & AI Boundaries	10%	Implicit trust is an attack surface

Total: 100%

A system MUST score at least 3.0 in Threat Modeling and Secure-by-Design to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

9 atomic criteria across 7 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Threat Modeling and Abuse Cases

SOURCE WEIGHT 20%

4.1.1
Proactive Threat Modeling

4.1.2
Abuse Case Development

MYTHOS-SEC-0004-EVAL-4.1.1

Proactive Threat Modeling

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Threat Modeling and Abuse Cases	**4.1.1**	**20%**	**Evidence-led**

Does the system architecture emerge from documented threat models?

0
No threat modeling; security is reactive

1
Basic threat modeling but not updated

2
STRIDE or equivalent applied to major components

3
Comprehensive threat models for all critical paths

4
Threat models drive architectural decisions and ADRs

5
Leading integration: threat models are living documents, continuously updated, and drive formal verification targets

ENGINEERING INTERPRETATION

This criterion evaluates whether proactive threat modeling 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-0004-EVAL-4.1.2

Abuse Case Development

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Threat Modeling and Abuse Cases	**4.1.2**	**20%**	**Evidence-led**

Are abuse cases explicitly defined and tested?

0
No abuse cases

1
Basic negative testing

2
Standard abuse cases (privilege escalation, injection)

3
AI-specific abuse cases (prompt injection, sycophancy)

4
Automated adversarial red-team suites based on abuse cases

5
Leading abuse case coverage: continuous red-team testing, formal verification of boundaries, and documented mitigations

ENGINEERING INTERPRETATION

This criterion evaluates whether abuse case development 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

Secure-by-Design and Fail-Closed

SOURCE WEIGHT 20%

4.2.1
Fail-Closed Posture

4.2.2
Least Privilege Enforcement

MYTHOS-SEC-0004-EVAL-4.2.1

Fail-Closed Posture

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Secure-by-Design and Fail-Closed	4.2.1	20%	Evidence-led

Does the system fail closed when security controls fail?

0
System fails open (grants access on policy engine failure)

1
Manual intervention required to restore security

2
Basic fail-closed but with high latency

3
Automated fail-closed with documented recovery

4
Fail-closed with zero trust and blast radius containment

5
Leading posture: formally verified fail-closed logic, zero implicit trust, and automatic quarantine

ENGINEERING INTERPRETATION

This criterion evaluates whether fail-closed posture 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-0004-EVAL-4.2.2

Least Privilege Enforcement

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Secure-by-Design and Fail-Closed	4.2.2	20%	Evidence-led

Is least privilege enforced at the architectural level?

0
Broad, shared credentials

1
Basic RBAC

2
Scoped roles but no dynamic adjustment

3
ABAC with dynamic scoping

4
Cryptographic capability grants with attenuation

5
Leading least privilege: formally verified boundaries, cryptographic isolation, and zero standing privileges

ENGINEERING INTERPRETATION

This criterion evaluates whether least privilege 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

- 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.

4.3 / CAPABILITY FAMILY

Formal Verification

SOURCE WEIGHT 15%

4.3.1
State Machine Verification

MYTHOS-SEC-0004-EVAL-4.3.1

State Machine Verification

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Formal Verification	4.3.1	15%	Evidence-led

Are critical state machines mathematically proven?

0
No formal verification

1
Basic unit tests

2
Property-based testing

3
Model checking (TLA+) for core logic

4
Formal verification of deadlock/livelock freedom

5
Leading verification: mathematical proof of all critical state transitions, continuously verified in CI

ENGINEERING INTERPRETATION

This criterion evaluates whether state machine verification 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.

4.4 / CAPABILITY FAMILY

Defense in Depth and Blast Radius

SOURCE WEIGHT 15%

4.4.1
Blast Radius Containment

MYTHOS-SEC-0004-EVAL-4.4.1

Blast Radius Containment

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Defense in Depth and Blast Radius	4.4.1	15%	Evidence-led

Is the blast radius of a compromised component minimized?

0
Flat architecture; one compromise = total takeover

1
Basic network segmentation

2
Application-level isolation

3
MicroVM/WASM isolation for untrusted code

4
Defense in depth: network, app, kernel, and hardware isolation

5
Leading containment: formally verified isolation boundaries, eBPF telemetry, and automatic quarantine

ENGINEERING INTERPRETATION

This criterion evaluates whether blast radius containment 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.

4.5 / CAPABILITY FAMILY

Cryptographic Architecture

SOURCE WEIGHT 10%

4.5.1
Crypto-Agility

MYTHOS-SEC-0004-EVAL-4.5.1

Crypto-Agility

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Cryptographic Architecture	4.5.1	10%	Evidence-led

Is the system architected for cryptographic agility?

0
Hardcoded algorithms

1
Config files but require code changes

2
Versioned algorithm registry

3
Dynamic algorithm rotation

4
PQC-ready with hybrid modes

5
Leading agility: formally verified crypto boundaries, AIBOM, and automated PQC migration triggers

ENGINEERING INTERPRETATION

This criterion evaluates whether crypto-agility 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

- cryptographic inventory, algorithm registry, key records, certificate data, and compatibility tests
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test rotation, hybrid negotiation, rollback, revocation, and algorithm-policy enforcement
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- hardcoded algorithms, unknown dependencies, missing rollback, and unverifiable keys
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- inventory coverage
- deprecated algorithm count
- rotation success 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.

4.6 / CAPABILITY FAMILY

Incident Response and Break-Glass

SOURCE WEIGHT 10%

4.6.1
Break-Glass Procedures

MYTHOS-SEC-0004-EVAL-4.6.1

Break-Glass Procedures

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Incident Response and Break-Glass	4.6.1	10%	Evidence-led

Are there governed, auditable break-glass procedures?

0
No break-glass procedures

1
Manual root access

2
Documented but unaudited

3
Automated with dual-control

4
Cryptographically signed break-glass with automatic evidence capture

5
Leading governance: dual-control, cryptographic attestation, zero-trust audit, and automatic rollback

ENGINEERING INTERPRETATION

This criterion evaluates whether break-glass procedures 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.

4.7 / CAPABILITY FAMILY

Zero-Trust and AI Boundaries

SOURCE WEIGHT 10%

4.7.1
AI Authority Isolation

MYTHOS-SEC-0004-EVAL-4.7.1

AI Authority Isolation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Zero-Trust and AI Boundaries	4.7.1	10%	Evidence-led

Is AI reasoning architecturally isolated from execution?

0
Model has direct execution authority

1
Basic interception but model can bypass

2
Policy engine exists but is advisory

3
Deterministic policy enforcement

4
Cryptographic capability grants and authority separation

5
Leading isolation: formally verified authority boundary, zero model execution authority, and independent verification

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 Security Architecture Suite (MSAS)

Traditional security benchmarks evaluate configuration. The MSAS evaluates architectural resilience.

5.1 Suite Composition

5.1.1 MSAS-ThreatModel

- **Adversarial Resilience:** 100+ tests simulating compromised components to verify blast radius containment.
- **Fail-Closed Verification:** 50+ tests killing policy engines to verify the system halts execution.

5.1.2 MSAS-Formal

- **Deadlock Detection:** 50+ TLA+ models verifying core state machines.
- **Blast Radius Simulation:** 100+ tests verifying compromised agents cannot escalate privileges.

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.