MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / CYBERSECURITY AND NETWORK SECURITY

Firewall Engineering, Policy Architecture, and Blast-Radius Standard

Defines firewall policy design, lifecycle, NAT, segmentation, change safety, and validation.

PERMANENT DOCUMENT ID
MYTHOS-SEC-0006

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
7 Atomic Criteria / 5 Families

PERMANENT PUBLICATION IDENTITY

Firewall Engineering, Policy Architecture, and Blast-Radius Standard

Defines firewall policy design, lifecycle, NAT, segmentation, change safety, and validation.

firewall-policy

DOCUMENT ID
MYTHOS-SEC-0006

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

7	5	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-0006 inside the cybersecurity and network security standard constellation

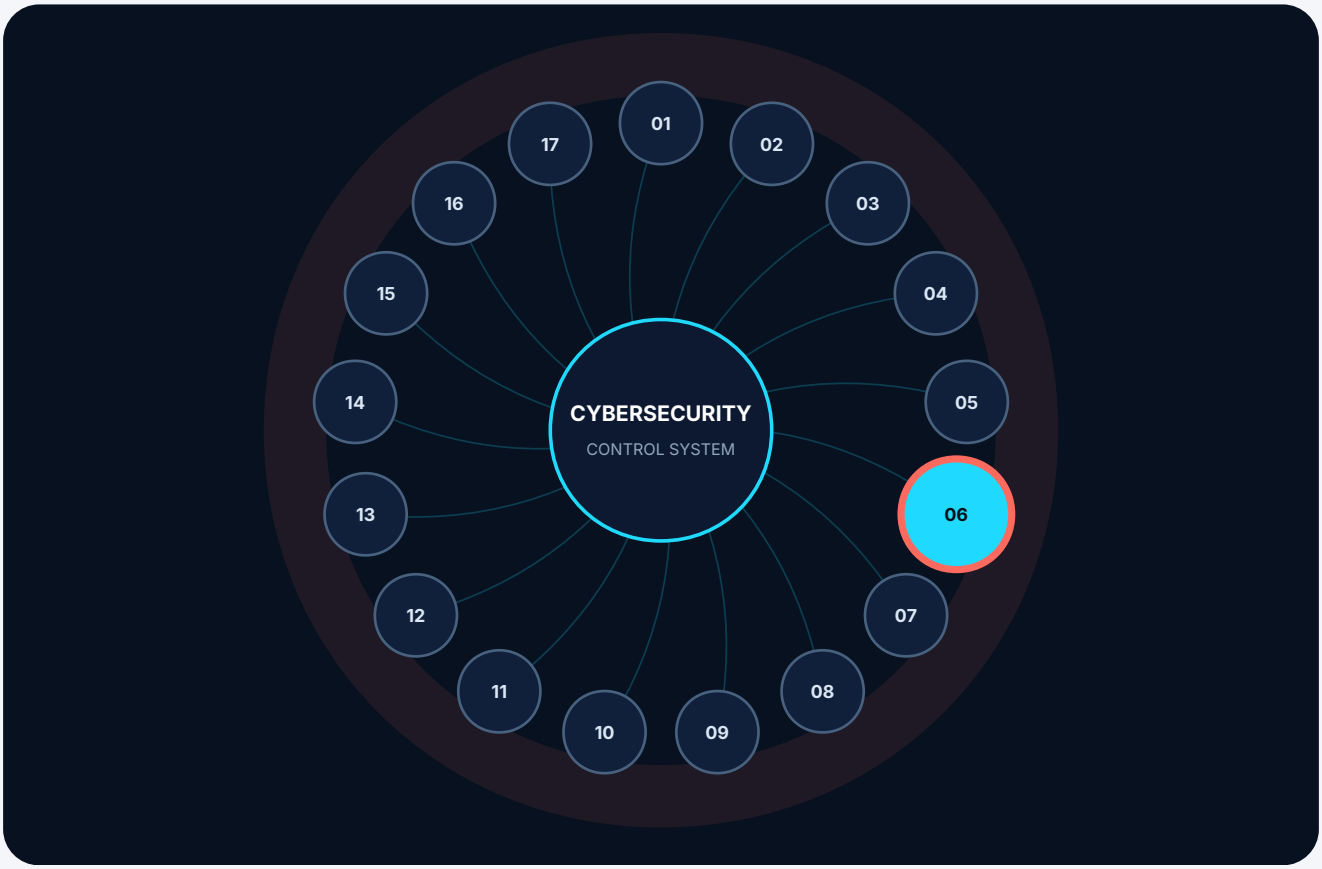

Connected assurance

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

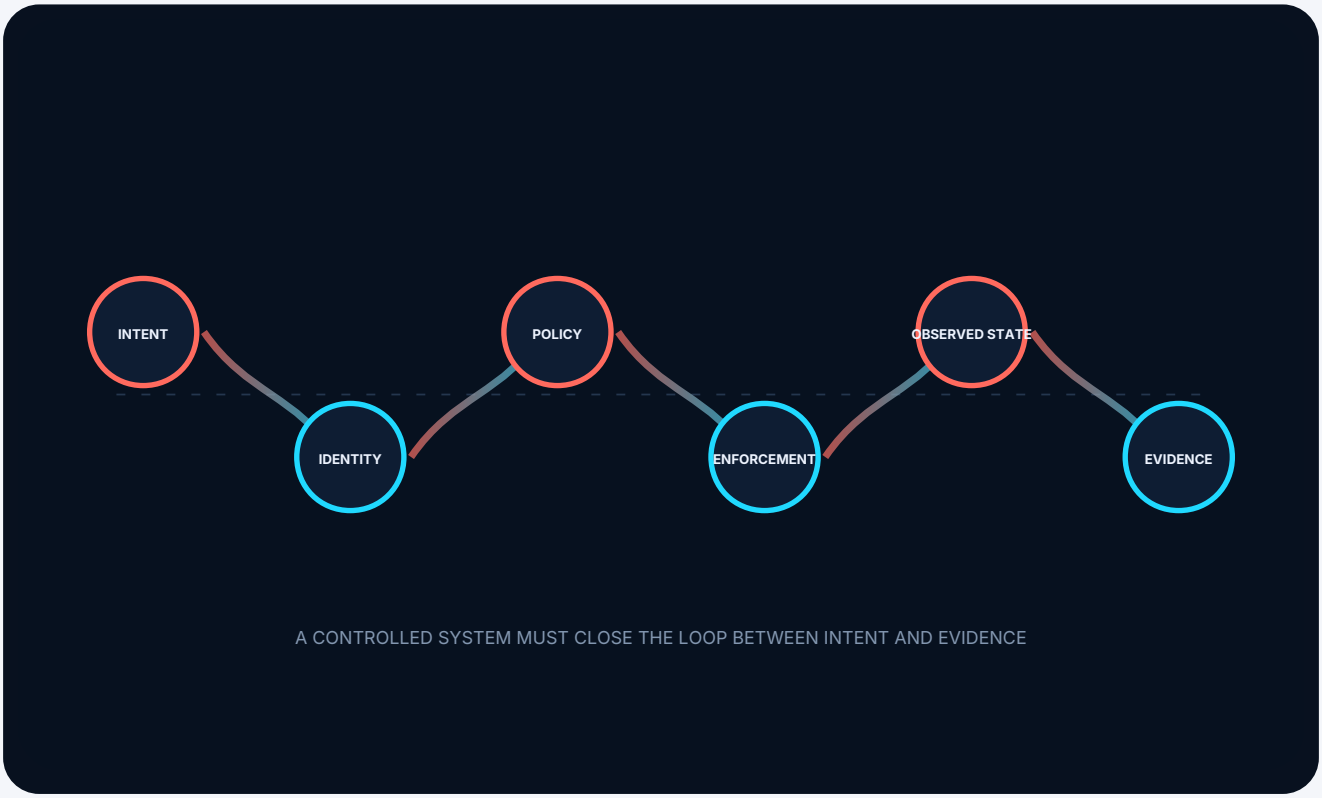

A CONTROLLED SYSTEM MUST CLOSE THE LOOP BETWEEN INTENT AND EVIDENCE

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-0006

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

0. Executive Mandate

Part I. Scope and Applicability

Part II. Definitions and Taxonomy

Part III. Evaluation Methodology

Evaluation system

4.1 - Pre-Change Simulation and Digital Twins

- 4.1.1 - Digital Twin Verification
- 4.1.2 - Reachability Verification

4.2 - Blast Radius Calculation and Enforcement

- 4.2.1 - Blast Radius Enforcement

4.3 - Policy-as-Code and Idempotency

- 4.3.1 - Idempotent Pushes
- 4.3.2 - Policy-as-Code

4.4 - Rule Lifecycle and Hygiene

- 4.4.1 - Shadowed Rule Detection

4.5 - AI Authority Isolation

- 4.5.1 - AI Rule Generation

Part V. The Mythos Firewall Suite (MFS)

ENGINEERING DOCTRINE

0. Executive Mandate

Traditional evaluation of firewall engineering is insufficient.

Organizations measure firewalls by throughput (terabits per second) and rule count (millions of lines). They do not measure whether the firewall's policy architecture is mathematically sound, whether rule changes are simulated before execution, or whether an AI agent can hallucinate a rule that opens the internal network to the internet.

This standard exists because:

1. **Throughput is Not Security:** A firewall that processes 10 Tbps but contains shadowed rules, overlaps, and logical contradictions is a liability. The primary metric of a firewall MUST be policy integrity and blast radius containment, not packet-per-second throughput.
2. **Manual Commits are Operational Suicide:** Relying on human operators to manually commit firewall rules across multi-vendor environments (Palo Alto, Fortinet, Cisco) guarantees drift, misconfiguration, and outages. Policy changes MUST be automated via deterministic, idempotent pipelines.
3. **AI Hallucination is a Critical Threat:** If an AI agent is allowed to propose firewall rules without formal verification, it will eventually hallucinate a rule that permits unauthorized traffic. AI proposals MUST be verified against a digital twin (e.g., Batfish) before they are ever applied.
4. **Blast Radius is the Only Metric That Matters:** A single misconfigured security rule can expose a production database to the internet. The system MUST calculate and enforce a maximum blast radius for every change, blocking any action that exceeds the threshold.

This document establishes the **Mythos Firewall Engineering Standard (MFES)**, a comprehensive, evidence-based methodology for evaluating the architecture, automation, and verification of security policy enforcement.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Firewall policy architecture (rule hierarchies, domains, device groups)
- Pre-change simulation and digital twin verification (Batfish, Forward Networks)
- Blast radius calculation and enforcement
- Idempotent network automation (desired-state vs. observed-state)
- AI-driven security policy generation and review
- Zero-trust segmentation and microsegmentation
- Multi-vendor firewall orchestration (PAN-OS, FortiOS, Cisco Secure Firewall)
- Rule lifecycle management (expiration, review, decomposition, shadowed rule detection)

1.2 Out of Scope

- General network automation (covered in MYTHOS-AI-0005)
- Zero-trust policy enforcement engines (covered in MYTHOS-SEC-0003)

1.3 Audience

- Network security engineers and architects designing policy fabrics
- Principal engineers selecting firewall automation tools
- Security teams evaluating network attack surfaces
- AI governance, risk, and compliance teams
- Standards bodies seeking a reference firewall engineering methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Firewall Dimensions

Dimension	Definition
Policy-as-Code	The management of firewall rules via declarative configuration files (HCL, YAML) rather than manual GUI clicks.
Blast Radius	The maximum potential impact (e.g., systems compromised, data exfiltrated) if a rule is misconfigured.
Digital Twin	A mathematical or emulated model of the network (e.g., Batfish) used to simulate the impact of a rule change.
Shadowed Rule	A rule that is never matched because a higher-priority rule encompasses its traffic, creating dead config.
Idempotent Push	A configuration change that can be applied repeatedly without introducing unintended side effects.

2.2 The Firewall Engineering Doctrine

> Throughput is not security. Manual commits are suicide. A rule without simulation is an outage. An AI that writes rules without verification 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; manual GUI management, no simulation
1	Basic automation but no simulation or idempotency
2	Functional but lacks blast radius enforcement
3	Solid production performance; policy-as-code, simulation, idempotency
4	Strong performance; full digital twin, AI-governed, blast radius limits
5	Best-in-class; sets the standard for firewall engineering

Weighting

Category	Weight	Rationale
Pre-Change Simulation & Digital Twins	25%	Untested changes cause outages
Blast Radius Calculation & Enforcement	20%	Unbounded changes expose the network
Policy-as-Code & Idempotency	15%	Manual pushes guarantee drift
Rule Lifecycle & Hygiene	15%	Stale/shadowed rules are attack surfaces
AI Authority Isolation	10%	AI must not directly commit rules
Zero-Trust & Microsegmentation	10%	Flat networks are indefensible
Operational Lifecycle	5%	Policies must be governed continuously

Total: 100%

A system MUST score at least 3.0 in Pre-Change Simulation and Blast Radius Calculation to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

7 atomic criteria across 5 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Pre-Change Simulation and Digital Twins

SOURCE WEIGHT 25%

4.1.1
Digital Twin Verification

4.1.2
Reachability Verification

MYTHOS-SEC-0006-EVAL-4.1.1

Digital Twin Verification

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Pre-Change Simulation and Digital Twins	**4.1.1**	**25%**	**Evidence-led**

Does the system simulate the network impact of a rule change before execution?

0
No simulation; executes directly in production

1
Basic dry-run mode (text diff)

2
Standard plan/diff mode showing textual changes

3
Graph-based blast radius analysis

4
Multi-domain simulation (e.g., Batfish for network reachability)

5
Leading simulation: full digital twin, cryptographic plan sealing, and rollback confidence scoring

ENGINEERING INTERPRETATION

This criterion evaluates whether digital twin 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.

MYTHOS-SEC-0006-EVAL-4.1.2

Reachability Verification

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Pre-Change Simulation and Digital Twins	4.1.2	25%	Evidence-led

Does the system verify that critical flows remain reachable after a change?

0
No reachability testing

1
Manual ping/traceroute after execution

2
Automated ping post-execution

3
Pre-execution simulation of reachability

4
Mathematical proof of reachability (Batfish)

5
Leading verification: formally verified reachability invariants, verified by adversarial suite

ENGINEERING INTERPRETATION

This criterion evaluates whether reachability 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.

4.2 / CAPABILITY FAMILY

Blast Radius Calculation and Enforcement

SOURCE WEIGHT 20%

4.2.1
Blast Radius Enforcement

MYTHOS-SEC-0006-EVAL-4.2.1

Blast Radius Enforcement

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Blast Radius Calculation and Enforcement	**4.2.1**	**20%**	**Evidence-led**

Is the blast radius of a change calculated and enforced before execution?

0
No blast radius calculation

1
Basic risk classification (low/medium/high)

2
Automated risk classification based on rule context

3
Automated blast radius graph analysis

4
Policy gates blocking changes that exceed blast radius limits

5
Leading enforcement: formally verified blast radius 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

- 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

Policy-as-Code and Idempotency

SOURCE WEIGHT 15%

4.3.1
Idempotent Pushes

4.3.2
Policy-as-Code

MYTHOS-SEC-0006-EVAL-4.3.1

Idempotent Pushes

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Policy-as-Code and Idempotency	4.3.1	15%	Evidence-led

Are firewall changes idempotent?

0
Destructive on retry

1
Partially idempotent

2
Idempotent for standard rules

3
Fully idempotent with idempotency keys

4
Idempotency verified by post-checks

5
Leading idempotency: mathematically proven with automatic compensation on failure

ENGINEERING INTERPRETATION

This criterion evaluates whether idempotent pushes 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

- rulebase export, ownership matrix, hit counts, recertification records, and flow logs
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test allow, deny, shadowing, expiration, asymmetric flow, and failover scenarios
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- orphan rules, broad services, hidden shadowing, stale objects, and weak ownership
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- owned-rule coverage
- stale-rule count
- policy verification pass 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-0006-EVAL-4.3.2

Policy-as-Code

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Policy-as-Code and Idempotency	4.3.2	15%	Evidence-led

Is firewall configuration managed as declarative code?

0
Manual GUI management

1
Basic CLI scripts

2
Version-controlled configuration files

3
Declarative IaC (Terraform/OpenTofu)

4
Policy-as-Code with formal logic verification

5
Leading PoC: formally verified, continuously reconciled, and AI-governed

ENGINEERING INTERPRETATION

This criterion evaluates whether policy-as-code 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

- rulebase export, ownership matrix, hit counts, recertification records, and flow logs
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test allow, deny, shadowing, expiration, asymmetric flow, and failover scenarios
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- orphan rules, broad services, hidden shadowing, stale objects, and weak ownership
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- owned-rule coverage
- stale-rule count
- policy verification pass 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

Rule Lifecycle and Hygiene

SOURCE WEIGHT 15%

4.4.1
Shadowed Rule Detection

MYTHOS-SEC-0006-EVAL-4.4.1

Shadowed Rule Detection

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Rule Lifecycle and Hygiene	4.4.1	15%	Evidence-led

Does the system automatically detect and flag shadowed rules?

0
No detection

1
Manual review

2
Basic automated scanning

3
Real-time detection on commit

4
Automated remediation suggestions

5
Leading hygiene: continuous formal verification, automatic decomposition, and expiration enforcement

ENGINEERING INTERPRETATION

This criterion evaluates whether shadowed rule detection 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

- telemetry coverage, detection source, test fixtures, version history, alerts, and case outcomes
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. replay benign and malicious fixtures; measure fidelity and operational response
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE – INTERVIEW – TEST

COMMON FAILURE PATTERNS

- untested rules, missing telemetry, alert-only metrics, and undocumented suppression
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- true-positive rate
- false-positive burden
- coverage by technique
- mean time to contain
- 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.

4.5 / CAPABILITY FAMILY

AI Authority Isolation

SOURCE WEIGHT 10%

4.5.1
AI Rule Generation

MYTHOS-SEC-0006-EVAL-4.5.1

AI Rule Generation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
AI Authority Isolation	4.5.1	10%	Evidence-led

Can an AI agent propose and execute firewall rules without human verification?

0
AI has direct commit access

1
AI proposes, but changes are auto-applied

2
AI proposes, human reviews text, human applies

3
AI proposes, digital twin simulates, human reviews simulation, human applies

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

5
Leading isolation: AI proposes, twin simulates, policy gates, dual-control, cryptographic attestation

ENGINEERING INTERPRETATION

This criterion evaluates whether ai rule generation 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

- rulebase export, ownership matrix, hit counts, recertification records, and flow logs
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test allow, deny, shadowing, expiration, asymmetric flow, and failover scenarios
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- orphan rules, broad services, hidden shadowing, stale objects, and weak ownership
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- owned-rule coverage
- stale-rule count
- policy verification pass 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.

LIFECYCLE AND ASSURANCE

Part V. The Mythos Firewall Suite (MFS)

Traditional firewall benchmarks measure throughput. The MFS evaluates policy integrity, simulation, and blast radius containment.

5.1 Suite Composition

5.1.1 MFS-Simulation

- **Blast Radius:** 100+ tests verifying the system blocks changes that exceed defined blast radius limits.
- **Reachability:** 100+ tests verifying critical flows remain reachable after simulated changes.

5.1.2 MFS-Hygiene

- **Shadow Detection:** 50+ tests verifying the system detects and flags shadowed/overlapping rules.
- **Idempotency:** 50+ tests verifying retries do not duplicate or corrupt configurations.

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.