MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / CYBERSECURITY AND NETWORK SECURITY

Network Security Architecture and Microsegmentation Standard

Defines secure zones, identity-aware policy, workload isolation, enforcement, and verification.

Response, and Assurance Evidence

PERMANENT DOCUMENT ID
MYTHOS-SEC-0007

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
7 Atomic Criteria / 5 Families

PERMANENT PUBLICATION IDENTITY

Network Security Architecture and Microsegmentation Standard

Defines secure zones, identity-aware policy, workload isolation, enforcement, and verification.

DOCUMENT ID
MYTHOS-SEC-0007

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-0007 inside the cybersecurity and network security standard constellation

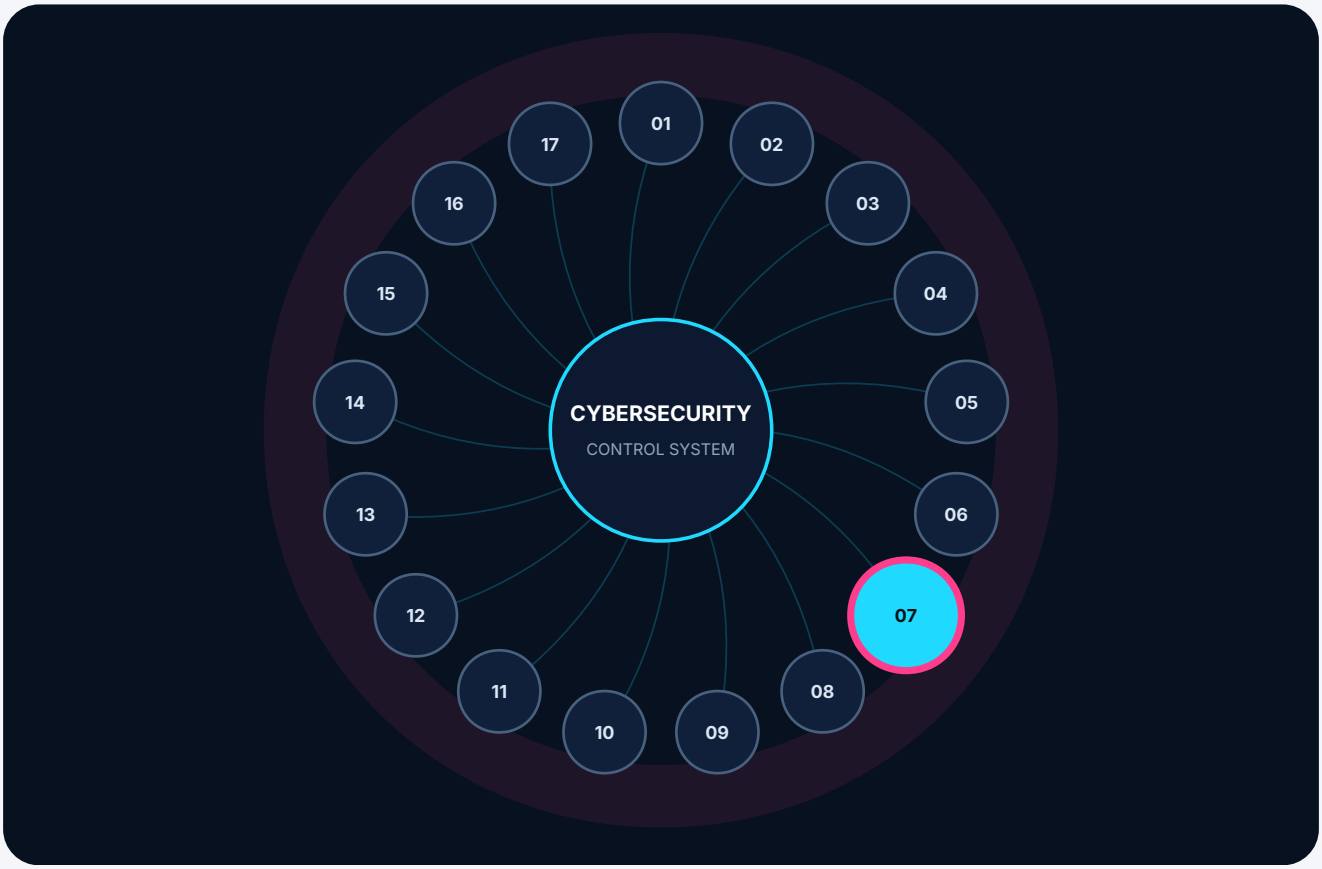

Connected assurance

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

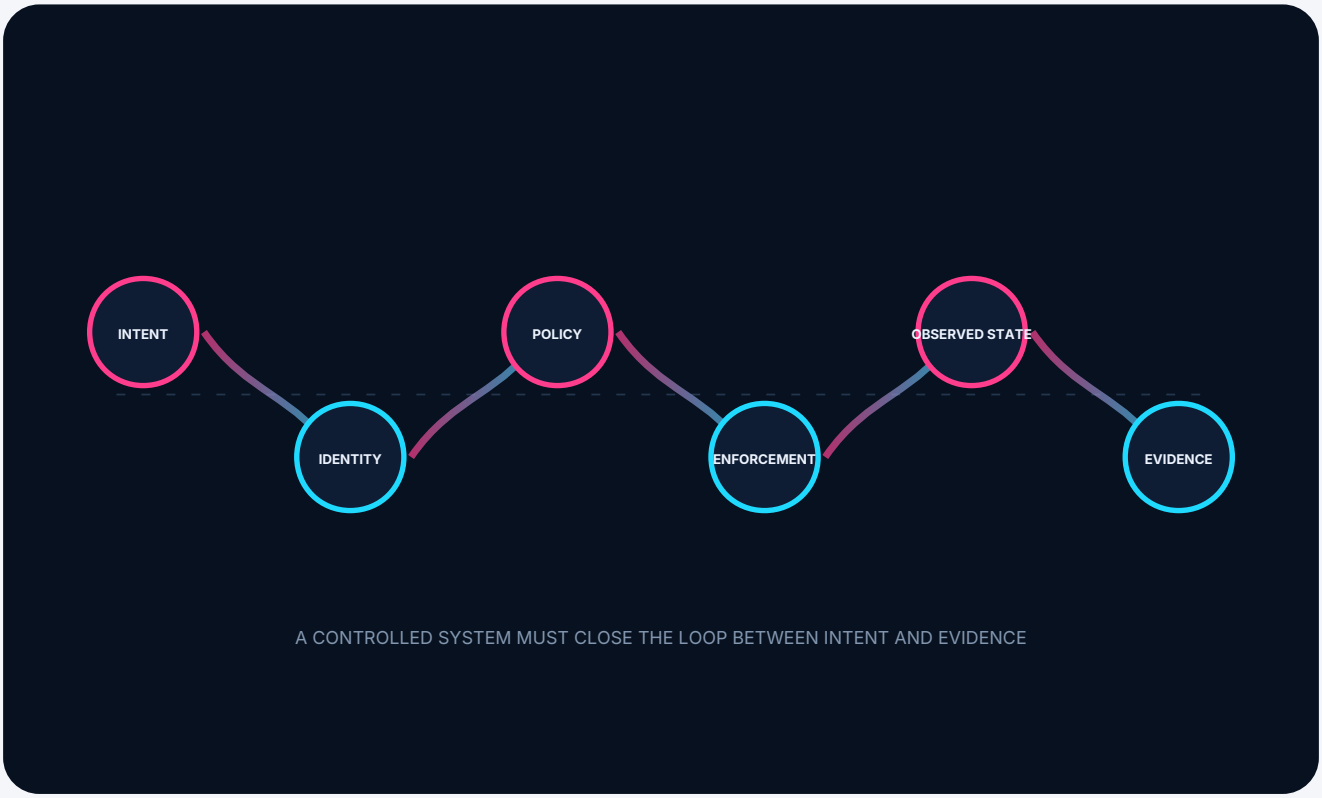

State separation

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-0007

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 network security architecture is insufficient.

Organizations measure security by the thickness of their perimeter (North-South firewalls) and the speed of their intrusion detection. They do not measure whether an attacker - or a hallucinating AI agent - can move laterally (East-West) across the internal network in seconds. The "castle-and-moat" architecture is dead. Once the perimeter is breached, the internal network is flat, undefended, and entirely compromised.

This standard exists because:

1. **Perimeters are Obsolete:** Cloud adoption, remote work, and AI agents have dissolved the network boundary. Security MUST be attached to the workload, not the perimeter.
2. **Lateral Movement is the #1 Threat:** Attackers (and rogue AI agents) do not break through firewalls; they steal credentials and move laterally. Microsegmentation - restricting communication between workloads - is the only architectural defense.
3. **AI Agents are High-Risk Workloads:** AI agents execute code, access secrets, and make decisions. They MUST be microsegmented. An agent compromised by prompt injection must not be able to access the production database.
4. **Static Rules Do Not Scale:** Manual firewall rules for microsegmentation are impossible to manage at scale. Policy MUST be identity-aware, dynamic, and driven by workload attributes (e.g., SPIFFE), not static IP addresses.
5. **Segmentation Must Be Proven:** "We think we are segmented" is not acceptable. The segmentation topology MUST be mathematically modeled and formally verified to ensure no unauthorized paths exist.

This document establishes the **Mythos Network Security & Microsegmentation Standard (MNSMS)**, a comprehensive, evidence-based methodology for evaluating zero-trust network architectures and lateral movement containment.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Microsegmentation architectures (e.g., Cilium, Istio, Cisco ACI)
- Zero-Trust Network Access (ZTNA) and Software-Defined Perimeters (SDP)
- Identity-aware policy enforcement (SPIFFE/SPIRE)
- East-West traffic control and lateral movement prevention
- L7 (application layer) policy enforcement
- Formal verification of segmentation graphs
- AI-driven network policy generation and review
- eBPF-based runtime observability and enforcement

1.2 Out of Scope

- General network architecture (covered in MYTHOS-NET-0001)
- Firewall engineering specifics (covered in MYTHOS-SEC-0006)

1.3 Audience

- Network security engineers and architects designing zero-trust fabrics
- Principal engineers selecting microsegmentation platforms
- Security teams evaluating lateral movement risks
- AI governance, risk, and compliance teams
- Standards bodies seeking a reference microsegmentation methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Architecture Dimensions

Dimension	Definition
Microsegmentation	The practice of isolating workloads (e.g., VMs, containers, AI agents) from one another using fine-grained, identity-aware policies, rather than relying on network perimeters.
East-West Traffic	Communication between internal workloads within a data center or cloud environment.
North-South Traffic	Communication entering or exiting the network perimeter.
ZTNA	Zero-Trust Network Access. Replaces VPNs with identity-aware, per-application access.
eBPF	Extended Berkeley Packet Filter. A kernel technology enabling high-performance, runtime observability and enforcement of network policies without kernel modifications.
Formal Verification	The use of mathematical logic to prove that a network topology contains no unauthorized paths.

2.2 The Network Security Doctrine

> Perimeters are obsolete. Lateral movement is the threat. IPs are not identity. A network that cannot prove its isolation is compromised.

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 network, perimeter-only defense
1	Basic VLAN segmentation but no microsegmentation
2	Functional but lacks identity-awareness or L7 enforcement
3	Solid production performance; microsegmented, identity-aware
4	Strong performance; L7 enforcement, eBPF observability, formal verification
5	Best-in-class; sets the standard for zero-trust architecture

Weighting

Category	Weight	Rationale
East-West Enforcement	25%	Lateral movement is the primary attack vector
Identity-Aware Policy	20%	IPs are easily spoofed; identity is mandatory
Formal Verification	15%	Unverified topology is assumed compromised
L7 Policy Enforcement	10%	L3/L4 is insufficient for modern apps
AI Agent Isolation	10%	AI agents are high-risk workloads
eBPF Observability	10%	Unobservable networks are undefendable
Operational Lifecycle	10%	Policies must be governed continuously

Total: 100%

A system MUST score at least 3.0 in East-West Enforcement and Identity-Aware Policy 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

East-West Enforcement

SOURCE WEIGHT 25%

4.1.1
Lateral Movement Prevention

4.1.2
Workload Isolation

MYTHOS-SEC-0007-EVAL-4.1.1

Lateral Movement Prevention

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
East-West Enforcement	4.1.1	25%	Evidence-led

> Does the system enforce strict default-deny for all internal traffic?

0
Flat network; any host can talk to any host

1
Basic VLAN segmentation; broad trust zones

2
Host-based firewalls but no centralized control

3
Centralized microsegmentation with default-deny

4
Real-time enforcement with eBPF/Cilium

5
Leading enforcement: formally verified default-deny, zero implicit trust, verified by lateral movement suite

ENGINEERING INTERPRETATION

This criterion evaluates whether lateral movement prevention 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-0007-EVAL-4.1.2

Workload Isolation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
East-West Enforcement	4.1.2	25%	Evidence-led

Are workloads (and AI agents) isolated from each other by default?

0
All workloads share a network namespace

1
Basic container networking isolation

2
Network policies exist but are not default-deny

3
Default-deny microsegmentation for all workloads

4
Dynamic isolation based on workload attributes (e.g., environment, sensitivity)

5
Leading isolation: formally verified boundaries, dynamic attribute-based isolation, and automatic quarantine

ENGINEERING INTERPRETATION

This criterion evaluates whether workload 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.

4.2 / CAPABILITY FAMILY

Identity-Aware Policy

SOURCE WEIGHT 20%

4.2.1
Cryptographic Identity

4.2.2
Dynamic Policy Evaluation

MYTHOS-SEC-0007-EVAL-4.2.1

Cryptographic Identity

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Identity-Aware Policy	4.2.1	20%	Evidence-led

Does the system use cryptographic identity (e.g., SPIFFE) instead of IPs?

0
Policies based on static IP addresses

1
Basic service accounts but no cryptographic identity

2
mTLS exists but policies still reference IPs

3
SPIFFE/SPIRE or equivalent for all workloads

4
Policies are 100% identity-aware; IPs are irrelevant

5
Leading identity: formally verified identity boundaries, cryptographic attestation, and zero IP-based rules

ENGINEERING INTERPRETATION

This criterion evaluates whether cryptographic identity 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.

MYTHOS-SEC-0007-EVAL-4.2.2

Dynamic Policy Evaluation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Identity-Aware Policy	4.2.2	20%	Evidence-led

Are policies evaluated dynamically based on real-time context?

0
Static rules

1
Manual rule updates

2
Automated updates but slow propagation

3
Real-time policy evaluation

4
Context-aware (e.g., threat intel, vulnerability status)

5
Leading evaluation: real-time, formally verified, and AI-governed context integration

ENGINEERING INTERPRETATION

This criterion evaluates whether dynamic policy evaluation 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
Topology Verification

MYTHOS-SEC-0007-EVAL-4.3.1

Topology Verification

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

Is the network topology mathematically proven to have no unauthorized paths?

0
No verification

1
Manual topology review

2
Automated topology scanning

3
Graph-based reachability analysis

4
Formal verification of reachability invariants

5
Leading verification: continuously verified topology, cryptographic sealing of policies, and mathematical proof of isolation

ENGINEERING INTERPRETATION

This criterion evaluates whether topology 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

L7 Policy Enforcement

SOURCE WEIGHT 10%

4.4.1
Application-Layer Control

MYTHOS-SEC-0007-EVAL-4.4.1

Application-Layer Control

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
L7 Policy Enforcement	4.4.1	10%	Evidence-led

Does the system enforce policy at Layer 7 (HTTP/gRPC)?

0
L3/L4 only

1
Basic URL filtering

2
L7 proxies but no policy enforcement

3
L7 policy enforcement (e.g., Istio AuthorizationPolicy)

4
L7 enforcement + eBPF bypass prevention

5
Leading L7: formally verified L7 policies, eBPF enforcement, and zero L3 bypass

ENGINEERING INTERPRETATION

This criterion evaluates whether application-layer control 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

AI Agent Isolation

SOURCE WEIGHT 10%

4.5.1
Agent Segmentation

MYTHOS-SEC-0007-EVAL-4.5.1

Agent Segmentation

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

Are AI agents microsegmented from production infrastructure?

0
AI agents have broad network access

1
Basic network ACLs for agents

2
Agents are in a separate subnet

3
Agents have SPIFFE identity and default-deny

4
Agents are dynamically isolated based on their task

5
Leading isolation: formally verified agent boundaries, zero lateral movement capability, and automatic quarantine on anomalous traffic

ENGINEERING INTERPRETATION

This criterion evaluates whether agent segmentation 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

- zone and identity model, policy graph, enforcement exports, and east-west telemetry
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test intended and prohibited flows across normal and degraded conditions
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- VLAN-only assumptions, transitive access, broad service accounts, and undocumented bypass
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- policy coverage
- prohibited-flow success rate
- exception age
- 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 Network Security Suite (MNSS)

Traditional network benchmarks measure throughput. The MNSS evaluates lateral movement containment and formal verification.

5.1 Suite Composition

5.1.1 MNSS-Lateral

- **Lateral Movement:** 100+ tests attempting to traverse from an untrusted workload to a production database.
- **Identity Spoofing:** 50+ tests attempting to bypass policy by spoofing IP/identity.

5.1.2 MNSS-Verification

- **Topology Integrity:** 50+ tests verifying the network graph matches the formal specification.
- **Policy Drift:** 50+ tests verifying policies have not drifted from the desired state.

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.