ENGINEERING STANDARDS LIBRARY / PORTFOLIO EDITION

Network Engineering Standards Portfolio

Permanent collection of MYTHOS-NET-0001 through MYTHOS-NET-0011.

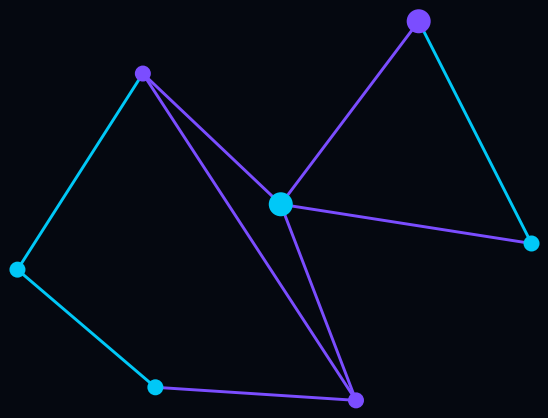

PERMANENT DOCUMENT ID

MYTHOS-NET-PORT-0001

Portfolio contents

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

Network Architecture, Design, and Topology Standard

Defines architecture views, topology, requirements, resilience, scale, documentation, and acceptance.

PERMANENT DOCUMENT ID
MYTHOS-NET-0001

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
10 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

Network Architecture, Design, and Topology Standard

Defines architecture views, topology, requirements, resilience, scale, documentation, and acceptance.

DOCUMENT ID
MYTHOS-NET-0001

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

10	6	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-NET-0001 inside the network engineering standard constellation

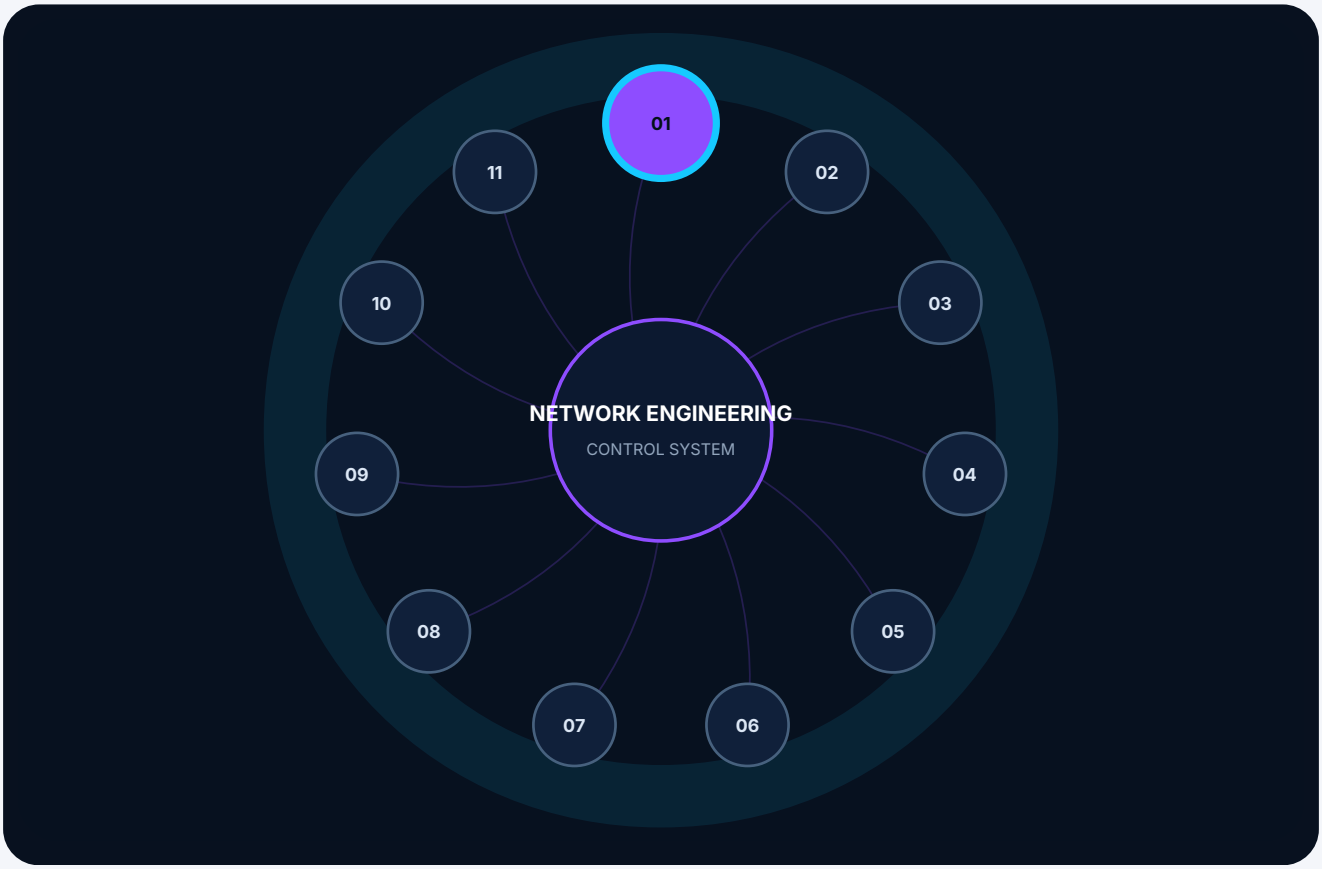

Connected assurance

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

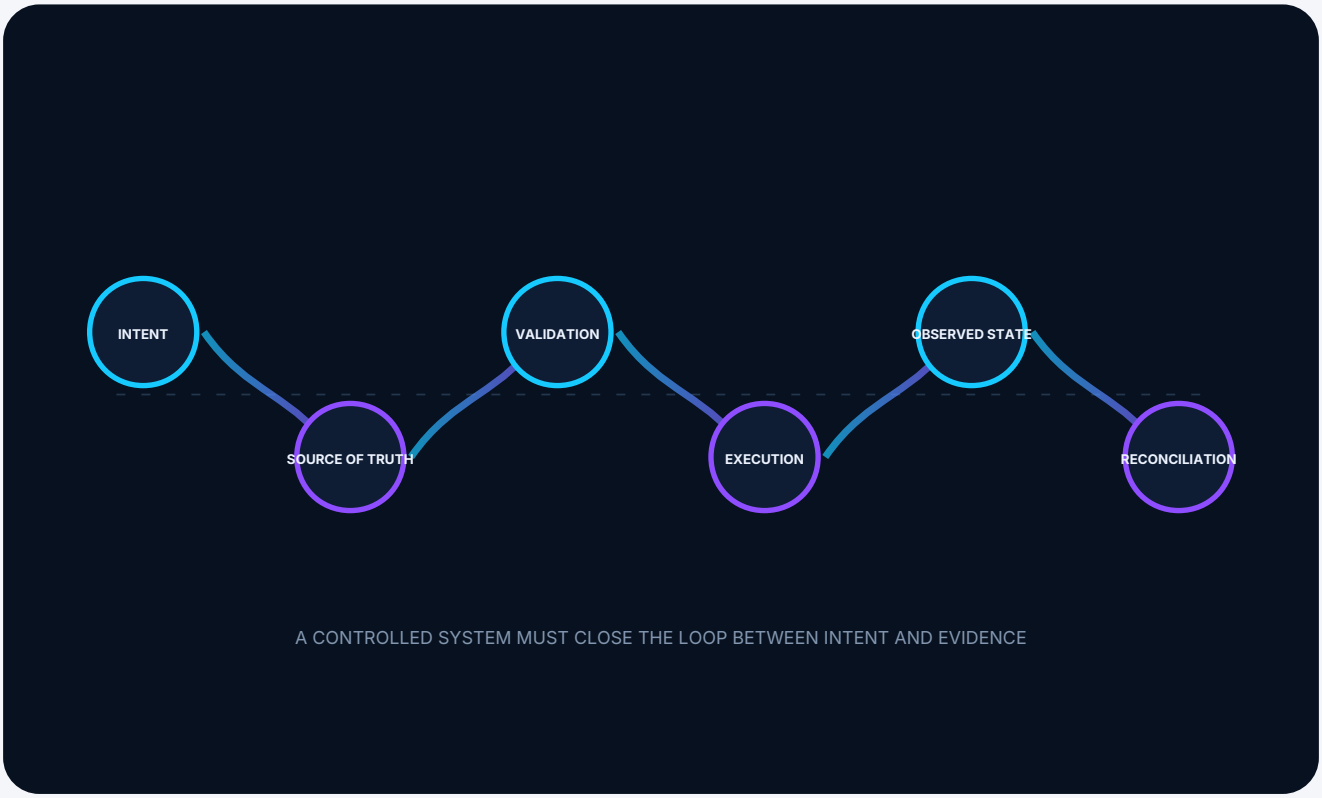

State separation

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-NET-0001

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 architecture is insufficient.

Organizations measure network health by whether the monitoring dashboard is green and whether the CLI still accepts commands. They document their architecture in static Visio diagrams that are obsolete the moment they are saved. They manage their state via artisanal CLI commands passed down through tribal knowledge. This is not engineering; it is digital folklore.

This standard exists because:

1. **A Diagram is Not Architecture:** A Visio drawing is a picture, not a source of truth. It cannot be tested, verified, or reconciled. Network architecture MUST be defined as executable, version-controlled intent.
2. **CLI is Not Intent:** Typing `router bgp 65001` into a terminal is an implementation detail, not an architectural specification. Intent-Based Networking (IBN) - where the desired state is declared and the system determines the configuration - is the preferred sustainable architectural paradigm.
3. **Reachability Must Be Proven, Not Guessed:** "I think the firewall allows that flow" is an unacceptable operational posture. Network architectures MUST be mathematically verifiable. Tools like Batfish must prove reachability before a single packet is forwarded.
4. **ClickOps is high-risk architectural constraint:** Manual configuration via vendor GUIs guarantees drift, inconsistency, and outages. The network MUST be governed by Network-as-Code (NaC) and GitOps pipelines.

This document establishes the **Mythos Network Architecture Standard (MNAS)**, a comprehensive, evidence-based methodology for evaluating network architectures against intent-alignment, formal verification, and automation readiness.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Network architecture topologies (Spine-Leaf, Fat-Tree, Hub-and-Spoke, Mesh)
- Intent-Based Networking (IBN) and Network-as-Code (NaC) paradigms
- Source-of-Truth (SoT) alignment (NetBox/Nautobot as authority)
- Formal verification of routing and reachability (Batfish, Forward Networks)
- Underlay and Overlay design (VXLAN/EVPN, SRv6, Segment Routing)
- Multi-vendor abstraction and normalization (OpenConfig, gNMI, YANG)
- Infrastructure-as-Code for network (Terraform, Ansible, Golden Config)
- Topology resilience, blast radius, and failure domain analysis

1.2 Out of Scope

- Specific routing protocol configurations (covered in MYTHOS-NET-0002)
- Digital twin simulation execution (covered in MYTHOS-NET-0003)
- Zero-trust network enforcement (covered in MYTHOS-SEC-0007)

1.3 Audience

- Network architects and principal engineers designing fabric topologies
- Platform engineers building Network-as-Code pipelines
- Security teams evaluating network segmentation and blast radius
- AI governance, risk, and compliance teams
- Standards bodies seeking a reference network architecture methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Architecture Dimensions

Dimension	Definition
Intent-Based Networking (IBN)	A network management paradigm where the operator defines the *what* (intent) and the system determines the *how* (configuration), continuously verifying that the observed state matches the intent.
Network-as-Code (NaC)	The practice of managing network state via version-controlled, declarative code rather than imperative CLI commands.
Source-of-Truth (SoT)	The authoritative, immutable record of network intent, IP allocations, and topology (e.g., NetBox, Nautobot).
Formal Verification	The mathematical proof that a network's routing and access control tables produce the intended reachability.
Topology Drift	The divergence between the intended network architecture (SoT) and the observed state of the network devices.

2.2 The Network Architecture Doctrine

> A diagram is documentation, not truth. A CLI is implementation, not intent. A network that cannot prove its reachability is compromised. A change that is not in Git does not exist.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; CLI-driven, undocumented, unverified
1	Documented but manual, no automation or verification
2	Functional but lacks intent-alignment or formal verification
3	Solid production performance; NaC, basic verification, SoT-aligned
4	Strong performance; IBN, formal verification, SRv6/overlay abstraction
5	Best-in-class; sets the standard for autonomous, verified network architecture

Weighting

Category	Weight	Rationale
Intent-Topology Alignment & SoT	20%	Observed state must match intent or the architecture is a lie
Formal Verification & Reachability	20%	Unverified reachability is an outage waiting to happen
Network-as-Code & GitOps	15%	CLI/ClickOps guarantee drift
Topology Resilience & Blast Radius	15%	Single points of failure cause catastrophic outages
Multi-Vendor Abstraction	15%	Vendor lock-in is high-risk architectural constraint
Underlay/Overlay & SRv6 Design	10%	Modern fabrics require modern transport
Operational Lifecycle	5%	Architecture must be governed continuously

Total: 100%

A system MUST score at least 3.0 in Intent-Topology Alignment and Formal Verification to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

10 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Intent-Topology Alignment and Source-of-Truth

SOURCE WEIGHT 20%

4.1.1
State Reconciliation

4.1.2
Intent Authority

MYTHOS-NET-0001-EVAL-4.1.1

State Reconciliation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Intent-Topology Alignment and Source-of-Truth	**4.1.1**	**20%**	**Evidence-led**

Does the observed network state continuously reconcile with the Source-of-Truth?

0
No Source-of-Truth; state is whatever is on the devices

1
Static SoT (Spreadsheet/Visio) but never reconciled

2
SoT exists (NetBox) but reconciliation is manual

3
Automated periodic reconciliation with drift alerts

4
Continuous reconciliation with automated drift remediation (desired-state push)

5
Leading alignment: formally verified continuous reconciliation, policy-defined drift tolerance, and cryptographic attestation of state

ENGINEERING INTERPRETATION

This criterion evaluates whether state reconciliation 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

- source-of-truth objects, Git history, observed-state snapshots, and reconciliation evidence
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. introduce controlled drift and verify detection, adjudication, and restoration
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- multiple authorities, silent manual changes, stale inventory, and destructive auto-remediation
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- drift detection time
- reconciliation success
- unowned object count
- 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-NET-0001-EVAL-4.1.2

Intent Authority

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Intent-Topology Alignment and Source-of-Truth	4.1.2	20%	Evidence-led

Is the SoT the designated authority for network intent?

0
Devices are the source of truth; changes made directly

1
SoT exists but out-of-band changes are common

2
SoT is authoritative but not enforced (drift accepted)

3
SoT is enforced; out-of-band changes trigger alerts and are overwritten

4
SoT is enforced via GitOps; all changes require PR and CI/CD pipeline

5
Leading authority: immutable SoT, zero out-of-band change capability, cryptographic signing of intent

ENGINEERING INTERPRETATION

This criterion evaluates whether intent authority 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

- source-of-truth objects, Git history, observed-state snapshots, and reconciliation evidence
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. introduce controlled drift and verify detection, adjudication, and restoration
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- multiple authorities, silent manual changes, stale inventory, and destructive auto-remediation
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- drift detection time
- reconciliation success
- unowned object count
- 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.

4.2 / CAPABILITY FAMILY

Formal Verification and Reachability

SOURCE WEIGHT 20%

4.2.1
Reachability Proof

4.2.2
Intent Verification

MYTHOS-NET-0001-EVAL-4.2.1

Reachability Proof

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Formal Verification and Reachability	**4.2.1**	**20%**	**Evidence-led**

Can the architecture mathematically prove its reachability?

0
No verification; reachability tested via ping/traceroute in production

1
Manual peer review of configurations

2
Automated syntax/semantic validation (linting)

3
Pre-deployment simulation (Batfish) for major changes

4
Continuous formal verification of all routing and ACL tables

5
Leading verification: mathematically proven reachability invariants, continuously verified in CI/CD, and cryptographically sealed topology maps

ENGINEERING INTERPRETATION

This criterion evaluates whether reachability proof 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-NET-0001-EVAL-4.2.2

Intent Verification

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Formal Verification and Reachability	**4.2.2**	**20%**	**Evidence-led**

Does the system verify that the deployed configuration actually fulfills the architectural intent?

0
No verification; hope the config works

1
Post-deployment manual testing

2
Automated post-deployment ping/ connectivity checks

3
Automated configuration compliance (Greenbone/NetMRI)

4
Formal intent verification (Batfish assertion: "Can A reach B on port 443?")

5
Leading verification: continuous intent assertion, automated rollback on violation, and mathematical proof of policy compliance

ENGINEERING INTERPRETATION

This criterion evaluates whether intent 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.

4.3 / CAPABILITY FAMILY

Network-as-Code and GitOps

SOURCE WEIGHT 15%

4.3.1
Configuration Generation

4.3.2
Change Control

MYTHOS-NET-0001-EVAL-4.3.1

Configuration Generation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Network-as-Code and GitOps	**4.3.1**	**15%**	**Evidence-led**

How is network configuration generated and deployed?

0
Manual CLI commands

1
Copy-paste templates

2
Scripts generating configs but not applied via pipeline

3
Infrastructure-as-Code (Terraform/ Ansible) via CI/CD

4
GitOps pull-model (ArgoCD/Flux reconciling network state)

5
Leading NaC: intent-driven generation, GitOps reconciliation, automated diff, and zero human CLI access

ENGINEERING INTERPRETATION

This criterion evaluates whether configuration 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

- 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-NET-0001-EVAL-4.3.2

Change Control

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

Are all network changes version-controlled and peer-reviewed?

0
No version control; changes made ad-hoc

1
Backups saved to a shared drive

2
Git repository exists but not used for all changes

3
All changes require PR and peer review before merge

4
PR triggers CI pipeline (lint, simulate, verify) before merge

5
Leading control: PR, CI verification, formal proof, automated deploy, and cryptographic audit trail

ENGINEERING INTERPRETATION

This criterion evaluates whether change 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.

4.4 / CAPABILITY FAMILY

Topology Resilience and Blast Radius

SOURCE WEIGHT 15%

4.4.1
Failure Domain Isolation

4.4.2
Redundancy Verification

MYTHOS-NET-0001-EVAL-4.4.1

Failure Domain Isolation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Topology Resilience and Blast Radius	**4.4.1**	**15%**	**Evidence-led**

Are failure domains minimized and isolated?

0
Flat network; single failure domain

1
Basic hierarchy but no fault isolation

2
Modular design (Spine-Leaf) but shared control plane

3
Isolated failure domains with independent control planes

4
Blast radius bounded by topology; automated failover verified by twin

5
Leading resilience: formally verified blast radius limits, sub-second convergence, and self-healing paths

ENGINEERING INTERPRETATION

This criterion evaluates whether failure domain 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.

MYTHOS-NET-0001-EVAL-4.4.2

Redundancy Verification

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Topology Resilience and Blast Radius	**4.4.2**	**15%**	**Evidence-led**

Is redundancy mathematically verified, not just designed?

0
No redundancy

1
Dual hardware but untested failover

2
Manual failover testing annually

3
Automated failover testing (chaos engineering)

4
Continuous verification of redundant paths via digital twin

5
Leading verification: formally proven N+1/N+2 redundancy, continuous path validation, and zero single points of failure

ENGINEERING INTERPRETATION

This criterion evaluates whether redundancy 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.

4.5 / CAPABILITY FAMILY

Multi-Vendor Abstraction

SOURCE WEIGHT 15%

4.5.1
Vendor Neutrality

MYTHOS-NET-0001-EVAL-4.5.1

Vendor Neutrality

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Multi-Vendor Abstraction	**4.5.1**	**15%**	**Evidence-led**

Is the architecture abstracted from vendor-specific CLI/dependencies?

0
Single vendor lock-in; proprietary CLI only

1
Multi-vendor but separate silos per vendor

2
Multi-vendor with standardized data models (OpenConfig/YANG)

3
Unified abstraction layer (gNMI/ RESTCONF) over multi-vendor

4
Intent-driven abstraction; vendor config is a derived artifact

5
Leading abstraction: trait-based interfaces, zero vendor-specific logic in intent, automated conformance testing

ENGINEERING INTERPRETATION

This criterion evaluates whether vendor neutrality 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.

4.6 / CAPABILITY FAMILY

Underlay/Overlay and SRv6 Design

SOURCE WEIGHT 10%

4.6.1
Transport Architecture

MYTHOS-NET-0001-EVAL-4.6.1

Transport Architecture

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Underlay/Overlay and SRv6 Design	4.6.1	10%	Evidence-led

Does the architecture use modern, programmable transport?

0
Legacy VLAN-based flat transport

1
MPLS/LDP or VXLAN with manual BGP EVPN

2
Automated VXLAN/EVPN fabric

3
SR-MPLS or SRv6 underlay with automated provisioning

4
SRv6 with programmable SIDs (traffic engineering, SLA)

5
Leading transport: SRv6, automated TE, formally verified path computation, and AI-driven optimization

ENGINEERING INTERPRETATION

This criterion evaluates whether transport architecture 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 Network Architecture Suite (MNAS)

Traditional network benchmarks measure convergence time. The MNAS evaluates intent alignment, formal verification, and resilience.

5.1 Suite Composition

5.1.1 MNAS-Intent

- **Drift Detection:** 100+ tests injecting out-of-band changes to verify they are detected and overwritten.
- **SoT Authority:** 50+ tests verifying no device configuration exists that is not derived from the SoT.

5.1.2 MNAS-Verification

- **Reachability Assertions:** 100+ Batfish assertions verifying critical flows (e.g., "Web tier cannot reach DB tier on port 22").
- **Failover Verification:** 50+ tests simulating link/node failures in the digital twin to verify path redundancy.

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.

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

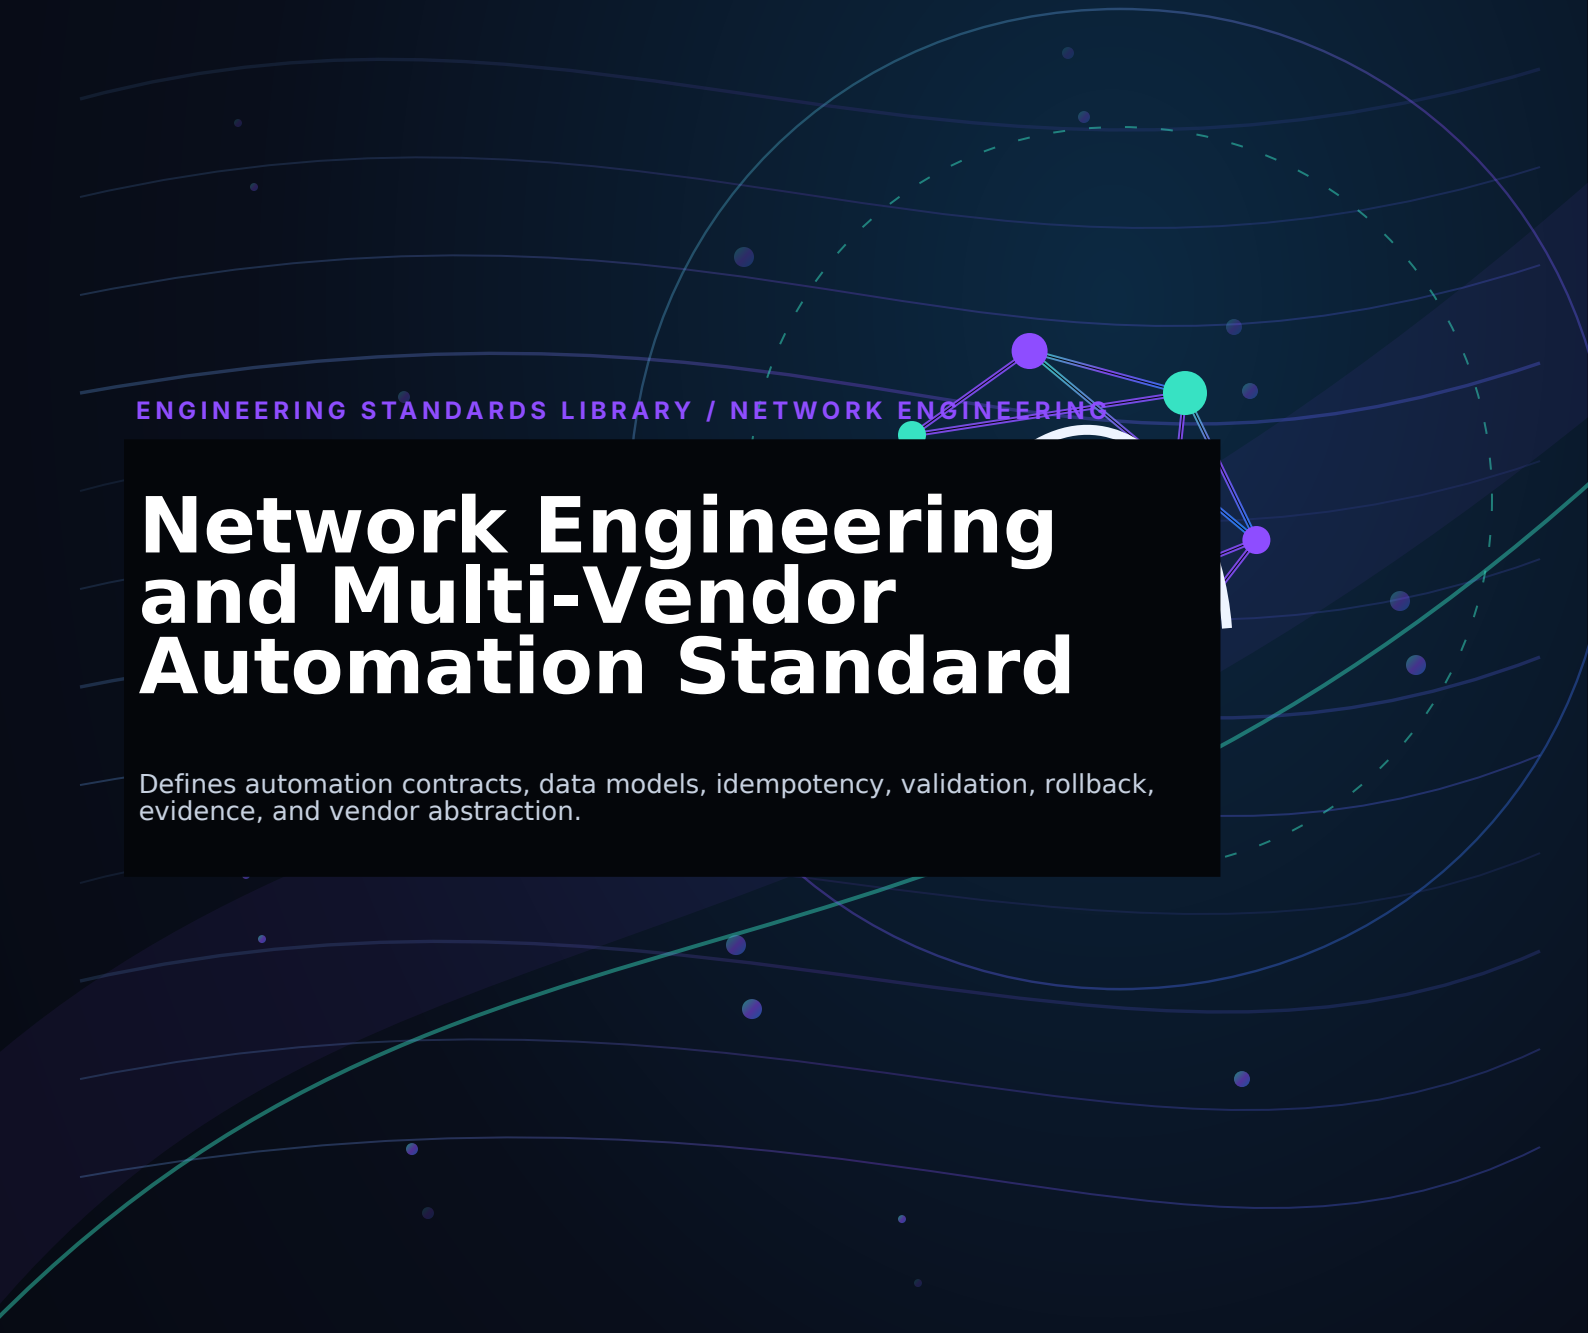

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

Network Engineering and Multi-Vendor Automation Standard

Defines automation contracts, data models, idempotency, validation, rollback, evidence, and vendor abstraction.

PERMANENT DOCUMENT ID
MYTHOS-NET-0002

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
10 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

Network Engineering and Multi-Vendor Automation Standard

Defines automation contracts, data models, idempotency, validation, rollback, evidence, and vendor abstraction.

DOCUMENT ID
MYTHOS-NET-0002

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

10	6	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-NET-0002 inside the network engineering standard constellation

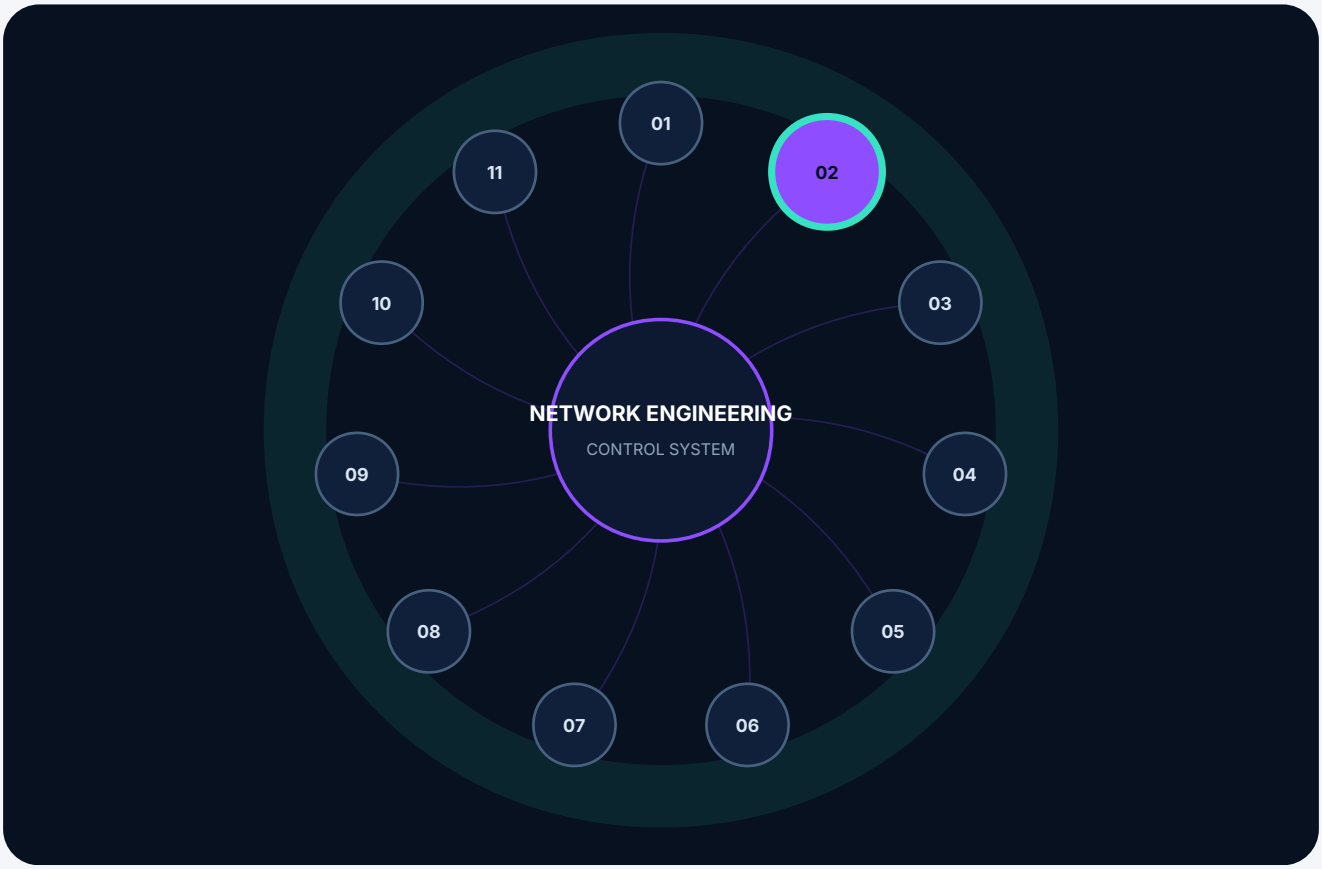

Connected assurance

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

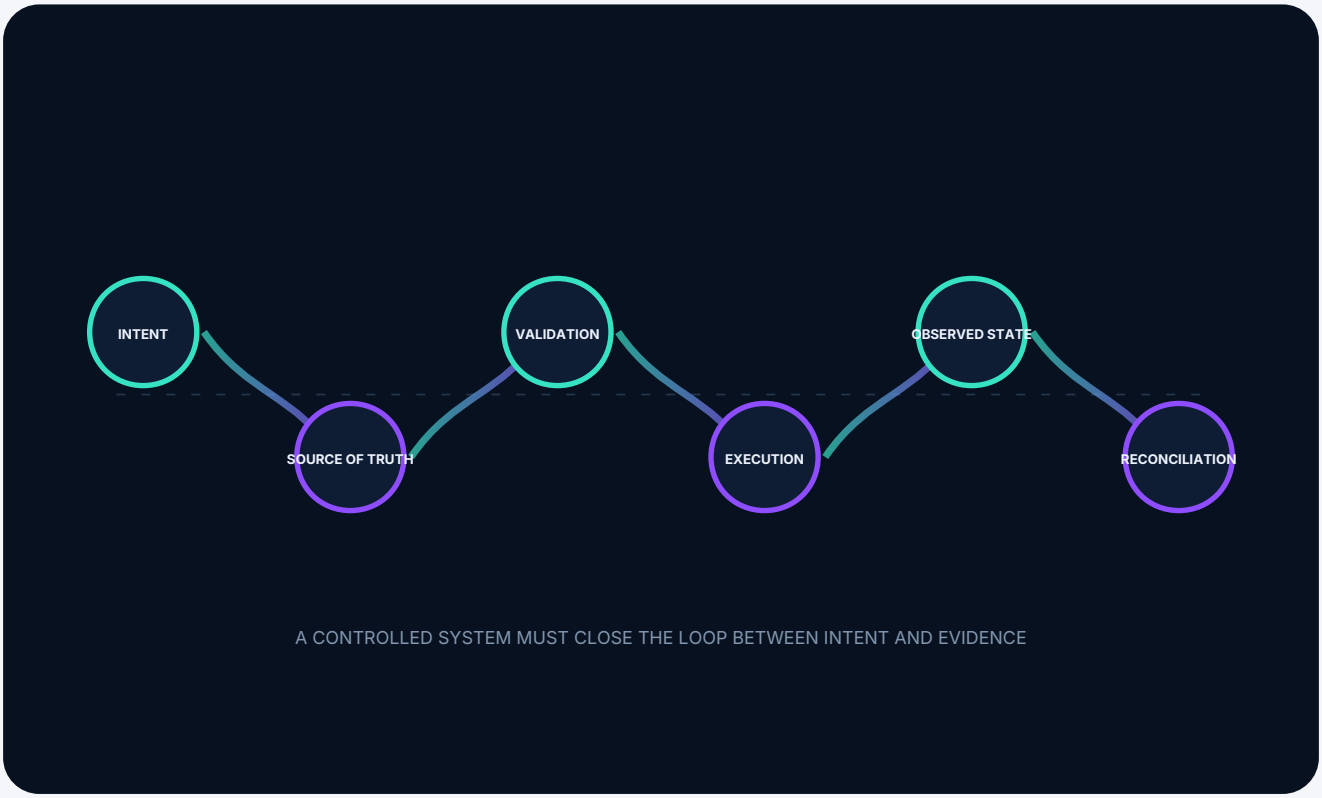

State separation

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-NET-0002

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

The practice of network engineering has failed to scale.

Organizations still manage networks via artisanal CLI commands and fragmented, vendor-specific scripts. This approach collapses under the weight of multi-vendor environments and completely breaks when faced with the strict lossless, high-bandwidth requirements of AI/GPU distributed training clusters. Furthermore, as network control planes shift to APIs, the security of those APIs - specifically token management and automation identity - has been catastrophically neglected.

This standard exists because:

1. **CLI is a Liability:** Typing commands is not engineering. It is unversioned, unaudited, non-idempotent, and fundamentally unscalable. Network engineering MUST be API-first, model-driven, and fully automated.
2. **AI/GPU Fabrics Demand a New Paradigm:** Standard Ethernet prioritization is insufficient for GPU distributed training (e.g., NCCL collectives). A single dropped packet or congestion event causes a training stride to stall, stranding millions of dollars in GPU compute. Networks MUST be engineered for RDMA (RoCEv2, InfiniBand) with mathematically verified lossless queues.
3. **API Security is the New Control Plane Security:** Automating the network via RESTCONF, gNMI, or vendor SDKs without robust token management, mTLS, and OAuth2/OIDC is equivalent to leaving the root password on a sticky note. Automation identities MUST be governed by zero-trust principles.
4. **Virtualization Must Evolve:** Heavyweight VM-based network functions (NFV) and kernel bypasses are being replaced by eBPF dataplanes and SR-IOV for hardware acceleration. The network operating system MUST support programmable dataplanes.

This document establishes the **Mythos Network Engineering & Automation Standard (MNEAS)**, a comprehensive, evidence-based methodology for evaluating multi-vendor automation, GPU fabric readiness, API security, and network virtualization.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Multi-vendor network automation and configuration management
- API-driven network control (RESTCONF, gNMI, NETCONF, OpenConfig)
- AI/GPU cluster fabric engineering (RoCEv2, InfiniBand, NVLink, CXL)
- Lossless Ethernet and congestion management (PFC, ECN, DCQCN)
- API security, token management, and zero-trust automation identity
- Network virtualization, eBPF dataplanes, and SR-IOV
- Idempotent configuration deployment and drift remediation

1.2 Out of Scope

- Network architecture topology design (covered in MYTHOS-NET-0001)
- Digital twin simulation execution (covered in MYTHOS-NET-0003)
- General IAM/Zero-Trust policy (covered in MYTHOS-SEC-0003)

1.3 Audience

- Network engineers and automation architects
- AI/ML infrastructure engineers managing GPU clusters
- Security engineers evaluating network API attack surfaces
- Platform engineering teams building network-as-code pipelines
- Standards bodies seeking a reference network automation methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Engineering Dimensions

Dimension	Definition
Model-Driven	Network configuration is derived from structured data models (YANG) rather than raw CLI strings.
RoCEv2	RDMA over Converged Ethernet version 2. Allows direct memory access between GPUs over an IP network, bypassing the OS kernel. Requires lossless Ethernet.
Lossless Fabric	A network architecture configured to guarantee zero packet loss due to congestion, mandatory for RDMA/AI workloads. Uses PFC (Priority Flow Control) and ECN (Explicit Congestion Notification).
gNMI	gRPC Network Management Interface. A high-performance, telemetry and configuration API using Protocol Buffers and gRPC.
eBPF	Extended Berkeley Packet Filter. Allows sandboxed programs to run in the Linux kernel without modifying kernel source, enabling high-performance, programmable dataplanes.

2.2 The Network Engineering Doctrine

> A CLI command is an anecdote. A data model is a contract. An API without auth is a backdoor. A GPU stalled on a network is a stranded asset. Automation is not optional; it is the only path to operational survival.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; CLI-driven, no API security, standard Ethernet only
1	Basic scripting but no data models or GPU fabric awareness
2	Functional API automation but lacks idempotency or token governance
3	Solid production performance; model-driven, secure APIs, basic RoCEv2
4	Strong performance; multi-vendor abstraction, eBPF, lossless verified
5	Best-in-class; sets the standard for autonomous, AI-fabric-ready networking

Weighting

Category	Weight	Rationale
Multi-Vendor Automation & Idempotency	20%	Vendor-specific automation is a silo; idempotent deployment is mandatory
AI/GPU Fabric Engineering (RoCEv2/IB)	20%	AI workloads demand lossless, low-latency fabrics
API Security & Token Management	20%	Unauthenticated APIs are critical attack surfaces
Network Virtualization & Dataplane	15%	Heavyweight NFV is obsolete; eBPF/SR-IOV is required
Telemetry & Observability (gNMI)	10%	Polling is dead; streaming telemetry is mandatory
Operational Lifecycle & Drift	15%	Configuration drift causes outages

Total: 100%

A system MUST score at least 3.0 in Multi-Vendor Automation, AI/GPU Fabric Engineering, and API Security to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

10 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Multi-Vendor Automation and Idempotency

SOURCE WEIGHT 20%

4.1.1
Data Model Abstraction

4.1.2
Idempotent Deployment

MYTHOS-NET-0002-EVAL-4.1.1

Data Model Abstraction

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Multi-Vendor Automation and Idempotency	**4.1.1**	**20%**	**Evidence-led**

Does the automation platform use vendor-agnostic data models (OpenConfig/YANG) rather than raw CLI templates?

0
CLI templates only; vendor-specific

1
Hybrid CLI/YANG but primarily CLI-driven

2
YANG models exist but require vendor-specific augmentation for standard tasks

3
Native OpenConfig/YANG support for multi-vendor configuration

4
Fully model-driven; CLI is never generated directly; vendor translation is abstracted

5
Leading abstraction: formal validation of YANG models, zero vendor-specific logic in automation code, automated conformance testing

ENGINEERING INTERPRETATION

This criterion evaluates whether data model abstraction 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-NET-0002-EVAL-4.1.2

Idempotent Deployment

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Multi-Vendor Automation and Idempotency	**4.1.2**	**20%**	**Evidence-led**

Are network configuration deployments idempotent and self-reconciling?

0
Deployments are imperative and destructive on retry

1
Basic scripts with manual rollback

2
Declarative (Terraform/Ansible) but drift is not reconciled

3
Declarative with periodic drift detection

4
Continuous reconciliation (GitOps) with automated drift remediation

5
Leading idempotency: formally verified desired-state, continuous reconciliation, and zero configuration divergence

ENGINEERING INTERPRETATION

This criterion evaluates whether idempotent deployment 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

AI/GPU Fabric Engineering (RoCEv2/ InfiniBand)

SOURCE WEIGHT 20%

4.2.1
Lossless Ethernet Readiness

4.2.2
AI Cluster Topology Awareness

MYTHOS-NET-0002-EVAL-4.2.1

Lossless Ethernet Readiness

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
AI/GPU Fabric Engineering (RoCEv2/ InfiniBand)	4.2.1	20%	Evidence-led

Can the automation platform correctly configure and verify lossless Ethernet for RoCEv2?

0
No awareness of RDMA or lossless requirements

1
Manual PFC/ECN configuration possible

2
Automated PFC/ECN deployment but no verification

3
Automated deployment with post-check verification of buffer thresholds

4
Continuous monitoring of PFC pauses and ECN marks; automated threshold tuning

5
Leading lossless fabric: intent-driven QoS, formally verified buffer allocation, real-time congestion avoidance (DCQCN), and zero-training-stride stalls

ENGINEERING INTERPRETATION

This criterion evaluates whether lossless ethernet readiness 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-NET-0002-EVAL-4.2.2

AI Cluster Topology Awareness

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
AI/GPU Fabric Engineering (RoCEv2/ InfiniBand)	4.2.2	20%	Evidence-led

Does the automation platform understand and optimize for AI/GPU topologies (Rail-Optimized, Rail-Aligned, NVLink, CXL)?

0
Unaware of GPU/NIC relationships

1
Basic server provisioning; no fabric optimization

2
Awareness of NIC placement but no topology-driven config

3
Automates rail-optimized leaf-spine configurations for GPU clusters

4
Integrates with GPU schedulers (Slurm/ K8s) to dynamically adjust fabric for workloads

5
Leading awareness: intent-driven AI fabric, CXL memory pooling integration, NVLink topology awareness, and automated path optimization for NCCL

ENGINEERING INTERPRETATION

This criterion evaluates whether ai cluster topology awareness 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

API Security and Token Management

SOURCE WEIGHT 20%

4.3.1
Authentication and Authorization

4.3.2
API Hardening

MYTHOS-NET-0002-EVAL-4.3.1

Authentication and Authorization

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
API Security and Token Management	**4.3.1**	**20%**	**Evidence-led**

How are automation service accounts and API tokens secured?

0
Static shared secrets or SNMP community strings

1
Local device accounts with long-lived passwords

2
Centralized TACACS+/RADIUS but long-lived API tokens

3
OAuth2.0/OIDC with short-lived tokens and role-based access

4
mTLS with SPIFFE/SPIRE identity for all automation agents

5
Leading zero-trust: PQC-agile mTLS, SPIFFE identity, per-session scoped tokens, cryptographic attestation of automation identity, and zero standing privileges

ENGINEERING INTERPRETATION

This criterion evaluates whether authentication and authorization 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

- API inventory, authorization policy, token metadata, gateway logs, and abuse controls
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test object and function authorization, token misuse, replay, expiration, and rate limits
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- bearer tokens without audience limits, hidden APIs, and inconsistent authorization
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- authorized API coverage
- token age
- replay rejection
- abuse-control efficacy
- 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-NET-0002-EVAL-4.3.2

API Hardening

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
API Security and Token Management	4.3.2	20%	Evidence-led

Are network device APIs (gNMI, RESTCONF) secured against injection and abuse?

0
APIs disabled or exposed without auth

1
Basic TLS but no rate limiting or input validation

2
Rate limiting and TLS 1.3 enforced

3
RBAC on API endpoints; strict payload validation (YANG schema)

4
API gateway with WAF, behavioral anomaly detection, and audit logging

5
Leading hardening: formal verification of API schemas, PQC TLS, zero-trust API gateway, and continuous red-team validation

ENGINEERING INTERPRETATION

This criterion evaluates whether api hardening 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

- API inventory, authorization policy, token metadata, gateway logs, and abuse controls
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test object and function authorization, token misuse, replay, expiration, and rate limits
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- bearer tokens without audience limits, hidden APIs, and inconsistent authorization
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- authorized API coverage
- token age
- replay rejection
- abuse-control efficacy
- 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

Network Virtualization and Dataplane

SOURCE WEIGHT 15%

4.4.1
Dataplane Architecture

4.4.2
Virtualization Overlay Integration

MYTHOS-NET-0002-EVAL-4.4.1

Dataplane Architecture

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Network Virtualization and Dataplane	**4.4.1**	**15%**	**Evidence-led**

Does the network OS leverage modern, programmable dataplanes (eBPF, SR-IOV, DPDK) rather than legacy kernel bridging?

0
Linux kernel bridging/routing only

1
OVS with kernel datapath

2
OVS with DPDK or SR-IOV for hardware offload

3
eBPF-accelerated networking (Cilium) for container/workload isolation

4
Programmable SmartNIC/DPU offload (P4/eBPF on hardware)

5
Leading dataplane: DPU-anchored isolation, eBPF programmable, SR-IOV for GPU Direct RDMA, zero CPU overhead for network I/O

ENGINEERING INTERPRETATION

This criterion evaluates whether dataplane architecture 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-NET-0002-EVAL-4.4.2

Virtualization Overlay Integration

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Network Virtualization and Dataplane	4.4.2	15%	Evidence-led

How well does the physical automation integrate with virtualization overlays (Kubernetes CNI, VMware NSX)?

0
Physical and virtual networks managed separately

1
Manual VLAN mapping between physical and virtual

2
Automated VLAN/VXLAN provisioning but no policy synchronization

3
Unified policy enforcement across physical and virtual (Cilium/NSX)

4
Intent-driven overlay-underlay stitching; automated MTU/BGP alignment

5
Leading integration: single source of truth, unified policy, formally verified path from hypervisor/Pod to spine

ENGINEERING INTERPRETATION

This criterion evaluates whether virtualization overlay 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

- 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

Telemetry and Observability (gNMI)

SOURCE WEIGHT 10%

4.5.1
Streaming Telemetry

MYTHOS-NET-0002-EVAL-4.5.1

Streaming Telemetry

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Telemetry and Observability (gNMI)	4.5.1	10%	Evidence-led

Has the system replaced SNMP polling with high-frequency streaming telemetry?

0
SNMP v2c/v3 polling only

1
Streaming telemetry available but not default

2
Default streaming telemetry (gNMI/gRPC) for critical metrics

3
Model-driven telemetry (OpenConfig paths) for all state

4
High-frequency (<1s) telemetry with on-change subscriptions

5
Leading observability: deterministic streaming, eBPF-extended metrics, real-time congestion/lossless state, and AI-driven anomaly detection

ENGINEERING INTERPRETATION

This criterion evaluates whether streaming telemetry 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.

4.6 / CAPABILITY FAMILY

Operational Lifecycle and Drift

SOURCE WEIGHT 15%

4.6.1
Configuration Drift Management

MYTHOS-NET-0002-EVAL-4.6.1

Configuration Drift Management

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Operational Lifecycle and Drift	4.6.1	15%	Evidence-led

How is configuration drift detected and remediated?

0
No drift detection; out-of-band changes are permanent

1
Manual diff checks

2
Automated nightly diff with alerts

3
Automated drift remediation (desired-state push)

4
Continuous reconciliation with GitOps; drift blocked by default

5
Leading lifecycle: formally verified state, continuous reconciliation, cryptographic attestation of running config, and zero tolerance for drift

ENGINEERING INTERPRETATION

This criterion evaluates whether configuration drift management 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 Network Engineering Suite (MNEAS)

Traditional network benchmarks measure protocol convergence. The MNEAS evaluates automation fidelity, API security, and AI fabric readiness.

5.1 Suite Composition

5.1.1 MNEAS-Fabric

- **Lossless Validation:** 50+ tests injecting congestion into RoCEv2 fabrics to verify PFC/ECN prevents drops.
- **GPU Topology:** 50+ tests verifying network paths align with GPU rail topology for distributed training.

5.1.2 MNEAS-Automation

- **Idempotency:** 100+ tests applying the same network intent 5 times to verify zero configuration duplication or change.
- **API Exploit:** 50+ tests attempting to inject malicious YANG payloads or use expired tokens against network controllers.

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.

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

Network Digital Twin and Simulation Verification Standard

Defines topology models, simulations, reachability, differential testing, evidence, and confidence.

PERMANENT DOCUMENT ID
MYTHOS-NET-0003

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
9 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

Network Digital Twin and Simulation Verification Standard

Defines topology models, simulations, reachability, differential testing, evidence, and confidence.

DOCUMENT ID
MYTHOS-NET-0003

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

9	6	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-NET-0003 inside the network engineering standard constellation

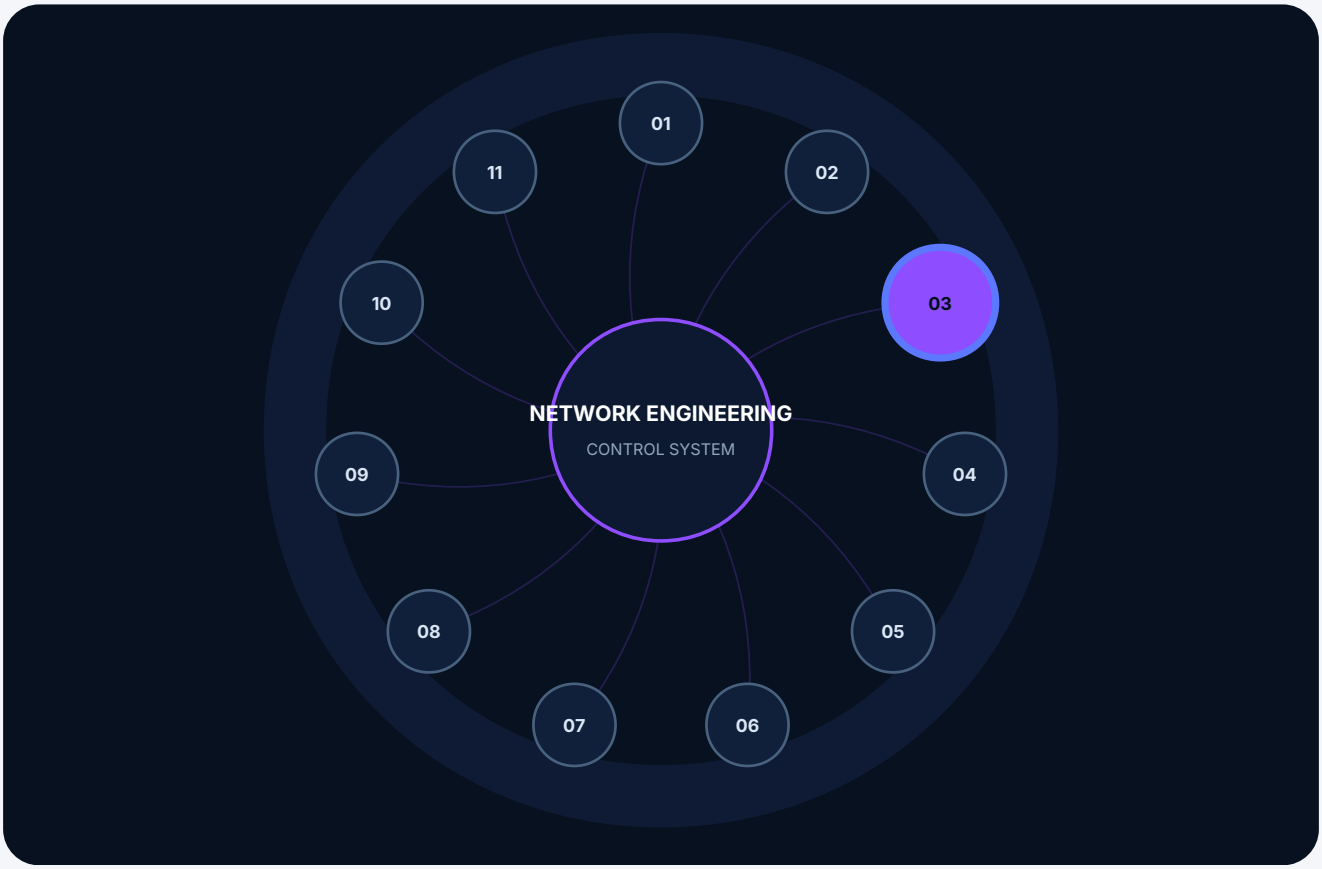

Connected assurance

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

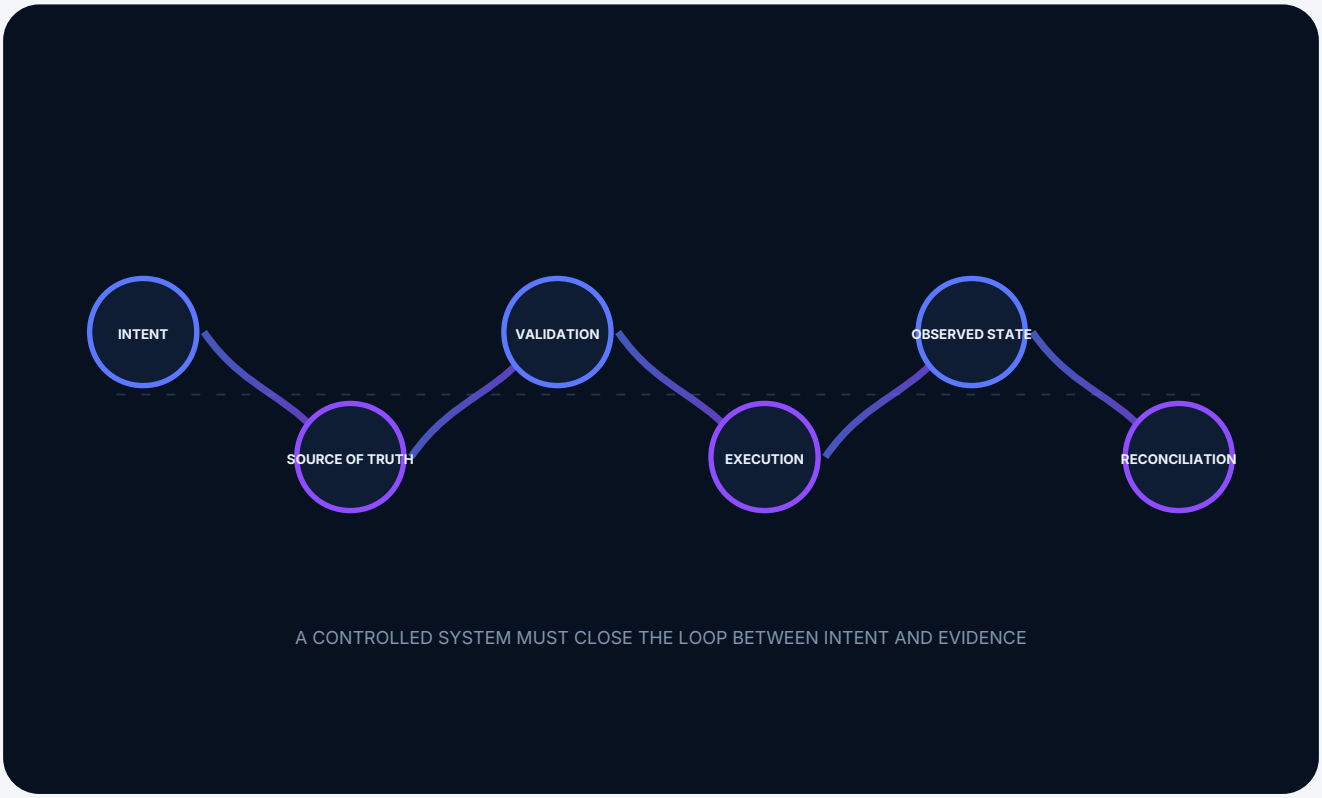

State separation

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-NET-0003

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

The validation of network changes has failed.

"Testing in production" is not a deployment strategy; it is a confession of operational failure. Organizations apply configuration changes to live networks and monitor dashboards to see if anything breaks. This reactive posture causes outages, security breaches, and catastrophic blast radius events that could have been prevented by mathematical simulation.

This standard exists because:

1. **Ping is Not Proof:** Reaching a destination via ICMP tells you nothing about the ACLs, asymmetric routing, or MTU mismatches that will trigger under real application traffic. Verification MUST be formal and mathematical.
2. **Diagrams are Not Twins:** A Visio drawing is a static picture. A digital twin is a live, stateful, executable model of the network that ingests real routing tables, ACLs, and configurations to prove reachability.
3. **Manual Peer Review is Insufficient:** Having a senior engineer read a diff is theater. Human cognition cannot compute the combinatorial explosion of BGP path selection, route-maps, and firewall rules across a 1,000-node topology. Machines must verify.
4. **Simulation Must Precede Deployment:** No configuration change should ever reach a production device unless it has passed a formal simulation verifying that intent is met and blast radius is bounded.

This document establishes the **Mythos Network Digital Twin Standard (MNDTS)**, a comprehensive, evidence-based methodology for evaluating the fidelity, verification capabilities, and CI/CD integration of network simulation platforms.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Network digital twin platforms (Forward Networks, Nokia BVT, Cisco WAE)
- Formal verification engines (Batfish,cenizer)
- Containerized network labs (Containerlab, Cisco Modelling Labs, EVE-NG, GNS3)
- Reachability assertion and blast radius simulation
- Failure injection and chaos engineering for network topologies
- Pre-deployment CI/CD pipeline integration for network changes
- AI-driven traffic prediction and anomaly simulation

1.2 Out of Scope

- Network architecture topology design (covered in MYTHOS-NET-0001)
- Network automation execution (covered in MYTHOS-NET-0002)
- General chaos engineering platforms (covered in MYTHOS-SRE)

1.3 Audience

- Network engineers and automation architects building validation pipelines
- Security teams evaluating pre-deployment blast radius
- Platform engineering teams integrating network verification into CI/CD
- AI governance, risk, and compliance teams
- Standards bodies seeking a reference network simulation methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Simulation Dimensions

Dimension	Definition
Digital Twin	A live, stateful, mathematical model of the network that ingests real device state (configs, routes, ACLs) to create an executable replica.
Formal Verification	The mathematical proof that a network's routing and access control tables produce a specific reachability outcome (e.g., Batfish assertions).
Reachability Assertion	A test that verifies a specific flow is permitted or denied (e.g., "Assert: Host A cannot reach Host B on Port 22").
Blast Radius Simulation	The mathematical modeling of the impact of a device or link failure, or a configuration change, on the broader network.
Topology Parity	The degree to which the digital twin's state matches the production network's actual state.

2.2 The Simulation Doctrine

> Testing in production is a failure. Ping is not proof. A diagram is not a twin. A change that is not simulated is an outage. Verification must be mathematical, not manual.

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 simulation, production testing only
1	Basic lab environment but manually maintained, no formal verification
2	Functional twin but lacks formal verification or CI/CD integration
3	Solid production performance; formal verification, basic CI/CD integration
4	Strong performance; full digital twin, blast radius, automated parity
5	Best-in-class; sets the standard for autonomous, verified network simulation

Weighting

Category	Weight	Rationale
Formal Reachability Verification	25%	Mathematical proof of flows is the primary value of simulation
Topology Parity & State Ingestion	20%	A twin that does not match reality is a hallucination
CI/CD Pipeline Integration	20%	Simulation must be a mandatory gate, not an optional tool
Failure Injection & Blast Radius	15%	Unvalidated resilience is assumed fragility
Multi-Vendor & Protocol Fidelity	10%	Inaccurate protocol models produce false proofs
Operational Lifecycle	10%	Twins must be governed continuously

Total: 100%

A system MUST score at least 3.0 in Formal Reachability Verification and Topology Parity to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

9 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Formal Reachability Verification

SOURCE WEIGHT 25%

4.1.1
Assertion Engine

4.1.2
Differential Analysis

MYTHOS-NET-0003-EVAL-4.1.1

Assertion Engine

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Formal Reachability Verification	**4.1.1**	**25%**	**Evidence-led**

Can the system mathematically prove or disprove reachability assertions?

0
No formal verification; relies on ping/ traceroute

1
Basic path visualization but no logical assertion

2
Can assert L3 reachability but ignores ACLs/Firewalls

3
Full L3/L4 reachability assertion including ACLs, NAT, and routing (Batfish-level)

4
Assertion includes L7 (application) awareness and packet-header tracing

5
Leading verification: formally proven reachability invariants, mathematical proof of policy compliance, and zero false positives/negatives

ENGINEERING INTERPRETATION

This criterion evaluates whether assertion engine 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-NET-0003-EVAL-4.1.2

Differential Analysis

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Formal Reachability Verification	**4.1.2**	**25%**	**Evidence-led**

Can the system compare the reachability of the current state vs. the proposed change?

0
No differential capability

1
Text-based config diff only

2
Manual comparison of two simulation snapshots

3
Automated differential analysis showing added/removed flows

4
Differential analysis bounded by blast radius policies (e.g., "No changes to DB flows permitted")

5
Leading differential: formally verified equivalence classes, automatic blocking of policy-violating diffs, cryptographic sealing of change impact

ENGINEERING INTERPRETATION

This criterion evaluates whether differential analysis 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

Topology Parity and State Ingestion

SOURCE WEIGHT 20%

4.2.1
Automated State Ingestion

4.2.2
Parity Validation

MYTHOS-NET-0003-EVAL-4.2.1

Automated State Ingestion

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Topology Parity and State Ingestion	**4.2.1**	**20%**	**Evidence-led**

How is the digital twin populated with network state?

0
Manual topology creation

1
Automated discovery but requires manual triggering

2
Automated ingestion of configurations via CLI/SSH scraping

3
Automated ingestion via APIs (gNMI/ RESTCONF) and streaming telemetry

4
Continuous state ingestion synchronized with Source-of-Truth (SoT)

5
Leading parity: real-time state ingestion, cryptographic attestation of twin-to-production fidelity, and zero manual intervention

ENGINEERING INTERPRETATION

This criterion evaluates whether automated state ingestion 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-NET-0003-EVAL-4.2.2

Parity Validation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Topology Parity and State Ingestion	**4.2.2**	**20%**	**Evidence-led**

Does the system verify that the twin accurately reflects production?

0
No parity validation; twin is assumed correct

1
Manual spot-checks of routing tables

2
Automated comparison of device counts and links

3
Automated comparison of routing tables and ACLs

4
Continuous parity validation with drift alerts

5
Leading validation: mathematical proof of state equivalence, automatic twin rebuild on drift, and zero simulation-reality divergence

ENGINEERING INTERPRETATION

This criterion evaluates whether parity validation 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

CI/CD Pipeline Integration

SOURCE WEIGHT 20%

4.3.1
Pre-Merge Simulation Gate

4.3.2
Post-Deployment Verification

MYTHOS-NET-0003-EVAL-4.3.1

Pre-Merge Simulation Gate

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
CI/CD Pipeline Integration	**4.3.1**	**20%**	**Evidence-led**

Is simulation a mandatory, automated gate in the CI/CD pipeline?

0
No pipeline integration; simulation is ad-hoc

1
Manual simulation step before deployment

2
Pipeline triggers simulation, but results are manually reviewed

3
Pipeline triggers simulation, automated pass/fail based on reachability assertions

4
Pipeline blocks merge on assertion failure; auto-generates blast radius report

5
Leading integration: formally verified gate, cryptographic signing of simulation results, and zero path to production without simulation proof

ENGINEERING INTERPRETATION

This criterion evaluates whether pre-merge simulation gate 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-NET-0003-EVAL-4.3.2

Post-Deployment Verification

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
CI/CD Pipeline Integration	4.3.2	20%	Evidence-led

Does the system verify production state matches the simulated intent after deployment?

0
No post-deployment verification

1
Manual ping/traceroute checks

2
Automated connectivity checks (synthetic monitoring)

3
Automated reachability assertion in production (comparing to twin)

4
Continuous production-vs-twin parity verification

5
Leading verification: continuous mathematical proof of production intent, automatic rollback on divergence

ENGINEERING INTERPRETATION

This criterion evaluates whether post-deployment 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

Failure Injection and Blast Radius

SOURCE WEIGHT 15%

4.4.1
Failure Simulation

MYTHOS-NET-0003-EVAL-4.4.1

Failure Simulation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Failure Injection and Blast Radius	**4.4.1**	**15%**	**Evidence-led**

Can the system simulate device, link, or process failures?

0
No failure simulation

1
Manual failure simulation in lab

2
Automated single-link/node failure simulation

3
Automated multi-failure (compound) simulation

4
Blast radius analysis showing impact of failures on critical flows

5
Leading simulation: formally verified failure domains, chaos engineering integration, AI-driven failure probability weighting, and automatic topology remediation

ENGINEERING INTERPRETATION

This criterion evaluates whether failure 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.

4.5 / CAPABILITY FAMILY

Multi-Vendor and Protocol Fidelity

SOURCE WEIGHT 10%

4.5.1
Behavioral Modeling

MYTHOS-NET-0003-EVAL-4.5.1

Behavioral Modeling

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Multi-Vendor and Protocol Fidelity	**4.5.1**	**10%**	**Evidence-led**

Does the simulation accurately model complex protocol behaviors (BGP path selection, EVPN Type-5 routes, PFC storms)?

0
Generic models; inaccurate protocol behavior

1
Basic protocol models (OSPF, BGP) but no vendor-specific nuances

2
Vendor-specific models for top 3 vendors

3
Accurate modeling of complex features (EVPN, SRv6, PFC)

4
Containerized vendor images (Containerlab/CML) for bit-exact behavioral parity

5
Leading fidelity: bit-exact emulation, formally verified protocol models, and zero behavioral divergence

ENGINEERING INTERPRETATION

This criterion evaluates whether behavioral 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.

4.6 / CAPABILITY FAMILY

Operational Lifecycle

SOURCE WEIGHT 10%

4.6.1
Twin Lifecycle Management

MYTHOS-NET-0003-EVAL-4.6.1

Twin Lifecycle Management

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Operational Lifecycle	4.6.1	10%	Evidence-led

Is the digital twin managed with the same rigor as production infrastructure?

0
Ad-hoc twin creation; no lifecycle management

1
Persistent twin but no version control

2
Version-controlled twin definitions

3
Automated twin lifecycle (spin up, ingest, verify, teardown) in CI/CD

4
Continuous persistent twin with automated scaling

5
Leading lifecycle: twin-as-code, cryptographic versioning, resource-optimized scheduling, and continuous fidelity auditing

ENGINEERING INTERPRETATION

This criterion evaluates whether twin lifecycle management 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 Network Twin Suite (MNTS)

Traditional network testing relies on ping. The MNTS evaluates mathematical proof, parity, and pipeline integration.

5.1 Suite Composition

5.1.1 MNTS-Verification

- **Reachability Assertions:** 100+ Batfish-style assertions verifying critical flows (e.g., "Web cannot reach DB on 3306").
- **Differential Checks:** 50+ proposed changes verifying no unintended flows are introduced.

5.1.2 MNTS-Parity

- **Topology Drift:** 50+ tests injecting production drift to verify the twin detects and alerts.
- **State Ingestion:** 50+ tests verifying routing table and ACL parity between twin and production.

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.

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

Network Security Architecture and Zero-Trust Networking Standard

Defines identity-aware networking, segmentation, enforcement, telemetry, and continuous validation.

PERMANENT DOCUMENT ID
MYTHOS-NET-0004

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
10 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

Network Security Architecture and Zero-Trust Networking Standard

Defines identity-aware networking, segmentation, enforcement, telemetry, and continuous validation.

DOCUMENT ID
MYTHOS-NET-0004

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

10	6	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-NET-0004 inside the network engineering standard constellation

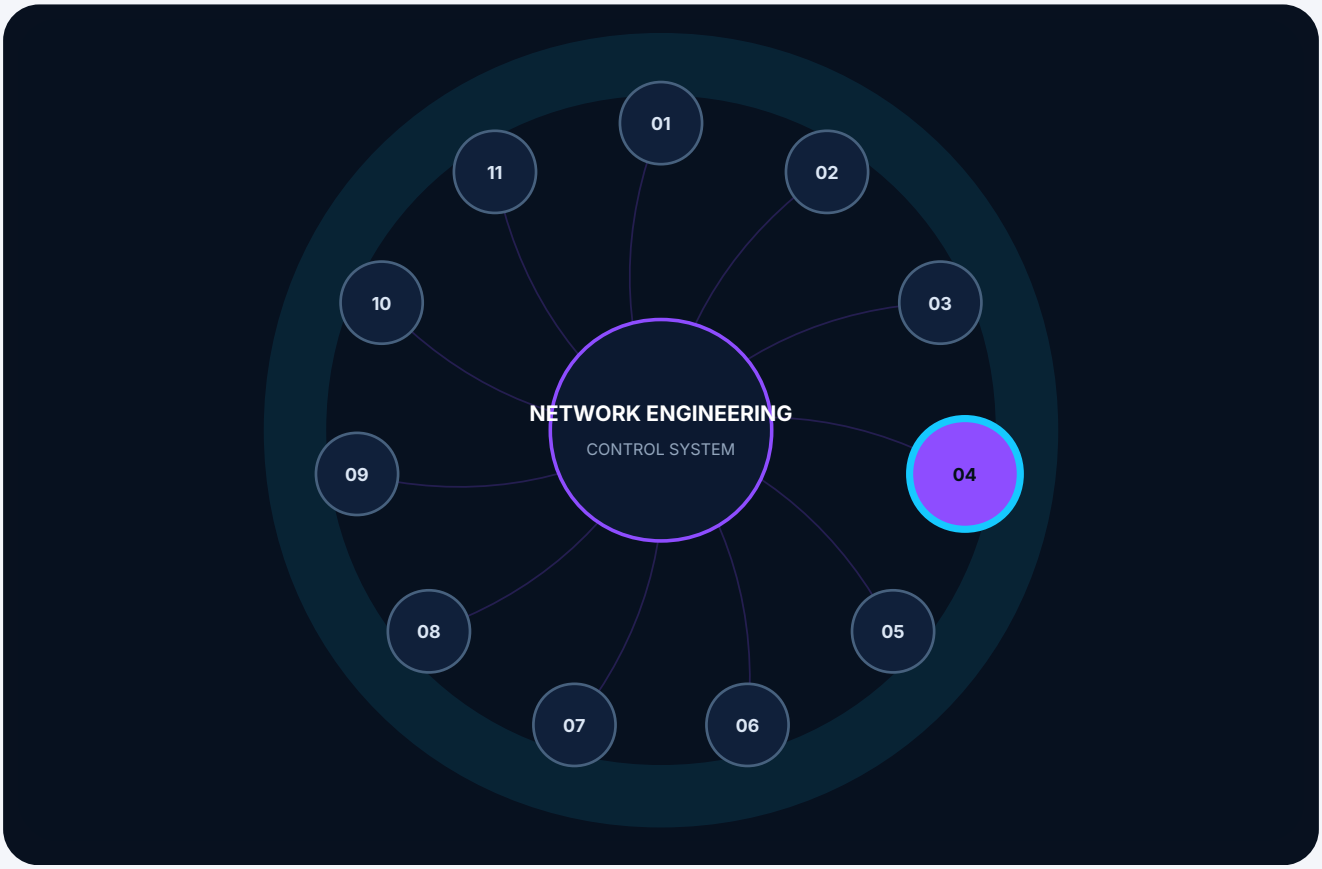

Connected assurance

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

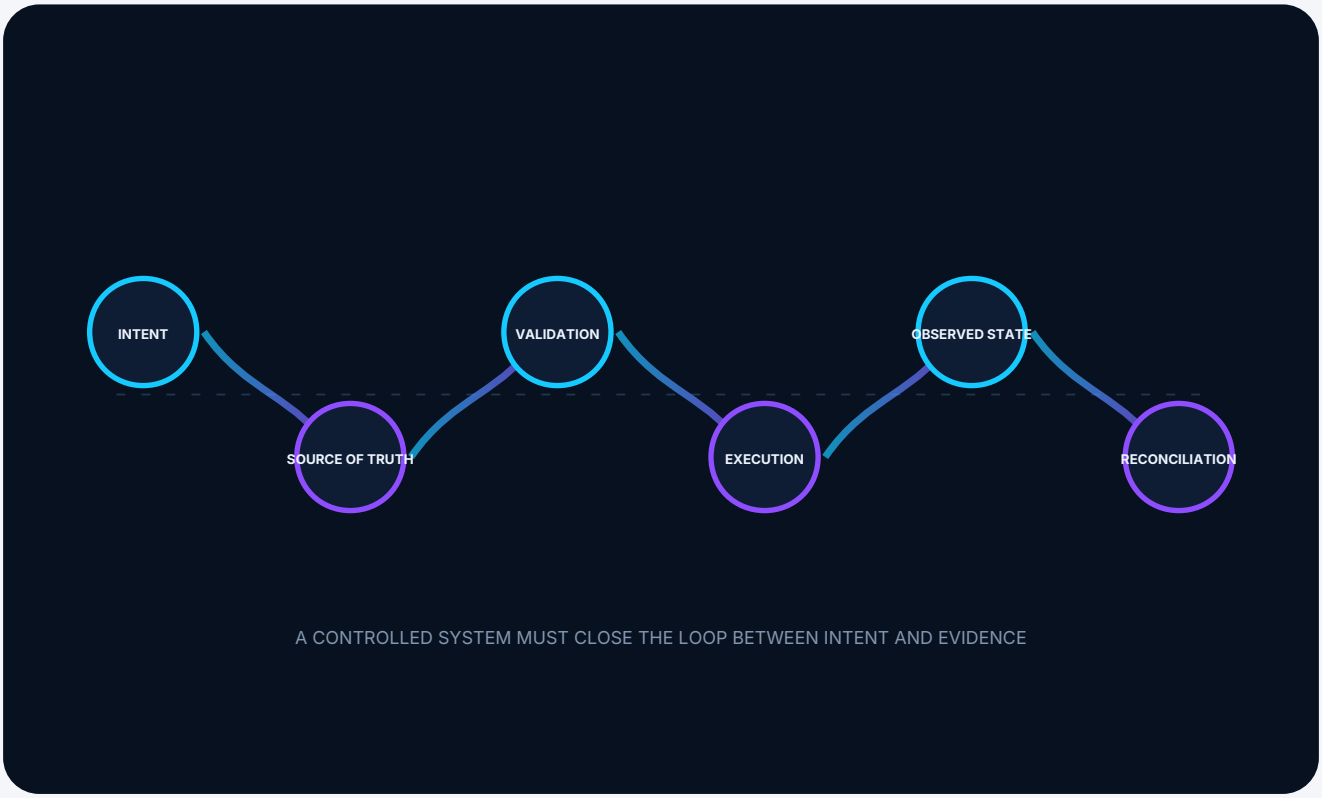

State separation

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-NET-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

The internal network is hostile. The perimeter has dissolved.

Organizations continue to architect their networks assuming a trusted interior, relying on North-South firewalls while ignoring the free-flowing East-West traffic of their data centers and clouds. This architecture is a relic. Once an attacker - or a compromised AI agent - breaches the perimeter, the flat internal network provides zero resistance to lateral movement.

This standard exists because:

1. **IP Addresses are Not Identity:** Firewalls and ACLs that filter on IP addresses are easily spoofed and fundamentally incapable of enforcing zero-trust principles. Network policy MUST be bound to cryptographic workload identity (SPIFFE/SPIRE) and user identity, not static IPs.
2. **The Network Must Encrypt Everything:** Unencrypted internal traffic (East-West) is an open book to anyone with a packet sniffer. MACsec for physical layers, IPSec/WireGuard for overlays, and PQC-TLS for API planes are mandatory, not optional.
3. **Network Access Control Must be Dynamic:** 802.1X and RADIUS must move from simple port-authentication to dynamic, posture-aware identity assignment. Devices and AI agents that fail posture checks must be quarantined automatically.
4. **AI Agents are Untrusted Workloads:** An AI agent processing untrusted data (e.g., RAG from the internet) must be treated as a potentially compromised entity. The network fabric MUST isolate agent execution planes from production data planes.

This document establishes the **Mythos Zero-Trust Networking Standard (MZTNS)**, a comprehensive, evidence-based methodology for evaluating the network fabric, identity integration, and encryption posture required for true zero-trust architecture.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Zero-Trust Network Access (ZTNA) and Software-Defined Perimeters (SDP)
- Network Access Control (NAC), 802.1X, RADIUS, and dynamic VLAN assignment
- Identity-aware network fabrics (SPIFFE/SPIRE integration)
- Encryption in transit (MACsec, IPSec, WireGuard, PQC-TLS)
- Network-level isolation and quarantine for AI agents
- DNS security (DoH, DoT, DNSSEC)

1.2 Out of Scope

- Microsegmentation policy and blast radius (covered in MYTHOS-SEC-0007)
- Zero-Trust policy engine authority and RBAC (covered in MYTHOS-SEC-0003)
- Post-Quantum Cryptography algorithm standardization (covered in MYTHOS-SEC-0001)

1.3 Audience

- Network security architects designing zero-trust fabrics
- Identity and Access Management (IAM) engineers
- Platform engineers managing SPIFFE/SPIRE or ZTNA overlays
- AI infrastructure teams securing agent network paths
- Standards bodies seeking a reference zero-trust networking methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Zero-Trust Networking Dimensions

Dimension	Definition
ZTNA	Zero-Trust Network Access. Replaces VPNs with identity-aware, per-application access, hiding the network from unauthorized users.
MACsec	Media Access Control Security. IEEE 802.1AE standard for encrypting data on Ethernet links, providing hop-by-hop confidentiality and integrity.
SPIFFE	Secure Production Identity Framework for Everyone. A set of open-source standards for securely identifying software systems in dynamic and heterogeneous environments.
802.1X	IEEE standard for port-based Network Access Control, authenticating devices before allowing network access.
Posture Assessment	The evaluation of a device's security state (OS patches, antivirus, configuration) before granting network access.

2.2 The Zero-Trust Networking Doctrine

> The internal network is hostile. IPs are not identity. Unencrypted traffic is a breach. A device without posture is a threat. An AI agent without network isolation is a loaded weapon.

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, no encryption, IP-based ACLs
1	Basic ZTNA for remote users, but internal network is flat
2	Functional but lacks identity integration or hardware encryption
3	Solid production performance; identity-aware, encrypted, NAC-enforced
4	Strong performance; SPIFFE-integrated, PQC-ready, agent-isolated
5	Best-in-class; sets the standard for autonomous, cryptographically-verified zero-trust networking

Weighting

Category	Weight	Rationale
Identity-Aware Networking (SPIFFE)	25%	IP-based trust is the root cause of lateral movement
Encryption in Transit & PQC Readiness	20%	Unencrypted East-West traffic is an open book
NAC & Dynamic Posture Enforcement	15%	Unassessed devices are attack vectors
AI Agent Network Isolation	15%	Agents are high-risk workloads
ZTNA & Application Access	15%	VPNs provide excessive network access
Operational Lifecycle & DNS Security	10%	DNS is the most abused protocol

Total: 100%

A system MUST score at least 3.0 in Identity-Aware Networking and Encryption in Transit to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

10 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Identity-Aware Networking

SOURCE WEIGHT 25%

4.1.1
Cryptographic Workload Identity

4.1.2
User-to-Workload Binding

MYTHOS-NET-0004-EVAL-4.1.1

Cryptographic Workload Identity

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Identity-Aware Networking	**4.1.1**	**25%**	**Evidence-led**

Does the network fabric bind policy to cryptographic identity (SPIFFE/SPIRE) rather than IP addresses?

0
All policies based on static IP addresses

1
Service accounts exist but network policy still uses IPs

2
mTLS used for application layer; network layer still IP-dependent

3
Network fabric ingests SPIFFE IDs for policy enforcement (e.g., Cilium)

4
Network policy is 100% identity-aware; IPs are dynamic and irrelevant

5
Leading identity: formally verified mapping of cryptographic identity to network flows, zero IP-based rules, and automatic policy rotation with SVIDs

ENGINEERING INTERPRETATION

This criterion evaluates whether cryptographic workload 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-NET-0004-EVAL-4.1.2

User-to-Workload Binding

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Identity-Aware Networking	4.1.2	25%	Evidence-led

Can the network trace a flow from a human user identity through to a specific workload?

0
No user-to-workload mapping

1
Logs correlate via IP addresses

2
Identity-aware proxies (Envoy) provide L7 user context

3
Network fabric tags packets with user identity context

4
Full chain: User OIDC token → Service SVID → Network Flow logged in Clio

5
Leading binding: cryptographic attestation of the entire user-to-workload path, zero unattributable network flows

ENGINEERING INTERPRETATION

This criterion evaluates whether user-to-workload binding 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

Encryption in Transit and PQC Readiness

SOURCE WEIGHT 20%

4.2.1
East-West Encryption

4.2.2
PQC Network Readiness

MYTHOS-NET-0004-EVAL-4.2.1

East-West Encryption

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Encryption in Transit and PQC Readiness	**4.2.1**	**20%**	**Evidence-led**

Is internal network traffic (East-West) encrypted at the link or network layer?

0
Unencrypted internal traffic

1
Application-layer encryption (TLS) only where enforced by app

2
Overlay encryption (IPSec/WireGuard) for cross-zone traffic

3
Hop-by-hop MACsec on all physical links within data centers

4
Combined MACsec (L2) + mTLS (L7) for defense in depth

5
Leading encryption: MACsec + mTLS + PQC hybrid key exchange, formally verified cryptographic boundaries, zero unencrypted internal traffic

ENGINEERING INTERPRETATION

This criterion evaluates whether east-west encryption 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-NET-0004-EVAL-4.2.2

PQC Network Readiness

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Encryption in Transit and PQC Readiness	4.2.2	20%	Evidence-led

Are network control planes and data planes ready for Post-Quantum Cryptography?

0
Classical crypto only (RSA, ECC)

1
PQC evaluated but not implemented

2
PQC-enabled TLS terminations at edge

3
Hybrid classical+PQC key exchange (e.g., X25519+ML-KEM) for control plane (BGP, gNMI)

4
Hybrid PQC for data plane overlays (IPSec/WireGuard)

5
Leading PQC readiness: full hybrid PQC for L2/L3/L7, crypto-agile algorithm registries, and PQC-attested MACsec

ENGINEERING INTERPRETATION

This criterion evaluates whether pqc network readiness 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.

4.3 / CAPABILITY FAMILY

NAC and Dynamic Posture Enforcement

SOURCE WEIGHT 15%

4.3.1
Dynamic Access Control

4.3.2
Quarantine Enforcement

MYTHOS-NET-0004-EVAL-4.3.1

Dynamic Access Control

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
NAC and Dynamic Posture Enforcement	**4.3.1**	**15%**	**Evidence-led**

Does the network dynamically assign access based on device posture and identity?

0
No NAC; any device can connect

1
Basic 802.1X but static VLAN assignment

2
Dynamic VLAN assignment based on RADIUS attributes

3
Dynamic VLAN + posture assessment (OS patch level, EDR status)

4
Continuous posture reassessment; dynamic quarantine on degradation

5
Leading NAC: continuous posture, cryptographic device identity (TPM/EK), automatic quarantine, and zero static network assignments

ENGINEERING INTERPRETATION

This criterion evaluates whether dynamic access 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.

MYTHOS-NET-0004-EVAL-4.3.2

Quarantine Enforcement

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
NAC and Dynamic Posture Enforcement	4.3.2	15%	Evidence-led

Can the network instantly isolate a compromised device without manual intervention?

0
Manual cable pull or switch port shutdown

1
Manual reassignment to quarantine VLAN

2
Automated quarantine triggered by SIEM alert

3
Automated quarantine triggered by EDR/XDR in real-time

4
Micro-quarantine: isolating only the compromised process/port, not the whole device

5
Leading quarantine: formally verified isolation, automatic forensic capture, and zero lateral movement capability

ENGINEERING INTERPRETATION

This criterion evaluates whether quarantine 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.

4.4 / CAPABILITY FAMILY

AI Agent Network Isolation

SOURCE WEIGHT 15%

4.4.1
Agent Egress Control

4.4.2
Agent Data Path Isolation

MYTHOS-NET-0004-EVAL-4.4.1

Agent Egress Control

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
AI Agent Network Isolation	4.4.1	15%	Evidence-led

Are AI agents restricted to minimum necessary network access?

0
Agents have broad, unmonitored network access

1
Agents are in a separate VLAN but with broad egress

2
Agents have explicit allow-list egress rules

3
Egress is identity-scoped per task (capability-based networking)

4
Agents are routed through ZTNA proxies; no direct network access

5
Leading isolation: ephemeral network identities, per-task egress scopes, cryptographic attestation of network flows, and automatic quarantine on anomaly

ENGINEERING INTERPRETATION

This criterion evaluates whether agent egress 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.

MYTHOS-NET-0004-EVAL-4.4.2

Agent Data Path Isolation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
AI Agent Network Isolation	4.4.2	15%	Evidence-led

Are the network paths used by AI agents to access sensitive data isolated from production traffic?

0
Agents share production data paths

1
Agents use separate VLANs

2
Agents use VRF-lite or separate VRFs

3
Agents use encrypted overlays (VXLAN/IPSec) with separate routing tables

4
Agents use dedicated compute + network fabrics (SR-IOV with isolated VFs)

5
Leading isolation: hardware-enforced paths (DPU offload), zero shared state, and formally verified path separation

ENGINEERING INTERPRETATION

This criterion evaluates whether agent data path 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.

4.5 / CAPABILITY FAMILY

ZTNA and Application Access

SOURCE WEIGHT 15%

MYTHOS-NET-0004-EVAL-4.5.1

Remote Access Architecture

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
ZTNA and Application Access	4.5.1	15%	Evidence-led

Has the organization replaced VPNs with Zero-Trust Network Access (ZTNA)?

0
Full-tunnel VPN providing excessive network access

1
Split-tunnel VPN but broad internal access

2
ZTNA for third-party/contractor access; VPN for employees

3
ZTNA for all human access; VPN deprecated

4
ZTNA with continuous posture checking and micro-segmented application access

5
Leading ZTNA: per-resource cryptographic identity, continuous behavioral analysis, and zero implicit network trust

ENGINEERING INTERPRETATION

This criterion evaluates whether remote access architecture 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.

4.6 / CAPABILITY FAMILY

Operational Lifecycle and DNS Security

SOURCE WEIGHT 10%

4.6.1
DNS Security

MYTHOS-NET-0004-EVAL-4.6.1

DNS Security

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Operational Lifecycle and DNS Security	4.6.1	10%	Evidence-led

Is DNS traffic encrypted and validated?

0
Unencrypted DNS (Port 53) with no validation

1
DNSSEC validation for external zones

2
DoH (DNS over HTTPS) or DoT (DNS over TLS) for clients

3
DoH/DoT + DNSSEC for all internal and external resolution

4
Encrypted DNS + zero-trust DNS policy (filtering, RPZ)

5
Leading DNS: PQC-encrypted DNS, cryptographic DNS identity, zero plaintext DNS queries, and formally validated DNS paths

ENGINEERING INTERPRETATION

This criterion evaluates whether dns security 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 Zero-Trust Networking Suite (MZTNS)

Traditional network benchmarks measure throughput and latency. The MZTNS evaluates identity binding, encryption coverage, and quarantine efficacy.

5.1 Suite Composition

5.1.1 MZTNS-Identity

- **IP Spoofing:** 100+ tests attempting to bypass network policy by spoofing IP addresses (must fail if identity-aware).
- **SVID Rotation:** 50+ tests verifying network policy dynamically updates when SPIFFE SVIDs rotate.

5.1.2 MZTNS-Encryption

- **Unencrypted Flow Detection:** 100+ tests verifying no internal traffic passes without MACsec or mTLS.
- **PQC Downgrade:** 50+ tests attempting to force classical crypto negotiation (must fail).

5.1.3 MZTNS-Quarantine

- **Rogue Device:** 50+ tests connecting unauthorized devices to verify NAC quarantine.
- **Agent Anomaly:** 50+ tests simulating agent network scanning to verify automatic egress quarantine.

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.

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

Network Source of Truth and Drift Reconciliation Standard

Defines authoritative state, discovery, intent, reconciliation, exceptions, and audit.

PERMANENT DOCUMENT ID
MYTHOS-NET-0005

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
9 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

Network Source of Truth and Drift Reconciliation Standard

Defines authoritative state, discovery, intent, reconciliation, exceptions, and audit.

DOCUMENT ID
MYTHOS-NET-0005

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

9	6	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-NET-0005 inside the network engineering standard constellation

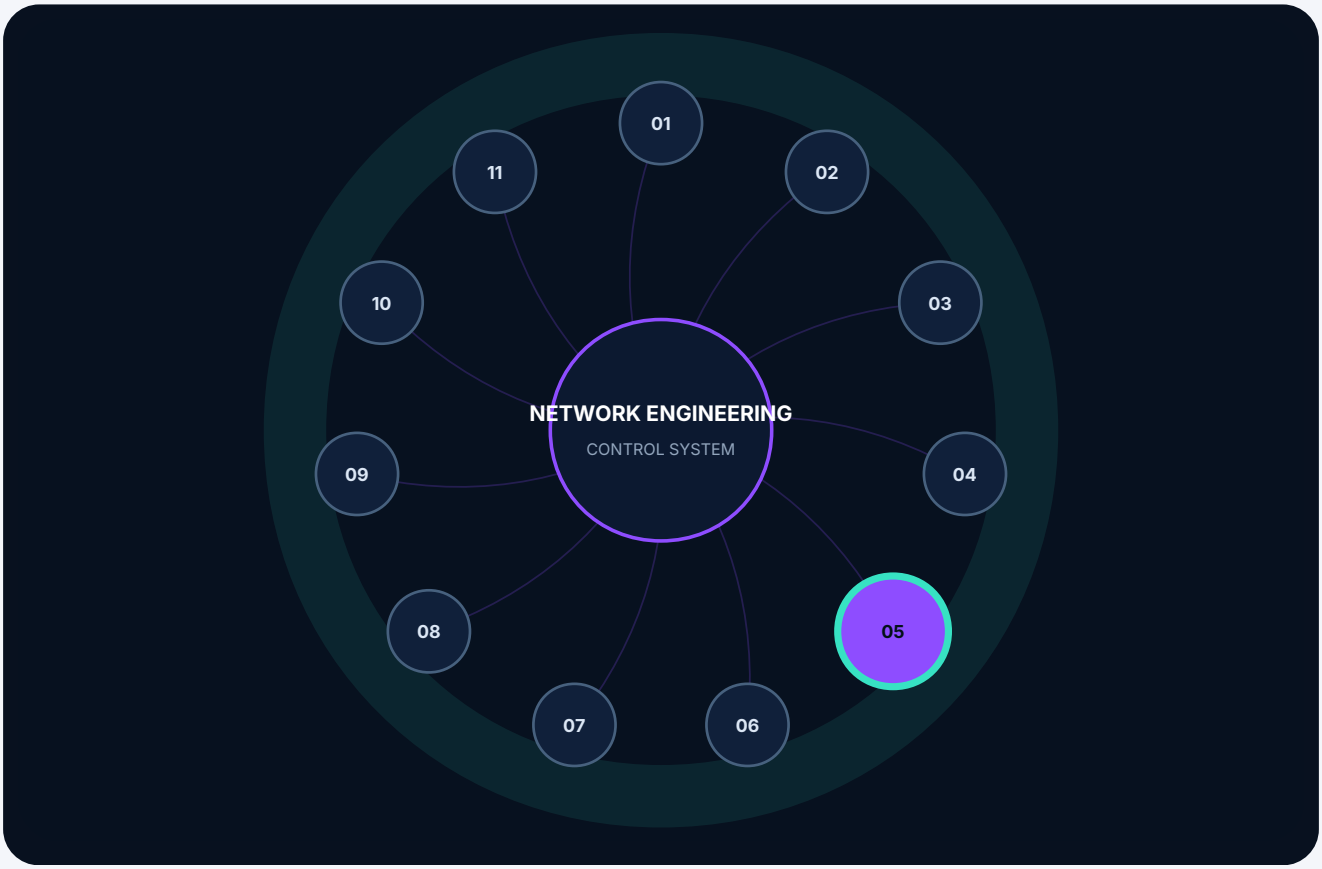

Connected assurance

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

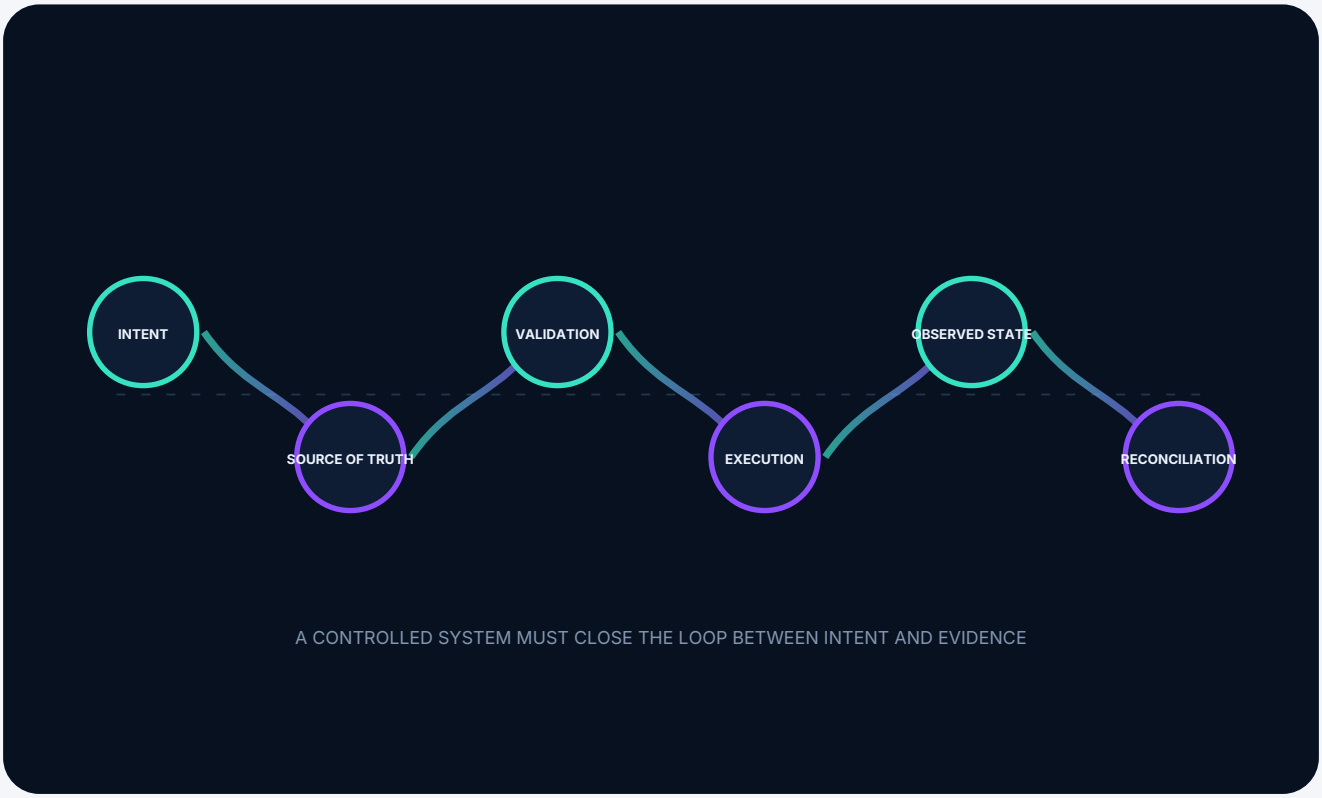

State separation

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-NET-0005

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

The management of network state has failed.

Organizations operate their networks using a fragmented patchwork of spreadsheets, tribal knowledge, and the device's running-configuration as the de facto source of truth. This results in phantom IP allocations, undocumented VLANs, and configuration drift that causes outages and security breaches. The network's actual state is a mystery until something breaks.

This standard exists because:

1. **A Spreadsheet is Not a Database:** Tracking critical infrastructure state in Excel is a liability. It cannot be queried programmatically, it cannot enforce schema, and it cannot be reconciled against reality.
2. **The Device is Not the Source of Truth:** The running-configuration on a switch is the *observed state*, not the *intent*. Treating the device as the SoT means every out-of-band CLI change becomes the new architecture, cementing drift into the fabric.
3. **Drift is a Bug, Not a Feature:** Configuration drift - any divergence between the intended state (SoT/Git) and the observed state (device) - is an operational failure. It MUST be detected, alerted, and automatically remediated.
4. **If It's Not in the SoT, It Doesn't Exist:** Any device, IP, VLAN, or circuit not recorded in the authoritative SoT is unauthorized. The SoT MUST be the absolute prerequisite for any automation, security policy, or simulation.

This document establishes the **Mythos Source-of-Truth & Drift Standard (MSTDS)**, a comprehensive, evidence-based methodology for evaluating the authority, completeness, ingestion fidelity, and reconciliation capabilities of network state management platforms.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Source-of-Truth (SoT) and Source-of-Record (SoR) platforms (NetBox, Nautobot)
- IP Address Management (IPAM) and Data Center Infrastructure Management (DCIM)
- Configuration drift detection and automated remediation
- GitOps reconciliation for network state (desired-state vs. observed-state)
- Real-time state ingestion (gNMI, streaming telemetry, eBPF)
- Data integrity, schema enforcement, and validation rules
- SoT API security and audit logging

1.2 Out of Scope

- Network automation execution engines (covered in MYTHOS-NET-0002)
- Digital twin simulation (covered in MYTHOS-NET-0003)
- General database architecture (covered in MYTHOS-DBS)

1.3 Audience

- Network architects and automation engineers designing state management
- Platform engineers building Network-as-Code pipelines
- Security teams evaluating IPAM and segmentation alignment
- AI governance, risk, and compliance teams
- Standards bodies seeking a reference network state methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 State Management Dimensions

Dimension	Definition
Source-of-Truth (SoT)	The authoritative, programmatically accessible data store that defines the intended state of the network (IPs, VLANs, topologies, contracts).
Observed State	The actual running configuration and operational status of network devices, as reported by the control plane.
Intended State	The state declared in the SoT and version-controlled Git repositories.
Drift	Any divergence between the Intended State and the Observed State.
Reconciliation	The automated process of detecting drift and realigning the observed state with the intended state.

2.2 The State Management Doctrine

> A spreadsheet is a liability. A device is an observer, not an authority. Drift is a bug. If it is not in the SoT, it does not exist. The SoT writes the network; the network never writes the SoT.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; spreadsheets, no SoT, manual state
1	Basic IPAM but incomplete, no reconciliation
2	Functional SoT but manual data entry, no drift remediation
3	Solid production performance; API-driven SoT, automated drift detection
4	Strong performance; continuous reconciliation, GitOps integration
5	Best-in-class; sets the standard for autonomous, cryptographically-verified network state

Weighting

Category	Weight	Rationale
SoT Authority & Schema Integrity	25%	An incomplete or unenforced SoT is useless
State Ingestion & Fidelity	20%	The SoT must reflect reality to be trusted
Drift Detection & Reconciliation	20%	Unremediated drift causes outages
GitOps & Version Control Integration	15%	State must be versioned and peer-reviewed
API Security & Data Governance	10%	The SoT is a critical attack surface
Operational Lifecycle	10%	State management must be continuous

Total: 100%

A system MUST score at least 3.0 in SoT Authority and Drift Detection to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

9 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

SoT Authority and Schema Integrity

SOURCE WEIGHT 25%

4.1.1
Single Source of Truth

4.1.2
Data Model Completeness

MYTHOS-NET-0005-EVAL-4.1.1

Single Source of Truth

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
SoT Authority and Schema Integrity	**4.1.1**	**25%**	**Evidence-led**

Is the SoT the designated authority for network intent and allocation?

0
No centralized SoT; state fragmented across spreadsheets and devices

1
Centralized IPAM but not used by all teams

2
SoT exists but out-of-band changes are common

3
SoT is the required pre-requisite for all automation; no manual IP allocation

4
SoT enforces schema (no orphaned IPs, no invalid VLANs); automation refuses to run without SoT context

5
Leading authority: formally verified schema, cryptographic attestation of state, zero out-of-band allocations, and SoT-as-code

ENGINEERING INTERPRETATION

This criterion evaluates whether single source of truth 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

- source-of-truth objects, Git history, observed-state snapshots, and reconciliation evidence
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. introduce controlled drift and verify detection, adjudication, and restoration
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- multiple authorities, silent manual changes, stale inventory, and destructive auto-remediation
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- drift detection time
- reconciliation success
- unowned object count
- 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-NET-0005-EVAL-4.1.2

Data Model Completeness

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
SoT Authority and Schema Integrity	4.1.2	25%	Evidence-led

Does the SoT capture the full intent of the network (physical, logical, virtual)?

0
Only tracks IP prefixes

1
Tracks IPs and basic device inventory

2
Tracks cables, interfaces, and VLANs

3
Tracks logical topologies (VRFs, overlays, routing policies)

4
Tracks security contracts (service meshes, ACL intent, microsegmentation tags)

5
Leading completeness: full infrastructure graph (physical, logical, security, AI workloads), formally verified referential integrity, and zero unmodeled state

ENGINEERING INTERPRETATION

This criterion evaluates whether data model completeness 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

State Ingestion and Fidelity

SOURCE WEIGHT 20%

4.2.1
Automated Discovery

4.2.2
Fidelity Validation

MYTHOS-NET-0005-EVAL-4.2.1

Automated Discovery

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
State Ingestion and Fidelity	**4.2.1**	**20%**	**Evidence-led**

How is the SoT populated and maintained?

0
Manual data entry

1
Automated network scanning (SNMP/SSH) on schedule

2
API-driven ingestion from automation tools

3
Streaming ingestion via gNMI/streaming telemetry

4
Real-time eBPF-driven state observation for fabric validation

5
Leading ingestion: continuous, real-time, multi-source correlation with cryptographic attestation of observed state

ENGINEERING INTERPRETATION

This criterion evaluates whether automated discovery 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-NET-0005-EVAL-4.2.2

Fidelity Validation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
State Ingestion and Fidelity	4.2.2	20%	Evidence-led

Does the system verify that the SoT matches the actual network?

0
No validation; SoT assumed correct

1
Manual spot checks

2
Automated nightly diff reports

3
Continuous parity validation between SoT and digital twin (Batfish)

4
Automated alerts on drift with suggested remediation

5
Leading fidelity: continuous mathematical proof of SoT-to-device parity, automatic flagging of undocumented state

ENGINEERING INTERPRETATION

This criterion evaluates whether fidelity validation 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

Drift Detection and Reconciliation

SOURCE WEIGHT 20%

4.3.1
Drift Detection Scope

4.3.2
Reconciliation Automation

MYTHOS-NET-0005-EVAL-4.3.1

Drift Detection Scope

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Drift Detection and Reconciliation	4.3.1	20%	Evidence-led

What types of drift does the system detect?

0
No drift detection

1
Detects configuration text changes only

2
Detects operational state changes (BGP neighbors, interface status)

3
Detects semantic drift (e.g., a VLAN exists but is not attached to the correct interface/subnet)

4
Detects intent drift (e.g., a device config complies with syntax but violates security policy)

5
Leading detection: formally verified semantic drift, zero-day policy violations, and cryptographic sealing of intended state

ENGINEERING INTERPRETATION

This criterion evaluates whether drift detection scope 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.

MYTHOS-NET-0005-EVAL-4.3.2

Reconciliation Automation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Drift Detection and Reconciliation	4.3.2	20%	Evidence-led

How is drift remediated?

0
Manual remediation

1
Automated alerts, manual push

2
Automated remediation for specific, low-risk changes

3
Automated remediation for all configuration drift (desired-state push)

4
GitOps reconciliation: drift triggers a PR to revert/overwrite, CI/CD pushes the fix

5
Leading reconciliation: continuous self-healing, formally verified convergence, and zero human intervention for standard drift

ENGINEERING INTERPRETATION

This criterion evaluates whether reconciliation automation 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.

4.4 / CAPABILITY FAMILY

GitOps and Version Control Integration

SOURCE WEIGHT 15%

4.4.1
State-as-Code

MYTHOS-NET-0005-EVAL-4.4.1

State-as-Code

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
GitOps and Version Control Integration	4.4.1	15%	Evidence-led

Is the intended network state version-controlled in Git?

0
No version control

1
Periodic backups to Git

2
Git repository exists, but SoT (NetBox) is the primary interface

3
Git is the Single Source of Truth; SoT is a derived cache (GitOps Pull)

4
Git PRs required for all changes; PR triggers simulation + SoT validation

5
Leading State-as-Code: cryptographic signing of commits, PQC-signed merge gates, immutable history, and formally verified Git-to-SoT-to-Device pipeline

ENGINEERING INTERPRETATION

This criterion evaluates whether state-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

- 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

API Security and Data Governance

SOURCE WEIGHT 10%

4.5.1
SoT Access Control

MYTHOS-NET-0005-EVAL-4.5.1

SoT Access Control

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
API Security and Data Governance	4.5.1	10%	Evidence-led

Is the SoT protected against unauthorized modification?

0 Open API, no auth	**1** Basic API keys	**2** OIDC/OAuth2 with RBAC
3 Granular object-level permissions (RBAC/ABAC)	**4** mTLS with SPIFFE identity for all automation clients	**5** Leading governance: PQC-mTLS, SPIFFE identity, cryptographic audit trails, and zero standing human write access

ENGINEERING INTERPRETATION

This criterion evaluates whether sot access 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

- API inventory, authorization policy, token metadata, gateway logs, and abuse controls
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test object and function authorization, token misuse, replay, expiration, and rate limits
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- bearer tokens without audience limits, hidden APIs, and inconsistent authorization
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- authorized API coverage
- token age
- replay rejection
- abuse-control efficacy
- 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

Operational Lifecycle

SOURCE WEIGHT 10%

4.6.1
State Hygiene

MYTHOS-NET-0005-EVAL-4.6.1

State Hygiene

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Operational Lifecycle	4.6.1	10%	Evidence-led

Does the system enforce hygiene (no stale data, no orphans)?

0
Stale data accumulates indefinitely

1
Manual cleanup

2
Automated alerts for stale objects

3
Automated garbage collection for orphaned objects

4
Strict referential integrity prevents orphaned objects from being created

5
Leading hygiene: formally verified referential integrity, automated lifecycle management (provisioning → active → decommissioned), and zero stale state

ENGINEERING INTERPRETATION

This criterion evaluates whether state hygiene 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 State Management Suite (MSTMS)

Traditional network benchmarks measure protocol uptime. The MSTMS evaluates SoT authority, fidelity, and reconciliation.

5.1 Suite Composition

5.1.1 MSTMS-Authority

- **Shadow IT Injection:** 50+ tests adding devices/IPs directly to the network without SoT entry; verify automation ignores/quarantines them.
- **Schema Validation:** 100+ tests attempting to create invalid objects (overlapping IPs, missing relations); verify SoT rejects them.

5.1.2 MSTMS-Reconciliation

- **Config Drift Injection:** 100+ tests modifying device configs out-of-band; verify system detects and auto-remediates.
- **State Fidelity:** 50+ tests verifying SoT topology matches digital twin topology.

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.

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

Segment Routing and Next-Generation Transport Standard

Defines SR-MPLS, SRv6, policy, traffic engineering, operations, migration, and security.

PERMANENT DOCUMENT ID
MYTHOS-NET-0006

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
10 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

Segment Routing and Next-Generation Transport Standard

Defines SR-MPLS, SRv6, policy, traffic engineering, operations, migration, and security.

DOCUMENT ID
MYTHOS-NET-0006

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

10	6	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-NET-0006 inside the network engineering standard constellation

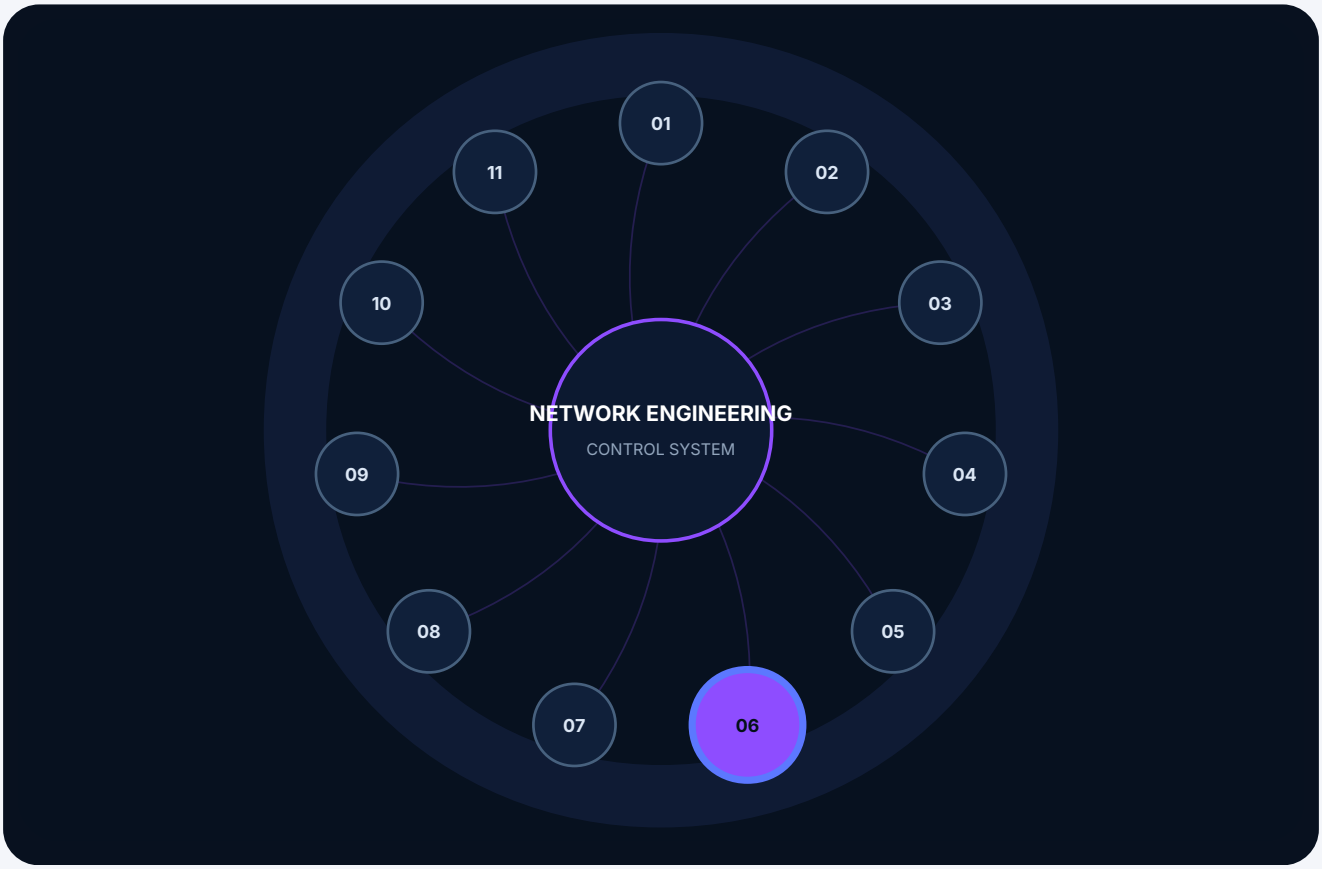

Connected assurance

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

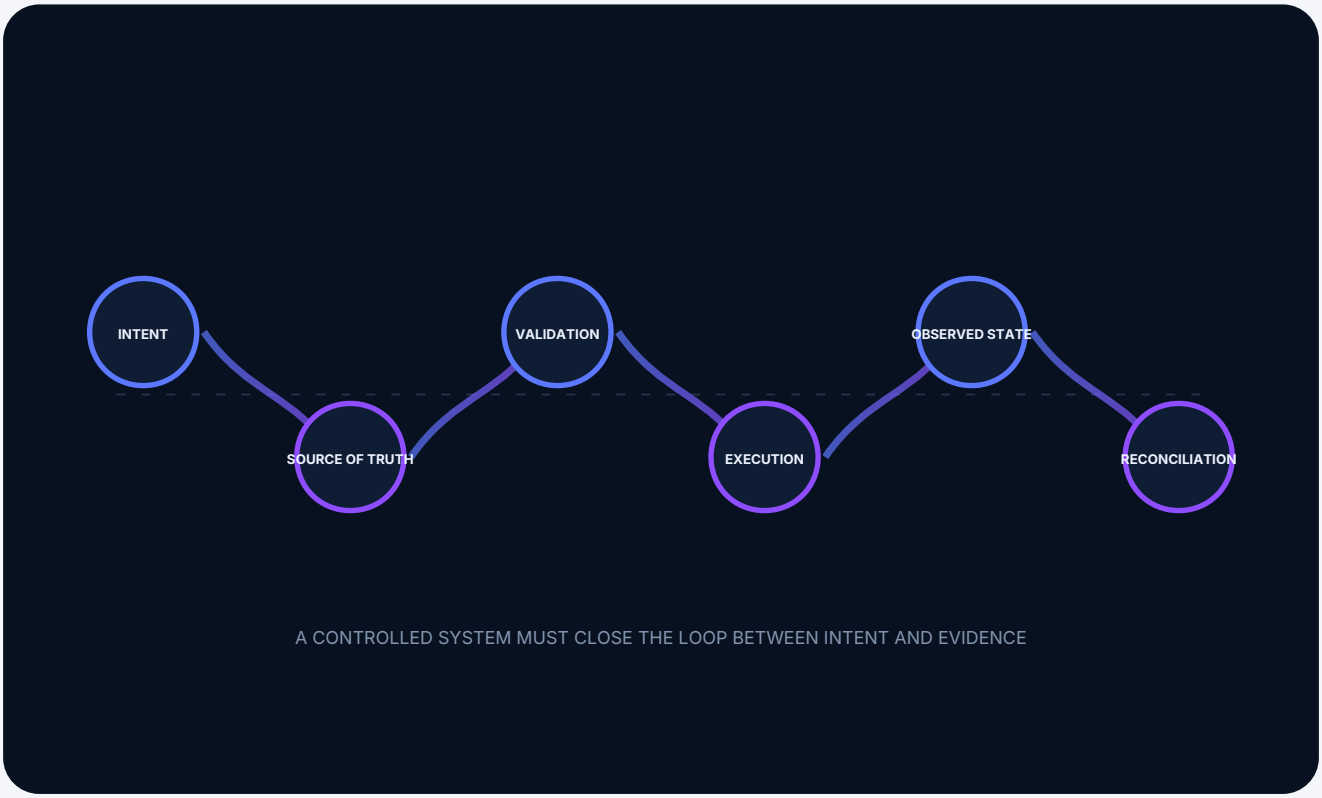

State separation

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-NET-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

ENGINEERING DOCTRINE

0. Executive Mandate

The transport layer has failed to keep pace with the demands of distributed AI and sovereign infrastructure.

Legacy transport architectures - built on LDP, RSVP-TE, and VLAN-based overlays - require massive per-hop state, scale poorly, and rely on ECMP hashing that fundamentally breaks when applied to AI/GPU distributed training flows (RDMA/NCCL). "Hoping" that a hash algorithm spreads traffic evenly across parallel links is not engineering; it is operational gambling.

This standard exists because:

1. **Stateful Core is Dead:** Maintaining per-flow state in every core router (RSVP-TE) does not scale. Segment Routing moves state to the ingress edge, encoding the path in the packet header (IPv6 Segment Routing Header). The core remains stateless.
2. **ECMP is Insufficient for AI:** Equal-Cost Multi-Path hashing cannot guarantee ordered delivery or deterministic latency for GPU collectives. SRv6 allows explicit, programmable path steering, ensuring flows traverse the exact rail-optimized topology required.
3. **Network Programming is Mandatory:** SRv6 is not just a routing protocol; it is a network programming language. SIDs (Segment IDs) can represent topological instructions, service functions, or appliance bypasses, enabling stateless service chaining without VLAN stitching.
4. **Resilience Must Be Sub-50ms:** Legacy convergence (seconds) causes TCP timeouts and RDMA disconnects. Topology-Independent Loop-Free Alternate (TI-LFA) provides sub-50ms failover, mathematically guaranteeing loop-free paths during convergence.

This document establishes the **Mythos Next-Gen Transport Standard (MNGTS)**, a comprehensive, evidence-based methodology for evaluating SRv6, traffic engineering, and deterministic transport capabilities.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Segment Routing over IPv6 (SRv6) and SR-MPLS architectures
- SRv6 Network Programming (SIDs, Behaviors, SR Policies)
- Traffic Engineering (TE) for AI/GPU fabrics and multi-cloud
- Fast Reroute and Resilience (TI-LFA, Micro-loop avoidance)
- Deterministic Networking (DetNet) for latency-sensitive workloads
- Service Function Chaining (SFC) via SRv6
- Hardware offloading (DPU/SmartNIC) for SRv6 encapsulation

1.2 Out of Scope

- General network architecture topology (covered in MYTHOS-NET-0001)
- Network automation and configuration (covered in MYTHOS-NET-0002)
- General zero-trust networking (covered in MYTHOS-NET-0004)

1.3 Audience

- Network architects designing spine-leaf and multi-cloud transports
- AI/ML infrastructure engineers managing GPU cluster networks
- Platform engineers building SRv6-based service meshes
- Standards bodies seeking a reference next-gen transport methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Transport Dimensions

Dimension	Definition
SRv6	Segment Routing over IPv6. A source-routing architecture that encodes the path and instructions in an IPv6 extension header (SRH), eliminating the need for per-flow state in the core.
SID	Segment Identifier. A 128-bit IPv6 address in SRv6 that represents a topological instruction (e.g., "go through node X") or a service instruction (e.g., "apply firewall").
TI-LFA	Topology-Independent Loop-Free Alternate. A fast-reroute mechanism that guarantees sub-50ms failover with zero loops, regardless of topology.
SR Policy	A defined path (list of SIDs) applied at the ingress node to steer traffic, replacing traditional RSVP-TE tunnels.
Network Programming	The SRv6 capability to combine topological and functional SIDs to create complex network behaviors without per-hop state.

2.2 The Transport Doctrine

> Stateful core is obsolete. ECMP is hoping; SRv6 is proving. A path must be programmable. A failure must be invisible to the application. Transport must serve the workload, not the other way around.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; legacy LDP/RSVP-TE, ECMP only
1	SR-MPLS deployed but no SRv6 or TE
2	Functional SRv6 but limited hardware support or TE
3	Solid production performance; SRv6, TI-LFA, basic TE
4	Strong performance; network programming, AI-driven TE, hardware offload
5	Best-in-class; sets the standard for deterministic, programmable transport

Weighting

Category	Weight	Rationale
SRv6 Architecture & Programmability	25%	Source routing and SIDs replace legacy stateful overlays
Traffic Engineering & AI Fabric Optimization	20%	ECMP fails AI workloads; explicit path steering is mandatory
Resilience & Sub-50ms Convergence (TI-LFA)	20%	RDMA/AI workloads cannot tolerate seconds of convergence
Multi-Cloud & Service Chaining (SFC)	15%	Stateless SFC replaces VLAN-stitching nightmares
Hardware Offload & Dataplane Efficiency	10%	SRv6 encapsulation requires line-rate hardware support
Operational Lifecycle & Telemetry	10%	Transport must be observable and governed continuously

Total: 100%

A system MUST score at least 3.0 in SRv6 Architecture and Resilience to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

10 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

SRv6 Architecture and Programmability

SOURCE WEIGHT 25%

4.1.1
Network Programming

4.1.2
Underlay Simplification

MYTHOS-NET-0006-EVAL-4.1.1

Network Programming

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
SRv6 Architecture and Programmability	**4.1.1**	**25%**	**Evidence-led**

Does the transport support full SRv6 network programming (End, End.X, End.T, End.DX2, etc.)?

0
No SRv6 support; LDP/RSVP-TE only

1
SR-MPLS only; no SRv6 data plane

2
Basic SRv6 (End, End.X) for topological routing only

3
Full SRv6 network programming (L3, L2, Service SIDs)

4
Custom/proprietary SID behaviors for DPU/SmartNIC offload

5
Leading programmability: formally verified SID behaviors, zero state in core, universal hardware support, and AI-generatable SR policies

ENGINEERING INTERPRETATION

This criterion evaluates whether network programming 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-NET-0006-EVAL-4.1.2

Underlay Simplification

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
SRv6 Architecture and Programmability	4.1.2	25%	Evidence-led

Has SRv6 eliminated the need for LDP, RSVP-TE, and NVGRE/VXLAN?

0
Full LDP/RSVP-TE stack still running

1
SR-MPLS replaces LDP; VXLAN still used for overlays

2
SRv6 replaces LDP; VXLAN still used

3
SRv6 provides both underlay routing and overlay (L2VPN/L3VPN) via SIDs

4
Complete elimination of VXLAN/NVGRE; SRv6 natively encodes tenant info

5
Leading simplification: single protocol (IS-IS/OSPFv3 for SRv6), zero legacy control planes, natively integrated overlay/underlay

ENGINEERING INTERPRETATION

This criterion evaluates whether underlay simplification 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

Traffic Engineering and AI Fabric Optimization

SOURCE WEIGHT 20%

4.2.1
Explicit Path Steering

4.2.2
GPU/Rail-Optimized Routing

MYTHOS-NET-0006-EVAL-4.2.1

Explicit Path Steering

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Traffic Engineering and AI Fabric Optimization	**4.2.1**	**20%**	**Evidence-led**

Can the system steer traffic along explicit, non-ECMP paths for AI workloads?

0
ECMP hashing only; no explicit path control

1
Static routing for specific flows

2
RSVP-TE tunnels for path control (stateful, unscalable)

3
SR Policies with explicit SID lists for topological steering

4
Dynamic SR Policies adjusted by AI/ML based on real-time congestion

5
Leading TE: AI-driven real-time path computation, per-flow steering for GPU collectives, zero ECMP polarization, formally verified bandwidth allocation

ENGINEERING INTERPRETATION

This criterion evaluates whether explicit path steering 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-NET-0006-EVAL-4.2.2

GPU/Rail-Optimized Routing

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Traffic Engineering and AI Fabric Optimization	4.2.2	20%	Evidence-led

Does the transport align with GPU rail topologies to prevent NCCL stalls?

0
Unaware of server/NIC internal topology

1
Application-level routing hints

2
Network-level static paths per rail

3
Automated SR Policies mapping GPU rails to network paths

4
Integration with NCCL/Slurm for dynamic path provisioning

5
Leading alignment: intent-driven rail routing, sub-microsecond re-balancing, formally verified deadlock freedom

ENGINEERING INTERPRETATION

This criterion evaluates whether gpu/rail-optimized routing 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

- routing policy, RIB/FIB snapshots, adjacency state, and path-analysis output
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. simulate route withdrawal, policy conflict, convergence, and rollback
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- route leaks, asymmetric paths, stale advertisements, and unbounded convergence
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- route-policy compliance
- convergence time
- unexpected path count
- 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

Resilience and Sub-50ms Convergence (TI-LFA)

SOURCE WEIGHT 20%

4.3.1
Fast Reroute

4.3.2
Micro-Loop Prevention

MYTHOS-NET-0006-EVAL-4.3.1

Fast Reroute

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Resilience and Sub-50ms Convergence (TI-LFA)	**4.3.1**	**20%**	**Evidence-led**

Does the transport guarantee sub-50ms failover for all failures?

0
Standard IGP convergence (seconds)

1
LFA/RLFA for some topologies

2
TI-LFA implemented but >50ms in practice

3
TI-LFA <50ms for link failures

4
TI-LFA <50ms for link and node failures; micro-loop avoidance

5
Leading resilience: mathematically proven <10ms TI-LFA, zero packet loss for RDMA flows, formally verified post-convergence paths

ENGINEERING INTERPRETATION

This criterion evaluates whether fast reroute 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-NET-0006-EVAL-4.3.2

Micro-Loop Prevention

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Resilience and Sub-50ms Convergence (TI-LFA)	**4.3.2**	**20%**	**Evidence-led**

Does the system prevent transient micro-loops during convergence?

0
Micro-loops occur frequently

1
Temporary micro-loops tolerated

2
Basic micro-loop avoidance mechanisms

3
SRv6 specific mechanisms (e.g., SR policy wait-timers)

4
Data-plane micro-loop avoidance (ordered FIB updates)

5
Leading prevention: formally verified loop-free convergence, zero transient drops

ENGINEERING INTERPRETATION

This criterion evaluates whether micro-loop 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.

4.4 / CAPABILITY FAMILY

Multi-Cloud and Service Chaining (SFC)

SOURCE WEIGHT 15%

4.4.1
Stateless Service Chaining

4.4.2
Inter-Domain Stitching

MYTHOS-NET-0006-EVAL-4.4.1

Stateless Service Chaining

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Multi-Cloud and Service Chaining (SFC)	**4.4.1**	**15%**	**Evidence-led**

Can the system perform service chaining (firewalls, load balancers) without VLAN stitching or stateful overlays?

0
Manual VLAN/subnet stitching for services

1
Static service graphs via routing

2
NSX/Service Mesh-based chaining (L7 overhead)

3
SRv6 SFC using Adjacency SIDs (End.X) for topology steering

4
SRv6 SFC using Service SIDs (End.AD) for stateless, transparent insertion

5
Leading SFC: dynamically programmable SRv6 chains, zero topology modification, DPU-offloaded service execution

ENGINEERING INTERPRETATION

This criterion evaluates whether stateless service chaining 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-NET-0006-EVAL-4.4.2

Inter-Domain Stitching

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Multi-Cloud and Service Chaining (SFC)	4.4.2	15%	Evidence-led

Can SR Policies span multiple domains (DC, WAN, Cloud) seamlessly?

0
Manual BGP peering and route-leaking

1
Static inter-domain tunnels

2
BGP-LU (Labeled Unicast) for SR-MPLS

3
BGP EPE (Egress Peer Engineering) + SRv6 SIDs for inter-domain

4
Seamless inter-domain SR Policies with automated SID stitching

5
Leading stitching: zero-touch multi-cloud SR Policies, formally verified end-to-end SLAs, and PQC-encrypted inter-domain SIDs

ENGINEERING INTERPRETATION

This criterion evaluates whether inter-domain stitching 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

Hardware Offload and Dataplane Efficiency

SOURCE WEIGHT 10%

4.5.1
Line-Rate SRv6

MYTHOS-NET-0006-EVAL-4.5.1

Line-Rate SRv6

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Hardware Offload and Dataplane Efficiency	4.5.1	10%	Evidence-led

Does the hardware support SRv6 encapsulation/decapsulation at line rate?

0
Software switching only

1
Hardware support for SR-MPLS only

2
Hardware support for SRv6 but with deep packet inspection penalties

3
Full line-rate SRv6 (End, End.X) on standard ASICs

4
Line-rate SRv6 with SmartNIC/DPU offload for SID processing

5
Leading efficiency: DPU-native SRv6, zero CPU overhead, SRv6 header compression (e.g., uSID) at line rate

ENGINEERING INTERPRETATION

This criterion evaluates whether line-rate srv6 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.

4.6 / CAPABILITY FAMILY

Operational Lifecycle and Telemetry

SOURCE WEIGHT 10%

4.6.1
Transport Telemetry

MYTHOS-NET-0006-EVAL-4.6.1

Transport Telemetry

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Operational Lifecycle and Telemetry	4.6.1	10%	Evidence-led

Is SRv6 policy state observable via streaming telemetry?

0
SNMP only

1
CLI show commands

2
Model-driven telemetry for SRv6 routes

3
Per-SID/per-Policy traffic and latency metrics (gNMI)

4
Real-time path tracing via SRv6 IOAM (In-Situ OAM)

5
Leading observability: deterministic IOAM, eBPF-validated path tracing, and AI-driven anomaly detection on SID metrics

ENGINEERING INTERPRETATION

This criterion evaluates whether transport telemetry 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 SRv6 Suite (M-SRv6-S)

Traditional transport benchmarks measure throughput and ECMP utilization. The M-SRv6-S evaluates programmability, TE, and deterministic resilience.

5.1 Suite Composition

5.1.1 M-SRv6-TE

- **GPU Flow Steering:** 50+ tests verifying AI/NCCL flows follow explicit SR Policies rather than ECMP hashing.
- **Dynamic Rebalancing:** 50+ tests injecting congestion to verify AI-driven TE adjusts SR Policies.

5.1.2 M-SRv6-Resilience

- **TI-LFA Failover:** 100+ link/node failure injections verifying <50ms convergence and zero RDMA disconnects.
- **Micro-Loop Check:** 50+ IGP flaps verifying zero transient loops during convergence.

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.

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

5G, 6G, RAN Automation, and Network Slicing Standard

Defines mobile architecture, orchestration, slicing, isolation, QoS, telemetry, and research boundaries.

PERMANENT DOCUMENT ID
MYTHOS-NET-0007

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
9 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

5G, 6G, RAN Automation, and Network Slicing Standard

Defines mobile architecture, orchestration, slicing, isolation, QoS, telemetry, and research boundaries.

DOCUMENT ID
MYTHOS-NET-0007

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

9	6	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-NET-0007 inside the network engineering standard constellation

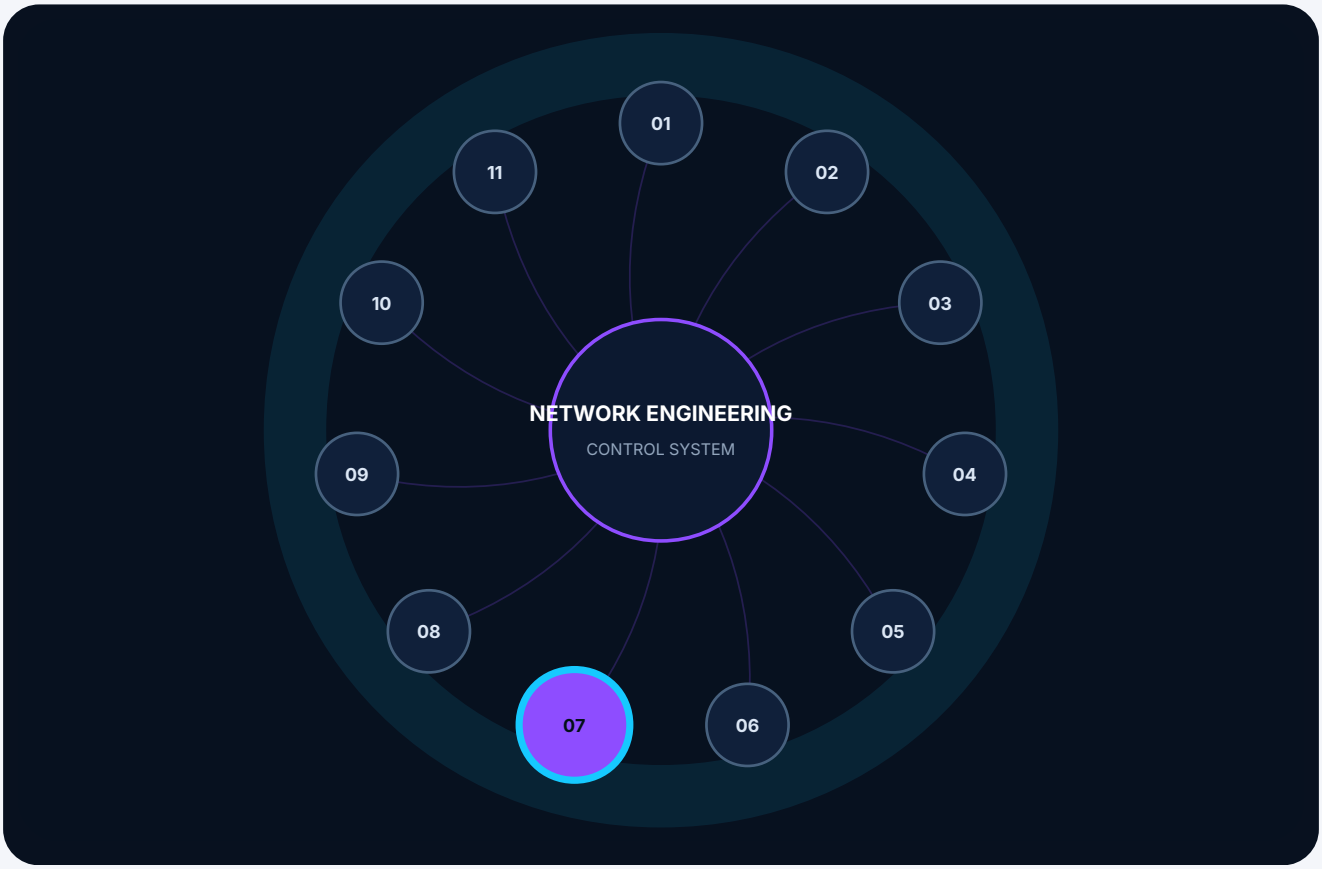

Connected assurance

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

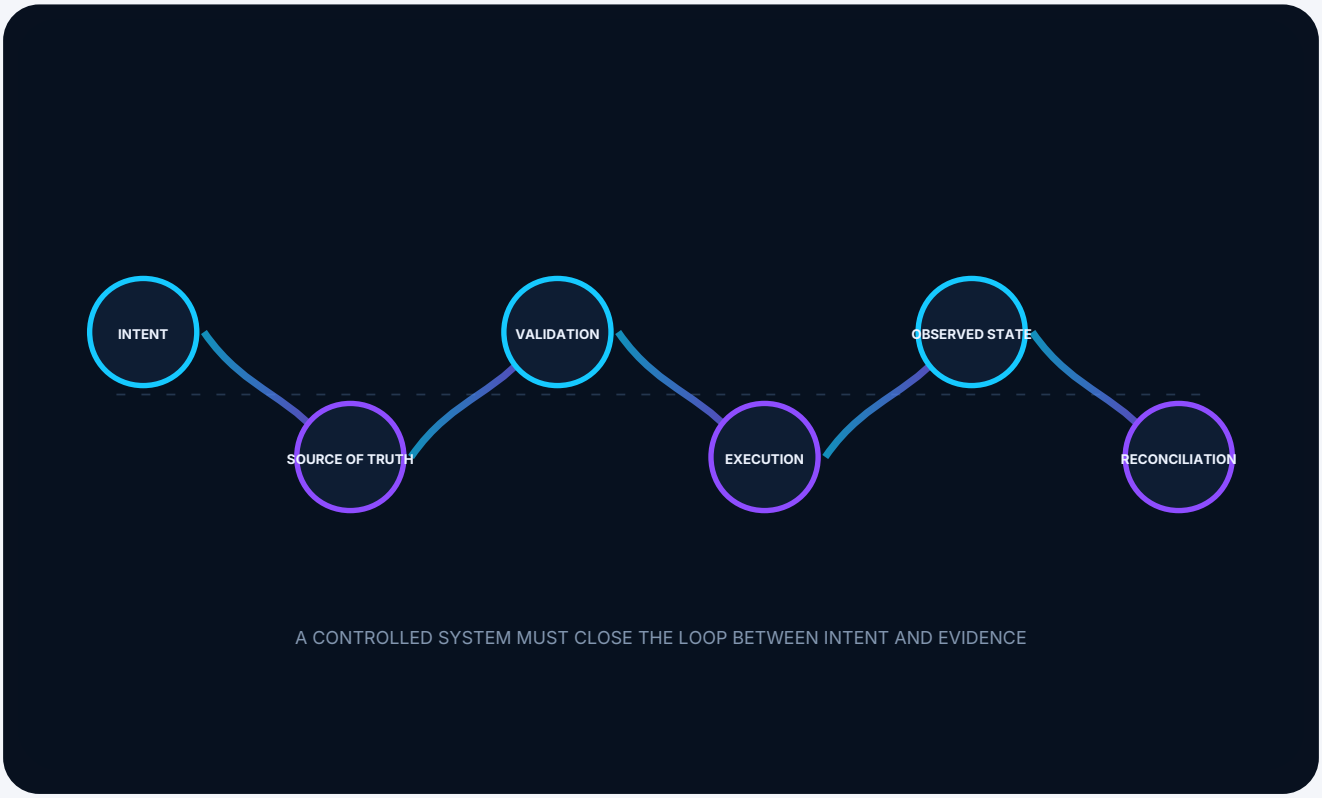

State separation

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-NET-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

The deployment of 5G and emerging 6G networks has failed to deliver on the promise of programmable wireless infrastructure.

Telecom operators have built 5G networks as fatter 4G pipes, utilizing standalone macro cells and monolithic, vendor-locked RAN stacks. Network slicing remains a theoretical construct confined to core network (5GC) QoS markers, completely ignoring the radio resource bottleneck. Best-effort wireless is fundamentally incompatible with the deterministic latency requirements of autonomous industrial systems (IIoT), real-time robotics, and sovereign edge AI.

This standard exists because:

1. **Best-Effort Wireless is a Liability:** Marking QCI/5QI values in the 5G Core does not guarantee radio resources. If the RAN is congested, URLLC packets drop. Slicing MUST extend all the way to the scheduler on the radio cell.
2. **Monolithic RAN is Vendor Lock-in:** Proprietary baseband units (BBUs) and remote radio heads (RRHs) from a single vendor stifle innovation and prevent the deployment of AI-driven radio optimization. O-RAN (Open RAN) decouples hardware from software, enabling RAN Intelligent Controllers (RICs) to optimize the radio in real-time.
3. **Edge AI Requires Deterministic Uplink:** AI inference at the edge (MEC) is useless if the sensor data cannot reach the edge node within strict latency bounds. 6G and 5G-Advanced must provide sub-millisecond jitter and 99.999% reliability for URLLC.
4. **Spectrum is a Finite Sovereign Asset:** Dynamic Spectrum Sharing (DSS) and CBRS-style shared spectra require AI-driven allocation to prevent interference and maximize utilization.

This document establishes the **Mythos RAN & Slicing Standard (MRSS)**, a comprehensive, evidence-based methodology for evaluating 5G/6G networks against slicing isolation, O-RAN programmability, and deterministic latency criteria.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- 5G/6G Network Slicing (Core, Transport, and RAN slicing)
- Open RAN (O-RAN) architecture (O-CU, O-DU, O-RU, SMO)
- RAN Intelligent Controllers (Non-RT RIC, Near-RT RIC, xApps, rApps)
- Ultra-Reliable Low-Latency Communication (URLLC) determinism
- Multi-access Edge Computing (MEC) and UPF breakout integration
- Spectrum management and Dynamic Spectrum Sharing (DSS)

1.2 Out of Scope

- General transport architecture (covered in MYTHOS-NET-0006)
- General edge compute sandboxing (covered in MYTHOS-EMT-0001)
- Broad telecom billing and roaming (reserved for future MYTHOS-TEL domain)

1.3 Audience

- Network architects designing private 5G/6G and O-RAN fabrics
- IIoT and operational technology (OT) engineers requiring deterministic wireless
- AI/ML engineers building RIC xApps/rApps
- Defense and sovereign infrastructure operators
- Standards bodies seeking a reference RAN/slicing methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 RAN and Slicing Dimensions

Dimension	Definition
Network Slice	An end-to-end logical network spanning Core (5GC), Transport, and RAN, providing specific QoS (e.g., URLLC, eMBB, mMTC) with isolated resources.
O-RAN	Open Radio Access Network. An architecture that disaggregates the RAN into O-CU (Centralized Unit), O-DU (Distributed Unit), and O-RU (Radio Unit) with open fronthaul interfaces.
RIC	RAN Intelligent Controller. An AI/ML-driven orchestration layer (Near-RT for sub-ms scheduling, Non-RT for policy) that optimizes RAN performance via xApps and rApps.
URLLC	Ultra-Reliable Low-Latency Communication. A 5G service category targeting sub-1ms latency and 99.999% reliability.
MEC	Multi-access Edge Computing. Hosting compute/AI resources at the cellular base station or aggregation site for local data breakout.

2.2 The Wireless Doctrine

> A pipe is not a platform. Best-effort is not a slice. Proprietary RAN is vendor lock-in. A slice that does not reserve radio resources is a VLAN. Wireless must be deterministic.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; monolithic RAN, best-effort only, no slicing
1	Basic 5GC QoS, but no RAN slicing or O-RAN
2	Functional slicing but lacks RAN isolation or RIC automation
3	Solid production performance; O-RAN, RIC, RAN slicing
4	Strong performance; deterministic URLLC, AI-driven RIC
5	Best-in-class; sets the standard for autonomous, 6G-ready wireless

Weighting

Category	Weight	Rationale
Network Slicing & Radio Isolation	25%	A slice without radio resources is a lie
O-RAN & RIC Automation	20%	Monolithic RAN cannot scale or optimize dynamically
URLLC & Deterministic Latency	20%	IIoT/Autonomy requires mathematical latency bounds
Edge Integration (MEC/UPF)	15%	Edge AI requires local breakout
Spectrum & Interference Management	10%	Shared spectrum requires AI governance
Operational Lifecycle & Security	10%	RIC/xApps are a new attack surface

Total: 100%

A system MUST score at least 3.0 in Network Slicing and URLLC to be considered for production use in regulated/industrial environments.

EVALUATION AND ASSURANCE SYSTEM

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

9 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Network Slicing and Radio Isolation

SOURCE WEIGHT 25%

4.1.1
End-to-End Slice Isolation

4.1.2
Slice Lifecycle Management

MYTHOS-NET-0007-EVAL-4.1.1

End-to-End Slice Isolation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Network Slicing and Radio Isolation	**4.1.1**	**25%**	**Evidence-led**

Does the slice guarantee isolated resources from the Core through the Transport to the RAN scheduler?

0
No slicing; best-effort QoS marking only

1
5GC slicing (UPF/AMF isolation) but shared RAN scheduler

2
Soft RAN slicing (QoS prioritization, but no hard resource reservation)

3
Hard RAN slicing (dedicated PRBs/ Resource Blocks per slice)

4
Hard RAN slicing + Transport slicing (SRv6/I-SID) + Core isolation

5
Leading isolation: mathematically proven end-to-end resource isolation, zero cross-slice interference under congestion, formally verified SLAs

ENGINEERING INTERPRETATION

This criterion evaluates whether end-to-end slice 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

- source-of-truth objects, Git history, observed-state snapshots, and reconciliation evidence
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. introduce controlled drift and verify detection, adjudication, and restoration
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- multiple authorities, silent manual changes, stale inventory, and destructive auto-remediation
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- drift detection time
- reconciliation success
- unowned object count
- 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-NET-0007-EVAL-4.1.2

Slice Lifecycle Management

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Network Slicing and Radio Isolation	**4.1.2**	**25%**	**Evidence-led**

Can slices be instantiated, modified, and terminated programmatically?

0
Manual provisioning

1
Template-based core slicing via OSS

2
API-driven core and transport slicing

3
Automated E2E slice provisioning including RAN parameters

4
Dynamic slice scaling based on real-time demand (RIC-driven)

5
Leading lifecycle: intent-driven slice creation, formally verified SLA compliance, AI-driven dynamic re-slicing in sub-second intervals

ENGINEERING INTERPRETATION

This criterion evaluates whether slice lifecycle management 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

O-RAN and RIC Automation

SOURCE WEIGHT 20%

4.2.1
RIC Architecture

4.2.2
xApp/rApp Ecosystem

MYTHOS-NET-0007-EVAL-4.2.1

RIC Architecture

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
O-RAN and RIC Automation	4.2.1	20%	Evidence-led

Does the network implement an O-RAN compliant RAN Intelligent Controller?

0
Monolithic vendor RAN; no RIC

1
Vendor-specific AI optimization (black box)

2
Non-RT RIC implemented for policy and rApps

3
Near-RT RIC implemented for sub-ms scheduling and xApps

4
Full O-RAN stack (SMO, Non-RT, Near-RT) with open xApp/rApp ecosystem

5
Leading RIC: formally verified xApp/rApp sandboxing (WASM), AI-driven intent translation, and zero vendor lock-in

ENGINEERING INTERPRETATION

This criterion evaluates whether ric architecture 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-NET-0007-EVAL-4.2.2

xApp/rApp Ecosystem

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
O-RAN and RIC Automation	4.2.2	20%	Evidence-led

Can third-party or custom AI applications (xApps/rApps) be safely deployed to the RIC?

0
No third-party extensibility

1
Vendor-approved apps only

2
Side-loaded apps with no isolation

3
Containerized apps with resource limits

4
WASM-sandboxed xApps with AIBOM and behavioral verification

5
Leading ecosystem: formally verified sandboxing, zero resource starvation, cryptographic attestation of app provenance, and automated rollback

ENGINEERING INTERPRETATION

This criterion evaluates whether xapp/rapp ecosystem 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

URLLC and Deterministic Latency

SOURCE WEIGHT 20%

4.3.1
Latency Guarantees

4.3.2
Time Sensitive Networking (TSN) Integration

MYTHOS-NET-0007-EVAL-4.3.1

Latency Guarantees

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
URLLC and Deterministic Latency	4.3.1	20%	Evidence-led

Can the network mathematically guarantee sub-ms latency for URLLC slices?

0
No latency guarantees; statistical best-effort

1
Target latency defined, but unverified under load

2
Latency verified in lab conditions

3
Latency verified under simulated congestion (digital twin)

4
Deterministic scheduling (e.g., 5G TSN integration, mini-slots)

5
Leading URLLC: formally verified sub-ms determinism, 99.999% reliability, zero jitter under adversarial load, and hardware-accelerated scheduling

ENGINEERING INTERPRETATION

This criterion evaluates whether latency guarantees 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-NET-0007-EVAL-4.3.2

Time Sensitive Networking (TSN) Integration

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
URLLC and Deterministic Latency	4.3.2	20%	Evidence-led

Does the 5G network integrate with industrial TSN networks for deterministic wired/wireless convergence?

0
No TSN awareness

1
Basic TSN translation at the edge

2
5G TSN Translator (DS-TT/NW-TT) compliant

3
Integrated TSN scheduling (802.1Qbv) with 5G slot coordination

4
Sub-microsecond time synchronization (IEEE 1588/802.1AS) across wireless and wired

5
Leading convergence: formally verified end-to-end determinism, zero timing drift, and AI-driven schedule optimization

ENGINEERING INTERPRETATION

This criterion evaluates whether time sensitive networking (tsn) 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

- RF survey, channel and power plan, authentication logs, and client telemetry
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test roaming, capacity, interference, authentication, and degraded-band behavior
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- coverage holes, sticky clients, co-channel interference, and insecure fallback
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- coverage compliance
- authentication success
- roaming interruption
- airtime utilization
- 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

Edge Integration (MEC/ UPF)

SOURCE WEIGHT 15%

4.4.1
UPF Breakout

MYTHOS-NET-0007-EVAL-4.4.1

UPF Breakout

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Edge Integration (MEC/ UPF)	4.4.1	15%	Evidence-led

Can traffic be selectively broken out at the edge (MEC) without traversing the core?

0
All traffic hairpins to central core

1
Static UPF at regional data centers

2
Dynamic UPF selection at the edge

3
UL CL (Uplink Classifier) for per-flow local breakout

4
Integrated MEC platform with local AI inference and RAN awareness

5
Leading edge: RAN-aware dynamic UPF breakout, sub-ms local routing, WASM-sandboxed edge AI, and sovereign data compliance enforcement

ENGINEERING INTERPRETATION

This criterion evaluates whether upf breakout 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

Spectrum and Interference Management

SOURCE WEIGHT 10%

4.5.1
Dynamic Spectrum Sharing

MYTHOS-NET-0007-EVAL-4.5.1

Dynamic Spectrum Sharing

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Spectrum and Interference Management	**4.5.1**	**10%**	**Evidence-led**

How efficiently is shared spectrum (e.g., CBRS) managed?

0
Static spectrum allocation only

1
Manual spectrum reconfiguration

2
SAS (Spectrum Access Server) compliant for CBRS

3
AI-driven spectrum sensing and allocation (RIC rApp)

4
Sub-ms dynamic spectrum sharing (DSS) between 4G/5G/6G

5
Leading management: AI-driven cognitive radio, formally verified interference avoidance, and zero spectrum waste

ENGINEERING INTERPRETATION

This criterion evaluates whether dynamic spectrum sharing 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.

4.6 / CAPABILITY FAMILY

Operational Lifecycle and Security

SOURCE WEIGHT 10%

4.6.1
RAN Security

MYTHOS-NET-0007-EVAL-4.6.1

RAN Security

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Operational Lifecycle and Security	4.6.1	10%	Evidence-led

Is the O-RAN fronthaul and RIC protected from supply chain and runtime attacks?

0
Unauthenticated fronthaul (eCPRI) and RIC APIs

1
TLS on management interfaces

2
mTLS on fronthaul and RIC interfaces

3
Signed firmware/images for O-CU/O-DU/O-RU; xApp signing

4
Hardware roots of trust in O-RU; PQC-ready control plane

5
Leading security: formally verified fronthaul integrity, PQC-mTLS, behavioral monitoring of xApps, and zero unauthenticated radio commands

ENGINEERING INTERPRETATION

This criterion evaluates whether ran security 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 RAN & Slicing Suite (MRSS)

Traditional RAN benchmarks measure throughput and coverage. The MRSS evaluates slicing isolation, deterministic latency, and RIC programmability.

5.1 Suite Composition

5.1.1 MRSS-Slicing

- **Cross-Slice Interference:** 50+ tests flooding eMBB slice to verify URLLC slice latency and packet loss remain within SLA.
- **E2E Provisioning:** 50+ tests automating slice creation via API, verifying RAN, Transport, and Core alignment.

5.1.2 MRSS-URLLC

- **Deterministic Jitter:** 100+ tests injecting RAN interference to verify sub-ms jitter bounds hold.
- **TSN Convergence:** 50+ tests verifying 1588 timing sync across wired/wireless boundaries.

5.1.3 MRSS-RIC

- **xApp Sandboxing:** 50+ tests deploying malicious/rogue xApps to verify WASM isolation and resource limits.
- **AI Optimization:** 50+ tests verifying RIC xApps improve spectral efficiency without violating slice SLAs.

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.

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

OSI Troubleshooting and Core Infrastructure Standard

Defines structured troubleshooting, telemetry, packet analysis, isolation, validation, and closure evidence.

PERMANENT DOCUMENT ID
MYTHOS-NET-0008

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
6 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

OSI Troubleshooting and Core Infrastructure Standard

Defines structured troubleshooting, telemetry, packet analysis, isolation, validation, and closure evidence.

DOCUMENT ID
MYTHOS-NET-0008

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

6	6	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-NET-0008 inside the network engineering standard constellation

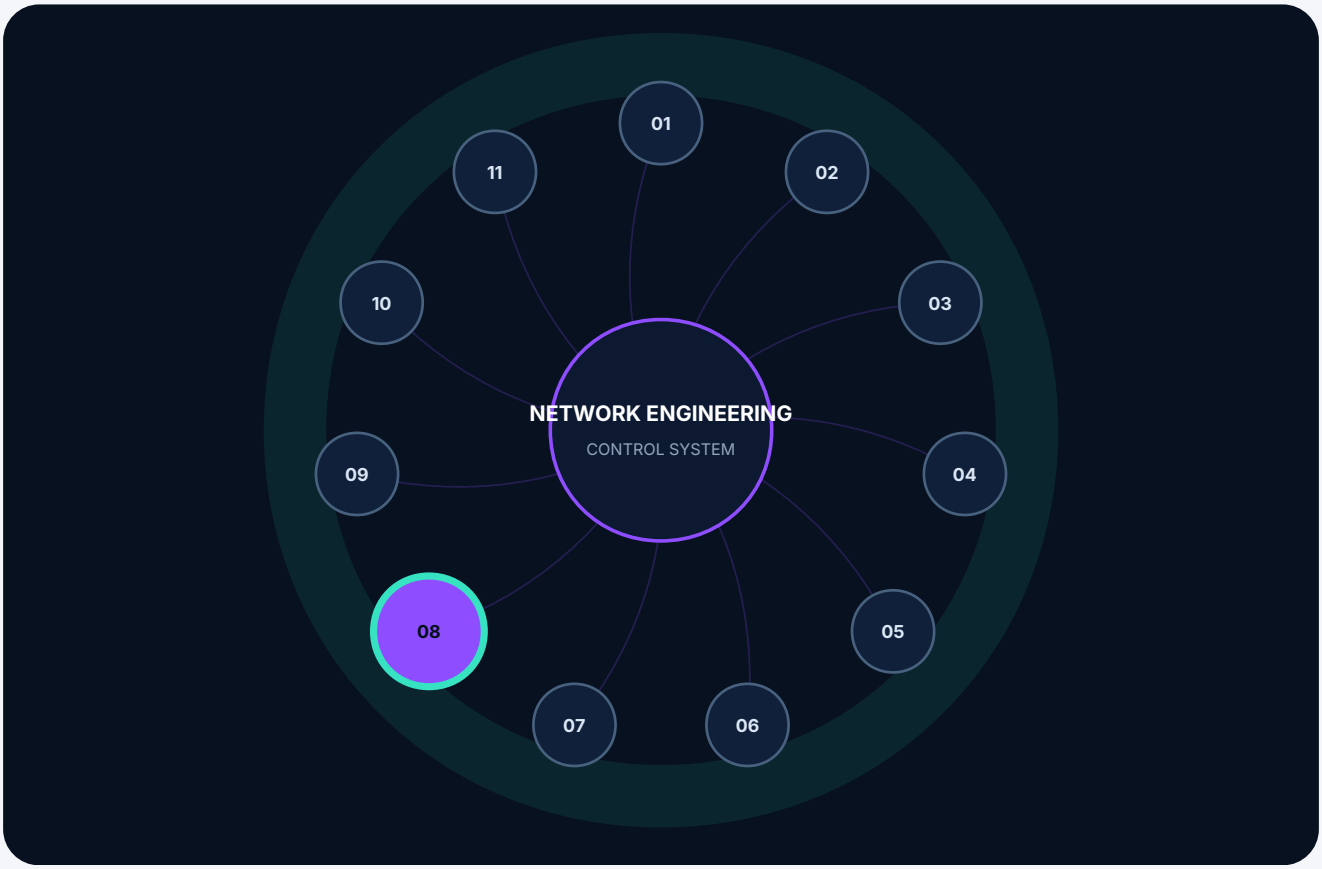

Connected assurance

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

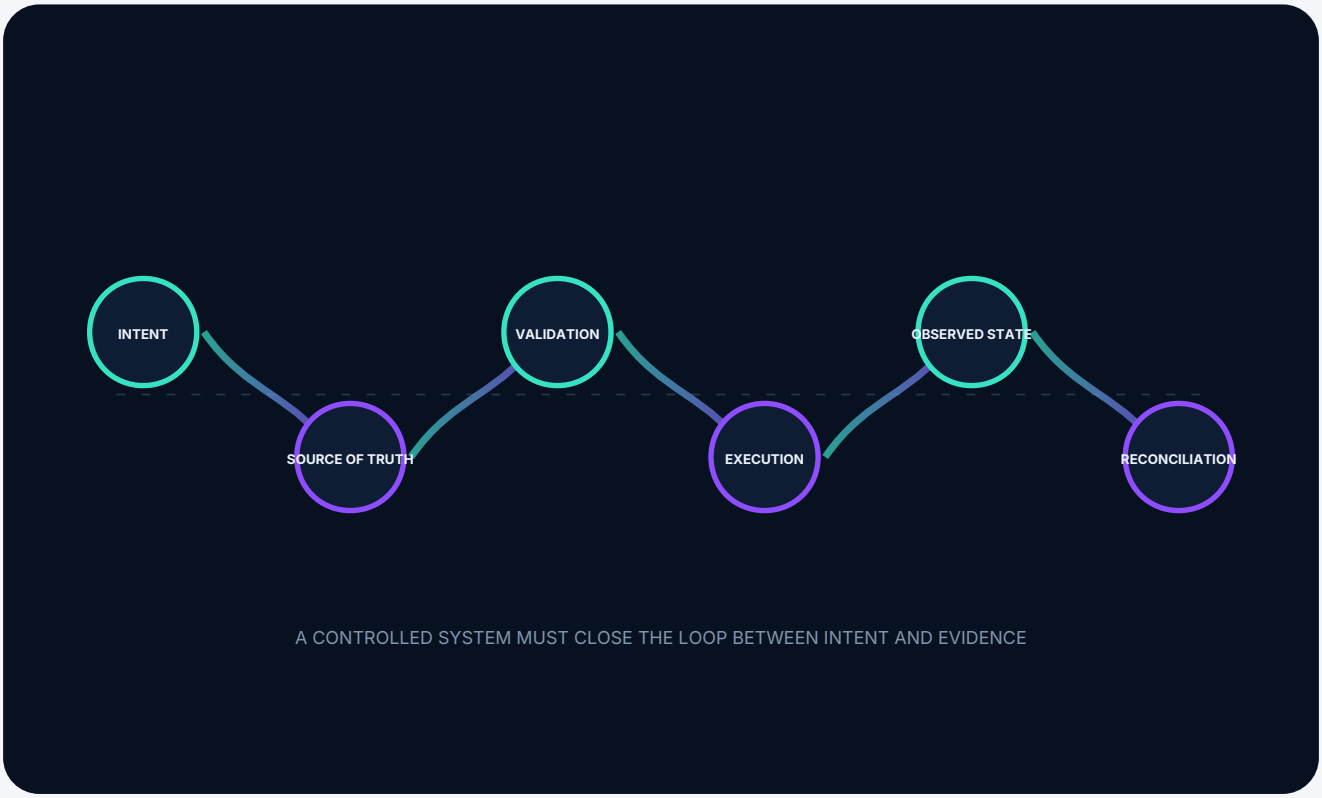

State separation

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-NET-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

The practice of network troubleshooting has failed.

When an application breaks, engineers immediately blame the network. Network engineers, in turn, reply with "I pinged it, the network is fine," without understanding how lower-layer MTU mismatches break upper-layer TCP sessions, or how asymmetric routing drops stateful firewalls. Troubleshooting is treated as tribal knowledge passed down via ancient runbooks, rather than a deterministic, scientific methodology.

This standard exists because:

1. **"Ping" is Not a Diagnostic Tool:** An ICMP echo reply tells you only that the Layer 3 control plane is alive between two endpoints. It tells you nothing about Layer 4 TCP state, Layer 2 MTU consistency, or Layer 7 application health. Engineers MUST utilize structured OSI model traversal to isolate faults.
2. **Core Services are Not Afterthoughts:** DNS, DHCP, and NTP are the foundational pillars of modern infrastructure. A 10-millisecond NTP drift breaks Kerberos and mTLS. A DNS timeout cascades into catastrophic application latency. Core services MUST be treated as zero-trust, highly available, critical infrastructure.
3. **Tribal Knowledge is Unscalable:** If a network can only be fixed by the engineer who built it, the architecture is a liability. Troubleshooting flows MUST be documented as deterministic algorithms (e.g., "If A cannot reach B on port 443, execute Step 1 through Step 5").
4. **Layer Isolation is Mandatory:** An MTU black hole (Layer 2) manifests as a hung TLS handshake (Layer 6). Engineers MUST be able to decouple the symptoms (Layer 7 timeout) from the root cause (Layer 2 fragmentation) using packet captures and eBPF tracing.

This document establishes the **Mythos OSI Troubleshooting Standard (MOTS)**, a comprehensive, evidence-based methodology for evaluating network operations and troubleshooting workflows against structural rigor, packet-level visibility, and core infrastructure resilience.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Structured troubleshooting methodologies (Bottom-Up, Top-Down, Divide-and-Conquer)
- OSI Layer 1-7 fault isolation and scenario flows
- Core infrastructure: DNS (DoH/DoT), DHCP/DDI, IPAM, NTP/PTP
- Packet flow analysis and stateful firewall troubleshooting
- MTU/MSS black hole detection and Path MTU Discovery (PMTUD)
- Network observability tools (tcpdump, Wireshark, eBPF, streaming telemetry)

1.2 Out of Scope

- Specific routing protocol configurations (covered in MYTHOS-NET-0010)
- Physical cabling standards and rack layouts (covered in MYTHOS-NET-0009)
- Zero-trust policy architecture (covered in MYTHOS-SEC-0007)

1.3 Audience

- Network engineers and NOC architects
- Cloud and platform engineers diagnosing cross-zone connectivity
- Application engineers needing to understand network failure modes
- Security analysts investigating lateral movement and packet drops
- Standards bodies seeking a reference troubleshooting methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Troubleshooting Dimensions

Dimension	Definition
Bottom-Up Methodology	A troubleshooting approach starting at Layer 1 (Physical/Cabling) and moving up to Layer 7 (Application), used when physical or environmental issues are suspected.
Top-Down Methodology	An approach starting at Layer 7 (Application/API) and moving down, used when the network is believed to be healthy but the application is failing.
Divide-and-Conquer	An approach starting at Layer 3/4 (Network/Transport) using tools like `ping` and `telnet` to quickly determine if the issue is above or below the network layer.
MTU Black Hole	A failure where a packet larger than the Maximum Transmission Unit of a link is silently dropped because the ICMP "Fragmentation Needed" message is blocked, causing TCP sessions to hang.
Asymmetric Routing	When the forward path (A to B) takes a different physical or logical route than the return path (B to A), often causing stateful firewalls to drop the traffic.

2.2 The Troubleshooting Doctrine

> A ping is not a diagnosis; it is a heartbeat. A timeout is a symptom, not a root cause. A network without time synchronization is a network without logs. If the ICMP "Fragmentation Needed" packet is blocked, the TCP session is dead.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; tribal knowledge, "ping" only, no structured methodology
1	Basic runbooks exist, but lack OSI layer isolation
2	Functional troubleshooting, but lacks eBPF/packet-level visibility
3	Solid production performance; structured OSI flows, DDI high-availability, eBPF tracing
4	Strong performance; automated MTU detection, PTP synchronization, zero-trust core services
5	Best-in-class; sets the standard for mathematically verified, deterministic fault isolation

Weighting

Category	Weight	Rationale
Methodology & Layer Isolation	20%	Guessing wastes uptime; structured OSI traversal is mandatory
Core Infrastructure (DNS/DHCP/NTP)	20%	The network is irrelevant if name resolution or time drifts
Packet Flow & State Analysis	15%	Firewalls drop based on state, not just port; asymmetry breaks apps
MTU & MSS Optimization	15%	Black holes are the most common cause of "random" app hangs
Observability & eBPF Tracing	15%	CLI commands are ephemeral; kernel-level tracing is required
Scenario Execution & Runbooks	15%	Troubleshooting must be deterministic and reproducible

Total: 100%

A system MUST score at least 3.0 in Methodology and Core Infrastructure to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

6 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Methodology and Layer Isolation

SOURCE WEIGHT 20%

4.1.1
Structural Approach

MYTHOS-NET-0008-EVAL-4.1.1

Structural Approach

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Methodology and Layer Isolation	**4.1.1**	**20%**	**Evidence-led**

Do engineers follow a deterministic OSI model traversal to isolate faults?

0
Ad-hoc guessing; "reboot the switch" mentality

1
Top-down only (blame the app, then blame the network)

2
Basic divide-and-conquer (ping/telnet), but no Layer 1/2 checks

3
Strict bottom-up or top-down methodology enforced via runbooks; Layer 1 checked before Layer 7

4
Automated layer isolation: scripts instantly test L1 (link state), L2 (ARP), L3 (Route), L4 (Port) sequentially

5
Leading methodology: formally verified fault isolation trees, zero guessing, mathematical proof of layer isolation

ENGINEERING INTERPRETATION

This criterion evaluates whether structural approach 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

Core Infrastructure (DNS, DHCP, NTP)

SOURCE WEIGHT 20%

4.2.1
Foundational Resilience

MYTHOS-NET-0008-EVAL-4.2.1

Foundational Resilience

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Core Infrastructure (DNS, DHCP, NTP)	4.2.1	20%	Evidence-led

Are name resolution, IP allocation, and time synchronization treated as critical infrastructure?

0
Single DNS server; NTP drift > 5 seconds; manual IP allocation

1
Basic redundant DNS/DHCP, but no monitoring; NTP via internet

2
DDI (DNS, DHCP, IPAM) suite used; internal NTP servers

3
High-availability DDI cluster; PTP (Precision Time Protocol) for microsecond sync; DoH/DoT for internal DNS

4
Core services placed in a zero-trust, isolated management plane; failover automated and tested

5
Leading resilience: formally verified uptime bounds, cryptographic time sync, zero DNS hijack potential

ENGINEERING INTERPRETATION

This criterion evaluates whether foundational resilience 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

Packet Flow and State Analysis

SOURCE WEIGHT 15%

4.3.1
Stateful Inspection Understanding

MYTHOS-NET-0008-EVAL-4.3.1

Stateful Inspection Understanding

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Packet Flow and State Analysis	**4.3.1**	**15%**	**Evidence-led**

Can engineers trace a packet through stateful firewalls and identify asymmetric routing?

0
No understanding of stateful inspection; "port is open, why doesn't it work?"

1
Basic understanding, but no tools to trace path

2
`traceroute` and `tcpdump` used, but cannot identify state drops

3
Engineers can perform bidirectional packet captures to identify asymmetric routing and state mismatches

4
Automated path analysis tools (eBPF/ digital twin) predict state drops before deployment

5
Leading analysis: formally verified state tables, zero asymmetric routing blind spots, real-time state visualization

ENGINEERING INTERPRETATION

This criterion evaluates whether stateful inspection understanding 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

- routing policy, RIB/FIB snapshots, adjacency state, and path-analysis output
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. simulate route withdrawal, policy conflict, convergence, and rollback
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- route leaks, asymmetric paths, stale advertisements, and unbounded convergence
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- route-policy compliance
- convergence time
- unexpected path count
- 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

MTU and MSS Optimization

SOURCE WEIGHT 15%

4.4.1
Black Hole Prevention

MYTHOS-NET-0008-EVAL-4.4.1

Black Hole Prevention

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
MTU and MSS Optimization	4.4.1	15%	Evidence-led

Is Path MTU Discovery (PMTUD) functional, or are ICMP blocks causing black holes?

0
ICMP completely blocked; TCP sessions hang randomly on large payloads

1
PMTUD broken, but MSS clamping manually configured on some routers

2
ICMP "Fragmentation Needed" allowed; PMTUD functions

3
Dynamic MSS clamping enforced at edge; MTU mismatches alerted via telemetry

4
Automated MTU path testing continuously validates all critical links

5
Leading optimization: formally verified payload sizes, zero black holes, mathematical proof of PMTUD success

ENGINEERING INTERPRETATION

This criterion evaluates whether black hole 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.

4.5 / CAPABILITY FAMILY

Observability and eBPF Tracing

SOURCE WEIGHT 15%

4.5.1
Kernel-Level Visibility

MYTHOS-NET-0008-EVAL-4.5.1

Kernel-Level Visibility

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Observability and eBPF Tracing	**4.5.1**	**15%**	**Evidence-led**

Can engineers trace packets from the NIC, through the kernel, to the application socket?

0
Only interface counters (ifErrors) available

1
`tcpdump` on interfaces, but no application-level visibility

2
Network and application logs correlated manually

3
eBPF used to trace packet lifecycle (NIC → TCP stack → Socket → App)

4
Streaming telemetry (sFlow/IPFIX) provides real-time flow visibility

5
Leading visibility: formally verified packet tracing, zero blind spots between wire and app

ENGINEERING INTERPRETATION

This criterion evaluates whether kernel-level visibility 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.

4.6 / CAPABILITY FAMILY

Scenario Execution and Runbooks

SOURCE WEIGHT 15%

4.6.1
Reproducibility

MYTHOS-NET-0008-EVAL-4.6.1

Reproducibility

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Scenario Execution and Runbooks	**4.6.1**	**15%**	**Evidence-led**

Is troubleshooting deterministic, or dependent on the specific engineer on call?

0
"Tribal knowledge" only; no documentation

1
High-level runbooks ("Check if DNS is working")

2
Detailed runbooks with specific CLI commands

3
Runbooks structured as algorithms (If X, then Y; else Z) mapped to OSI layers

4
Runbooks automated via infrastructure-as-code and SOAR platforms

5
Leading reproducibility: formally verified troubleshooting algorithms, zero human variance in fault isolation

ENGINEERING INTERPRETATION

This criterion evaluates whether reproducibility 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 Troubleshooting Suite (MOTS)

Traditional network benchmarks measure convergence time. The MOTS evaluates fault isolation speed, layer understanding, and core infrastructure resilience.

5.1 Suite Composition

5.1.1 MOTS-Isolation

- **MTU Black Hole Injection:** 100+ tests dropping ICMP "Fragmentation Needed" packets, verifying the NOC identifies the L2 MTU issue rather than blaming the L7 application.
- **Asymmetric Route Injection:** 50+ tests injecting a routing change that causes return traffic to bypass the stateful firewall, verifying engineers identify the state drop.

5.1.2 MOTS-Core

- **DNS Poisoning Simulation:** 100+ tests injecting a rogue DNS response, verifying DoH/DoT or DNSSEC prevents the hijack.
- **NTP Drift Test:** 50+ tests inducing clock drift on a node, verifying PTP corrects it before mTLS/Kerberos breaks.

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.

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

Physical Layer, Cabling, Optics, and Data-Center Topology Standard

Defines cabling, optics, facilities interfaces, topology, labeling, testing, and lifecycle.

PERMANENT DOCUMENT ID
MYTHOS-NET-0009

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
6 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

Physical Layer, Cabling, Optics, and Data-Center Topology Standard

Defines cabling, optics, facilities interfaces, topology, labeling, testing, and lifecycle.

DOCUMENT ID
MYTHOS-NET-0009

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

6	6	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-NET-0009 inside the network engineering standard constellation

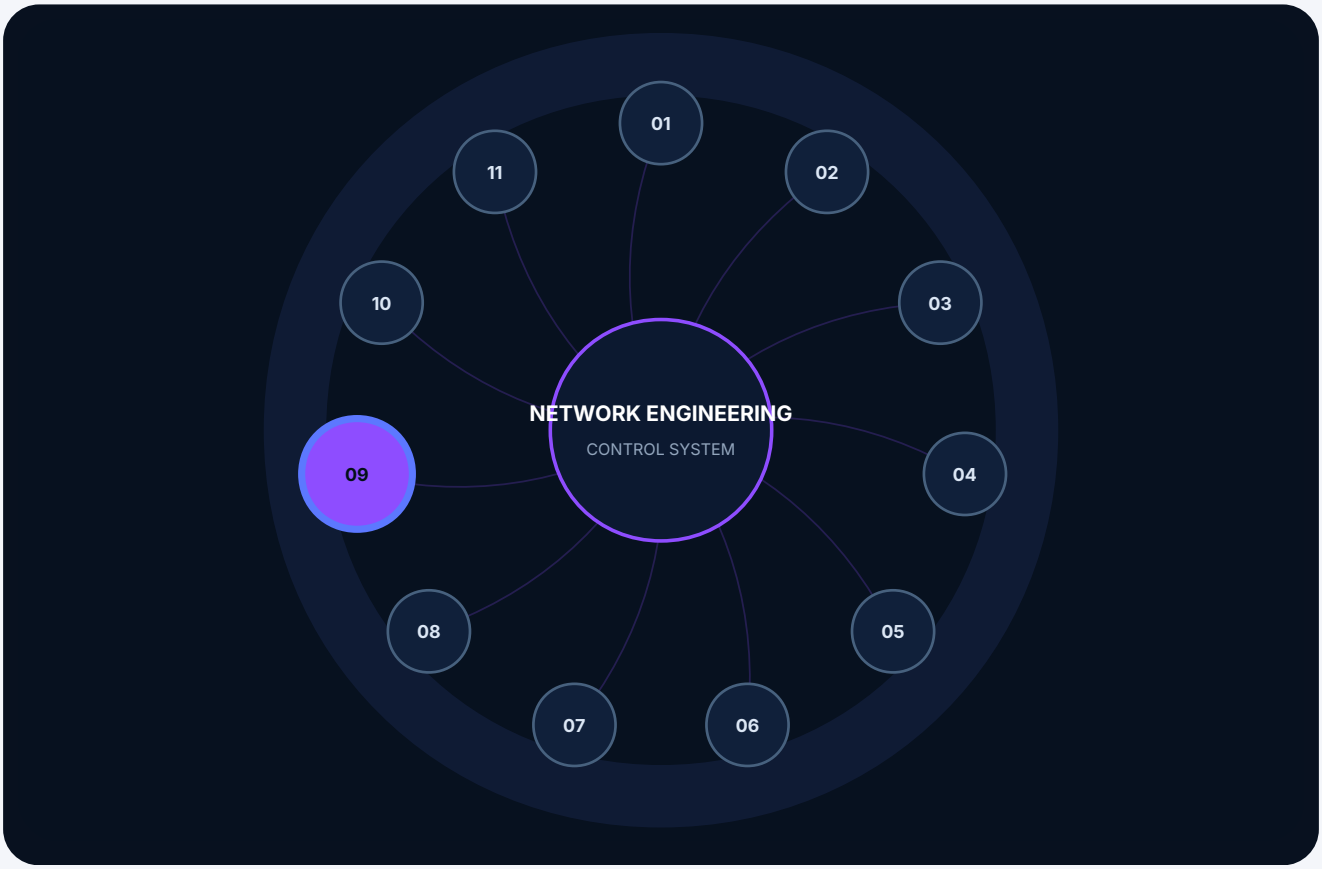

Connected assurance

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

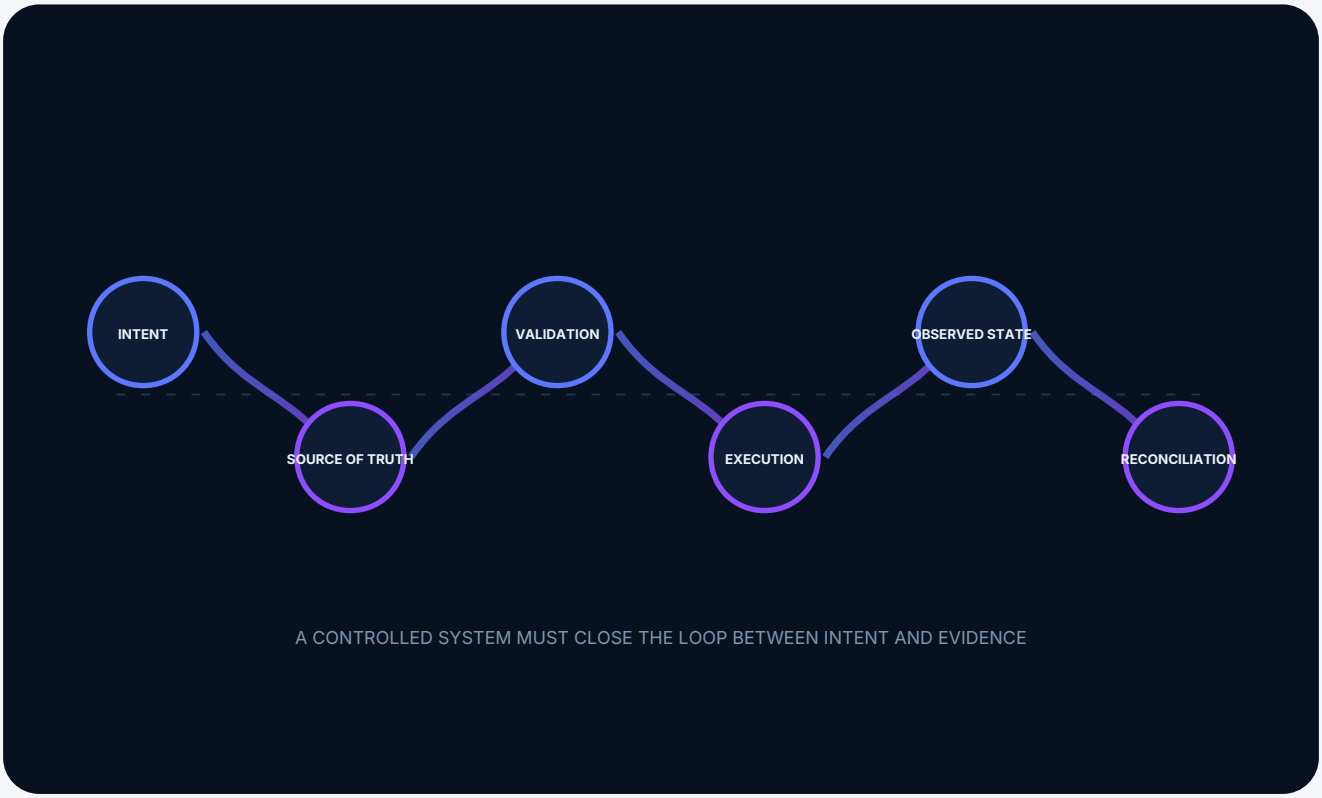

State separation

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-NET-0009

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

The architecture of the physical datacenter has failed.

Organizations spend millions on high-availability cloud architectures and distributed state machines, only to house them in datacenters where a single misrouted cable takes down a core switch, a blocked CRAC unit causes thermal throttling, and unlabeled fiber strands result in hours of downtime. Physical infrastructure is treated as mere "plumbing" rather than the foundational Layer 1 upon which all zero-trust and AI compute paradigms rely.

This standard exists because:

1. **Spaghetti Cabling is a Security and Operational Threat:** An unorganized, unlabeled cable plant guarantees extended outages. When a patch cable fails, engineers must trace it through a rat's nest of wires, increasing Mean Time to Recovery (MTTR) from seconds to hours. Structured cabling with strict color-coding and mandatory labeling is a prerequisite for production.
2. **Thermals Dictate Compute Density:** Modern AI GPU clusters (e.g., NVIDIA Hopper/Blackwell) draw 10kW+ per rack. Traditional datacenter cooling (perforated tiles, hot/cold aisle) is insufficient. Organizations MUST implement containment strategies, in-row cooling, or direct-to-chip liquid cooling to prevent thermal throttling and hardware degradation.
3. **Single-Point Power is high-risk architectural constraint:** A server plugged into a single PDU, or a rack fed by a single power feed, is a ticking time bomb. Production systems MUST have A/B power diversity, from the UPS down to the server's redundant power supplies, capable of surviving a single-phase or single-feed failure.
4. **The Physical Layer Must Be Software-Defined:** Treating the datacenter as a static physical entity is obsolete. Organizations MUST utilize DCIM (Datacenter Infrastructure Management) tools to monitor power draw, temperature, and humidity in real-time, treating power and cooling as consumable, programmable resources.

This document establishes the **Mythos Physical Infrastructure Standard (MPIS)**, a comprehensive, evidence-based methodology for evaluating datacenter topologies, cabling rigor, and environmental resilience.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Structured cabling (Copper Cat6a/Cat8, Fiber OM4/OM5/OS2)
- Cable color-coding, labeling (TIA-606-C), and management (hook-and-loop vs. zip ties)
- Datacenter topology (Hot/Cold aisle containment, rack layouts)
- Power distribution (A/B feeds, PDUs, UPS, generators)
- Cooling and thermal management (In-row, liquid cooling, blanking panels)
- Environmental monitoring (DCIM, leak detection, temperature/humidity sensors)

1.2 Out of Scope

- Logical network topology and routing protocols (covered in MYTHOS-NET-0001)
- Hardware specifications (covered in MYTHOS-HW-0001)
- General disaster recovery (covered in MYTHOS-SRE-0002)

1.3 Audience

- Datacenter facility managers and mechanical engineers
- Network engineers and cable plant technicians
- IT operations and infrastructure architects
- Hardware procurement teams
- Standards bodies seeking a reference physical infrastructure methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Physical Dimensions

Dimension	Definition
Structured Cabling	A standardized architecture for telecommunications cabling, using dedicated pathways, patch panels, and strict labeling to ensure modularity and ease of management.
A/B Power Diversity	An architectural paradigm where critical infrastructure receives power from two completely independent electrical feeds (UPS A and UPS B), terminating on dual power supplies in the server.
Hot/Cold Aisle Containment	A datacenter cooling strategy where server racks are arranged in alternating rows (hot and cold), and physical barriers are used to prevent hot exhaust air from mixing with cold intake air.
DCIM	Datacenter Infrastructure Management. Software tools that monitor, measure, and manage datacenter resources (power, cooling, space, connectivity) in real-time.
OM5 Fiber	Wideband multimode fiber optimized for Short Wavelength Division Multiplexing (SWDM), allowing multiple wavelengths over a single fiber pair for high-density datacenter backbones.

2.2 The Physical Infrastructure Doctrine

> An unlabeled cable is a liability. A single power feed is a single point of failure. Mixed airflow is a thermal bottleneck. If the physical layer is not structured, the logical layer is an illusion.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; spaghetti cabling, single power feed, no containment, thermal hotspots
1	Basic labeling exists, but inconsistent power and cooling
2	Functional A/B power, but lacks strict color-coding or DCIM monitoring
3	Solid production performance; strict TIA-606-C labeling, A/B power, hot/cold containment, DCIM
4	Strong performance; liquid cooling for high-density, automated environmental alerts
5	Best-in-class; sets the standard for mathematically verified, autonomous physical resilience

Weighting

Category	Weight	Rationale
Cabling Structure & Labeling	20%	Untraceable cables cause massive MTTR during outages
Power Redundancy & Diversity	20%	A single power event should never take down a production node
Cooling & Airflow Containment	20%	High-density AI compute requires strict thermal management
Environmental Monitoring (DCIM)	15%	Blindness to power/thermal data leads to cascading failures
Physical Security & Access	15%	Unlocked racks allow physical tampering and cable tapping
Rack Density & Space Utilization	10%	Poor space planning halts scaling capabilities

Total: 100%

A system MUST score at least 3.0 in Cabling Structure and Power Redundancy to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

6 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Cabling Structure and Labeling

SOURCE WEIGHT 20%

4.1.1
Structured Plant Rigor

MYTHOS-NET-0009-EVAL-4.1.1

Structured Plant Rigor

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Cabling Structure and Labeling	**4.1.1**	**20%**	**Evidence-led**

Is the cabling plant strictly organized, color-coded, and traceable?

0
Tangled "spaghetti" cabling; zip ties used; no labels

1
Basic patch panels used, but inconsistent labeling and color coding

2
TIA-606-C compliant labels at both ends; hook-and-loop straps used

3
Strict color-coding enforced (e.g., Red = Security, Blue = Copper, Orange = Internal Fiber, Yellow = Out-of-Band)

4
Cable paths documented in DCIM; fiber optic OTDR (Optical Time-Domain Reflectometer) traces mapped

5
Leading structure: formally verified cable topologies, zero undocumented strands, automated patch-port validation

ENGINEERING INTERPRETATION

This criterion evaluates whether structured plant rigor 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

Power Redundancy and Diversity

SOURCE WEIGHT 20%

4.2.1
Electrical Resilience

MYTHOS-NET-0009-EVAL-4.2.1

Electrical Resilience

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Power Redundancy and Diversity	**4.2.1**	**20%**	**Evidence-led**

Can a server survive the loss of a single power feed or UPS?

0
Single PDU per rack; single power supply in servers

1
Dual PDUs, but fed from the same UPS

2
A/B power feeds from diverse UPS systems; servers have dual PSUs

3
A/B feeds from diverse generator/ transformer sources; automated transfer switches (ATS) tested regularly

4
Rack-level power monitoring (per outlet) with automated capping to prevent breaker trips

5
Leading resilience: formally verified power diversity, zero single points of failure from grid to server, sub-second failover

ENGINEERING INTERPRETATION

This criterion evaluates whether electrical resilience 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

Cooling and Airflow Containment

SOURCE WEIGHT 20%

4.3.1
Thermal Management

MYTHOS-NET-0009-EVAL-4.3.1

Thermal Management

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Cooling and Airflow Containment	4.3.1	20%	Evidence-led

Is the datacenter capable of cooling high-density AI racks without throttling?

0
Mixed airflow (some racks face front, some back); no blanking panels

1
Basic hot/cold aisle alignment, but no containment

2
Physical containment (doors/end-of-row panels) separating hot and cold aisles

3
In-row cooling deployed for high-density AI/GPU racks (>15kW)

4
Direct-to-chip liquid cooling (DLC) for extreme-density clusters (>40kW)

5
Leading cooling: formally verified Computational Fluid Dynamics (CFD), zero thermal hotspots, dynamic fan/pump scaling based on real-time CPU/GPU thermals

ENGINEERING INTERPRETATION

This criterion evaluates whether thermal management 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

Environmental Monitoring (DCIM)

SOURCE WEIGHT 15%

4.4.1
Facility Observability

MYTHOS-NET-0009-EVAL-4.4.1

Facility Observability

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Environmental Monitoring (DCIM)	4.4.1	15%	Evidence-led

Is the physical environment monitored with the same rigor as the application stack?

0
No environmental monitoring; relies on server internal sensors

1
Basic temperature sensors in the room

2
Per-rack temperature and humidity sensors

3
Comprehensive DCIM platform tracking power, cooling, and space; leak detection under raised floors

4
Integration of DCIM with ITSM (e.g., auto-generating tickets when a UPS battery fails)

5
Leading observability: formally verified sensor grids, AI-driven predictive cooling, zero blind spots in facility telemetry

ENGINEERING INTERPRETATION

This criterion evaluates whether facility observability 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

Physical Security and Access

SOURCE WEIGHT 15%

4.5.1
Access Control

MYTHOS-NET-0009-EVAL-4.5.1

Access Control

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Physical Security and Access	**4.5.1**	**15%**	**Evidence-led**

How is physical access to the datacenter and individual racks governed?

0
Unlocked server room; open racks

1
Single badge reader on the datacenter door

2
Biometric (fingerprint/retina) access on datacenter entry; locked cabinets

3
Mantrap (interlocking doors) at datacenter entry; CCTV covering all aisles

4
Per-rack electronic locks tied to centralized identity (e.g., audited access logs per cabinet)

5
Leading security: formally verified zero-trust physical access, cryptographic attestation of rack entry, zero tailgating potential

ENGINEERING INTERPRETATION

This criterion evaluates whether access 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.

4.6 / CAPABILITY FAMILY

Rack Density and Space Utilization

SOURCE WEIGHT 10%

4.6.1
Scalability

MYTHOS-NET-0009-EVAL-4.6.1

Scalability

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Rack Density and Space Utilization	**4.6.1**	**10%**	**Evidence-led**

Is the physical space utilized efficiently to allow for future scaling?

0
Ad-hoc rack layouts; wasted space; no weight limits considered

1
Basic rack standards, but no power/ cooling capacity planning

2
Standardized rack elevations documented in DCIM

3
Power and cooling capacity mapped to physical space; new deployments validated against constraints

4
High-density zones established for GPU clusters; standard zones for storage/ network

5
Leading utilization: formally verified space/power/cooling algorithms, zero stranded capacity, optimal workload placement

ENGINEERING INTERPRETATION

This criterion evaluates whether scalability 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 Physical Infrastructure Suite (MPIS)

Traditional facility audits measure uptime. The MPIS evaluates cable traceability, power failover, and thermal resilience under load.

5.1 Suite Composition

5.1.1 MPIS-Cable

- **Blind Trace Test:** 100+ tests requiring a technician to trace a specific port to its origin without visual inspection, verifying the DCIM and labeling system allows instant identification.
- **OTDR Validation:** 50+ tests running Optical Time-Domain Reflectometer scans to verify fiber pairs match documented paths and have no micro-bends.

5.1.2 MPIS-PowerThermal

- **Feed Failure Simulation:** 50+ tests simulating the loss of Power Feed A, verifying all servers seamlessly failover to Feed B without dropping packets.
- **Thermal Soak Test:** 50+ tests running GPU clusters at 100% utilization for 4 hours, verifying containment and cooling systems prevent thermal throttling.

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.

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

Routing, DNS, DHCP, and IPAM Standard

Defines routing architecture, naming, addressing, DDI, policy, operations, and assurance.

PERMANENT DOCUMENT ID
MYTHOS-NET-0010

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
6 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

Routing, DNS, DHCP, and IPAM Standard

Defines routing architecture, naming, addressing, DDI, policy, operations, and assurance.

DOCUMENT ID
MYTHOS-NET-0010

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

6	6	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-NET-0010 inside the network engineering standard constellation

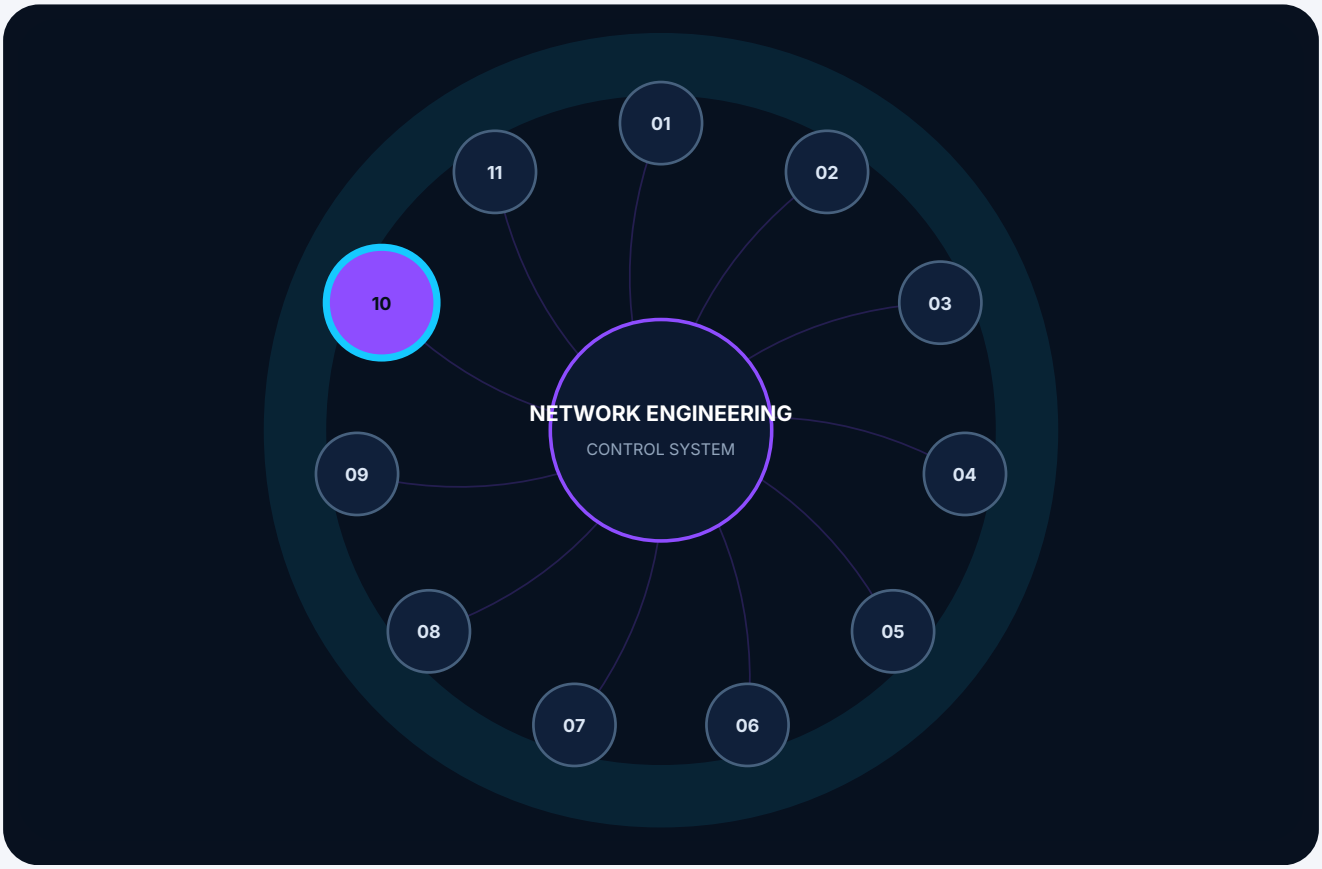

Connected assurance

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

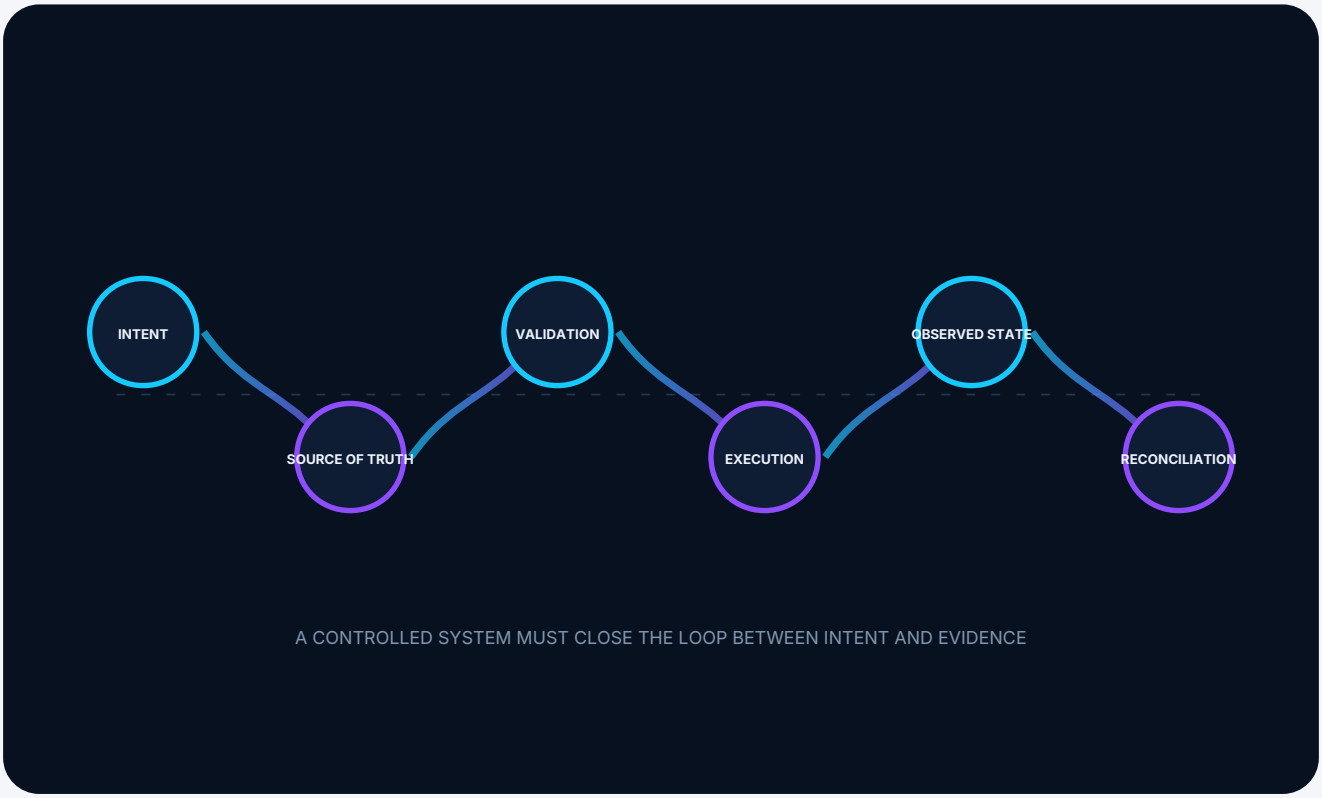

State separation

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-NET-0010

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

The architecture of core network services has failed.

Organizations attempt to manage millions of IP addresses across multi-cloud and on-premises environments using static Excel spreadsheets. They accept BGP routes from transit providers based on "trust" rather than cryptographic validation, leaving the global routing table vulnerable to hijacking. They expose internal DNS queries in plaintext, allowing attackers to map network topologies and exfiltrate data via DNS tunneling.

This standard exists because:

1. **Spreadsheets are Not IPAM:** Tracking IP allocations in a spreadsheet guarantees IP conflicts, orphaned addresses, and broken automation. IPAM MUST be a dynamic, API-first system integrated with DNS and DHCP (DDI), serving as the absolute Source of Truth for network intent.
2. **BGP is an Unauthenticated Anarchy:** The default global routing infrastructure assumes all participants are honest. Without RPKI (Resource Public Key Infrastructure) and Route Origin Validation (ROV), any attacker can hijack an IP prefix and reroute global traffic. Production routing MUST cryptographically validate routes.
3. **DNS is an Attack Surface, Not a Phonebook:** Plaintext DNS allows ISPs and attackers to inspect, block, or poison name resolutions. Furthermore, internal DNS queries reveal topological secrets. Systems MUST enforce DNSSEC for data integrity and DNS over HTTPS/TLS (DoH/DoT) for transport privacy.
4. **Manual DHCP is a Bottleneck:** In modern, ephemeral infrastructure, IP allocation must happen at machine speed. DDI MUST be fully automated, integrating with cloud providers and Kubernetes to allocate, reserve, and reclaim addresses without human intervention.

This document establishes the **Mythos Routing, DNS & IPAM Standard (MRDIS)**, a comprehensive, evidence-based methodology for evaluating core network services against cryptographic routing validation, zero-trust name resolution, and automated address lifecycle management.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- DDI (DNS, DHCP, IPAM) platforms (BlueCat, Infoblox, NetBox, Nautobot)
- Routing protocols and architectures (BGP, OSPF, IS-IS, MPLS, VXLAN/EVPN)
- Route security (RPKI, Route Origin Validation, BGPsec, MANRS)
- DNS architecture (Split-horizon, DNSSEC, DoH, DoT, DNS tunneling prevention)
- DHCP automation and IP lifecycle management (allocation, reclamation, conflict detection)

1.2 Out of Scope

- Physical layer cabling and topology (covered in MYTHOS-NET-0009)
- High-level network design and intent (covered in MYTHOS-NET-0001)
- Zero-trust network access (ZTNA) and microsegmentation policy (covered in MYTHOS-SEC-0007)

1.3 Audience

- Network engineers and architects designing core infrastructure
- Cloud and platform engineers managing IP allocation at scale
- Security engineers evaluating DNS attack surfaces and route hijacking
- DevSecOps teams automating network configuration
- Standards bodies seeking a reference DDI and routing methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Core Service Dimensions

Dimension	Definition
DDI	The convergence of DNS, DHCP, and IPAM into a single, unified management platform.
RPKI (Resource Public Key Infrastructure)	A cryptographic framework that binds IP address prefixes to their rightful Autonomous System (ASN), allowing routers to validate BGP routes.
Route Origin Validation (ROV)	The process of a BGP router dropping routes that do not match RPKI records, preventing prefix hijacking.
Split-Horizon DNS	An architecture where internal and external clients receive different IP address responses for the same domain name, hiding internal topology.
DNSSEC	A suite of cryptographic extensions to DNS that authenticates the origin of DNS data, verifying it has not been modified in transit.
IPAM Source of Truth	The principle that the IPAM system is the sole authority for IP allocation; no IP is assigned via CLI or cloud console without IPAM API approval.

2.2 The Core Services Doctrine

> A spreadsheet is not a source of truth. An unvalidated BGP route is a hijack. A plaintext DNS query is a public broadcast. If the IP lifecycle is not automated, the network cannot scale.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; spreadsheets for IPAM, unvalidated BGP, plaintext DNS
1	Basic DDI exists, but manual processes and no RPKI enforcement
2	Functional routing and DNS, but lacks cryptographic validation or automation
3	Solid production performance; API-driven DDI, RPKI/ROV enforced, DNSSEC, split-horizon DNS
4	Strong performance; BGPsec, DoH/DoT internally, automated IP reclamation
5	Best-in-class; sets the standard for mathematically verified, zero-trust core services

Weighting

Category	Weight	Rationale
IPAM & DDI Automation	25%	Manual IP tracking breaks at cloud scale and causes conflicts
BGP & Routing Security (RPKI)	20%	Unvalidated routes allow global traffic hijacking
DNS Security & Privacy	20%	Plaintext DNS is easily spoofed and mined for intelligence
DHCP & Address Lifecycle	15%	Ephemeral compute requires sub-second, automated allocation
High Availability & Resilience	10%	Core services are single points of failure; must be highly available
Observability & Drift Detection	10%	IPAM must reconcile with actual network state

Total: 100%

A system MUST score at least 3.0 in IPAM Automation and BGP/Routing Security to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

6 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

IPAM and DDI Automation

SOURCE WEIGHT 25%

4.1.1
Address Management Authority

MYTHOS-NET-0010-EVAL-4.1.1

Address Management Authority

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
IPAM and DDI Automation	4.1.1	25%	Evidence-led

Is the IPAM system the absolute, automated Source of Truth for all IP addresses?

0
IP addresses tracked in spreadsheets; manual CLI configuration on devices

1
Basic IPAM tool used, but requires manual sync with DNS/DHCP

2
Integrated DDI platform, but lacks API integration with cloud/Kubernetes

3
API-first DDI: all IP allocations (on-prem, cloud, k8s) automated via API; zero manual CLI allocation

4
Automated IP reclamation: DDI detects orphaned IPs and reclaims them dynamically

5
Leading authority: formally verified IP lifecycle, zero IP conflicts possible, cryptographic attestation of allocations

ENGINEERING INTERPRETATION

This criterion evaluates whether address management authority 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

- source-of-truth objects, Git history, observed-state snapshots, and reconciliation evidence
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. introduce controlled drift and verify detection, adjudication, and restoration
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- multiple authorities, silent manual changes, stale inventory, and destructive auto-remediation
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- drift detection time
- reconciliation success
- unowned object count
- 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.

4.2 / CAPABILITY FAMILY

BGP and Routing Security (RPKI)

SOURCE WEIGHT 20%

4.2.1
Route Validation

MYTHOS-NET-0010-EVAL-4.2.1

Route Validation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
BGP and Routing Security (RPKI)	**4.2.1**	**20%**	**Evidence-led**

Is the global routing table cryptographically protected against hijacking?

0
BGP operates on pure trust; no RPKI; accepts any route from peers

1
RPKI signed on outbound, but no inbound Route Origin Validation (ROV)

2
ROV enabled, but in "monitoring" mode (does not drop invalid routes)

3
ROV enforced in "dropping" mode; rejects unauthorized BGP routes

4
BGPsec implemented cryptographically validating the entire AS_PATH

5
Leading security: formally verified routing boundaries, zero ability to inject unauthorized prefixes, full MANRS compliance

ENGINEERING INTERPRETATION

This criterion evaluates whether route validation 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

- routing policy, RIB/FIB snapshots, adjacency state, and path-analysis output
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. simulate route withdrawal, policy conflict, convergence, and rollback
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- route leaks, asymmetric paths, stale advertisements, and unbounded convergence
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- route-policy compliance
- convergence time
- unexpected path count
- 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

DNS Security and Privacy

SOURCE WEIGHT 20%

4.3.1
Name Resolution Integrity

MYTHOS-NET-0010-EVAL-4.3.1

Name Resolution Integrity

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
DNS Security and Privacy	4.3.1	20%	Evidence-led

Are DNS queries protected against spoofing, interception, and surveillance?

0
Open recursive resolvers; plaintext DNS; no DNSSEC

1
Internal DNS servers, but queries sent in plaintext

2
DNSSEC implemented on authoritative zones, but not validated by resolvers

3
DNSSEC validated end-to-end; split-horizon DNS hiding internal topology

4
DoH/DoT (DNS over HTTPS/TLS) enforced for all internal and external queries

5
Leading integrity: formally verified cryptographic resolution, zero plaintext DNS, AI-driven DNS tunneling detection

ENGINEERING INTERPRETATION

This criterion evaluates whether name resolution integrity 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

DHCP and Address Lifecycle

SOURCE WEIGHT 15%

4.4.1
Dynamic Allocation

MYTHOS-NET-0010-EVAL-4.4.1

Dynamic Allocation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
DHCP and Address Lifecycle	**4.4.1**	**15%**	**Evidence-led**

How are addresses allocated to ephemeral and physical devices?

0
Static IP assignments; no DHCP

1
Basic DHCP server, but with finite scopes and manual management

2
DHCP integrated with IPAM, but lacks rapid reclamation

3
Dynamic DHCP with automated lease tracking and conflict detection (ARP ping)

4
Ephemeral compute integration: Kubernetes and cloud instances allocated via DHCP relay to centralized DDI

5
Leading lifecycle: formally verified lease boundaries, zero IP exhaustion risk, sub-millisecond allocation at scale

ENGINEERING INTERPRETATION

This criterion evaluates whether dynamic allocation 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

High Availability and Resilience

SOURCE WEIGHT 10%

4.5.1
Core Service Uptime

MYTHOS-NET-0010-EVAL-4.5.1

Core Service Uptime

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
High Availability and Resilience	4.5.1	10%	Evidence-led

Can the DDI and routing infrastructure survive node failures?

0
Single DNS server; single router; no failover

1
Primary/backup DNS; basic router redundancy (VRRP)

2
Anycast DNS deployed; routing control planes redundant (NSF/GR)

3
Active-active DDI clusters across multiple datacenters; seamless failover

4
Multi-region routing autonomy; DDI survives region loss without caching delays

5
Leading resilience: formally verified zero-downtime failover, mathematical proof of continuous name resolution

ENGINEERING INTERPRETATION

This criterion evaluates whether core service uptime 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

- routing policy, RIB/FIB snapshots, adjacency state, and path-analysis output
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. simulate route withdrawal, policy conflict, convergence, and rollback
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- route leaks, asymmetric paths, stale advertisements, and unbounded convergence
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- route-policy compliance
- convergence time
- unexpected path count
- 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

Observability and Drift Detection

SOURCE WEIGHT 10%

4.6.1
State Reconciliation

MYTHOS-NET-0010-EVAL-4.6.1

State Reconciliation

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Observability and Drift Detection	**4.6.1**	**10%**	**Evidence-led**

Does the actual network state continuously match the IPAM Source of Truth?

0
No monitoring; IPAM drifts from actual device configurations

1
Periodic manual audits of IP allocations

2
Automated scanning (Nmap) compared against IPAM, but requires manual remediation

3
Continuous reconciliation: DDI pulls live state from routers/cloud APIs and alerts on drift

4
Automated drift remediation: unauthorized IP allocations are immediately quarantined or overwritten

5
Leading observability: formally verified state parity, policy-defined drift tolerance, real-time topological visualization

ENGINEERING INTERPRETATION

This criterion evaluates whether state reconciliation 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

- source-of-truth objects, Git history, observed-state snapshots, and reconciliation evidence
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. introduce controlled drift and verify detection, adjudication, and restoration
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- multiple authorities, silent manual changes, stale inventory, and destructive auto-remediation
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- drift detection time
- reconciliation success
- unowned object count
- 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 Routing & DDI Suite (MRDIS)

Traditional benchmarks measure DHCP lease throughput. The MRDIS evaluates route hijacking resistance, DNS spoofing immunity, and IPAM automation.

5.1 Suite Composition

5.1.1 MRDIS-Routing

- **Prefix Hijack Simulation:** 100+ tests announcing unauthorized BGP routes to edge routers, verifying RPKI/ROV immediately drops them.
- **Route Leak Test:** 50+ tests simulating a transit provider leaking full tables, verifying route maps and max-prefix limits drop the leak.

5.1.2 MRDIS-DNS_IPAM

- **Cache Poisoning Test:** 100+ tests injecting spoofed DNS responses, verifying DNSSEC cryptographic validation rejects them.
- **Drift Injection:** 50+ tests manually configuring an IP on a switch interface without IPAM approval, verifying the system detects the drift and alerts/remediates.

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.

MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

Access, NAC, Wireless, and RF Engineering Standard

Defines access switching, 802.1X, profiling, wireless design, RF, roaming, and validation.

PERMANENT DOCUMENT ID
MYTHOS-NET-0011

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
6 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

Access, NAC, Wireless, and RF Engineering Standard

Defines access switching, 802.1X, profiling, wireless design, RF, roaming, and validation.

DOCUMENT ID
MYTHOS-NET-0011

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

6	6	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-NET-0011 inside the network engineering standard constellation

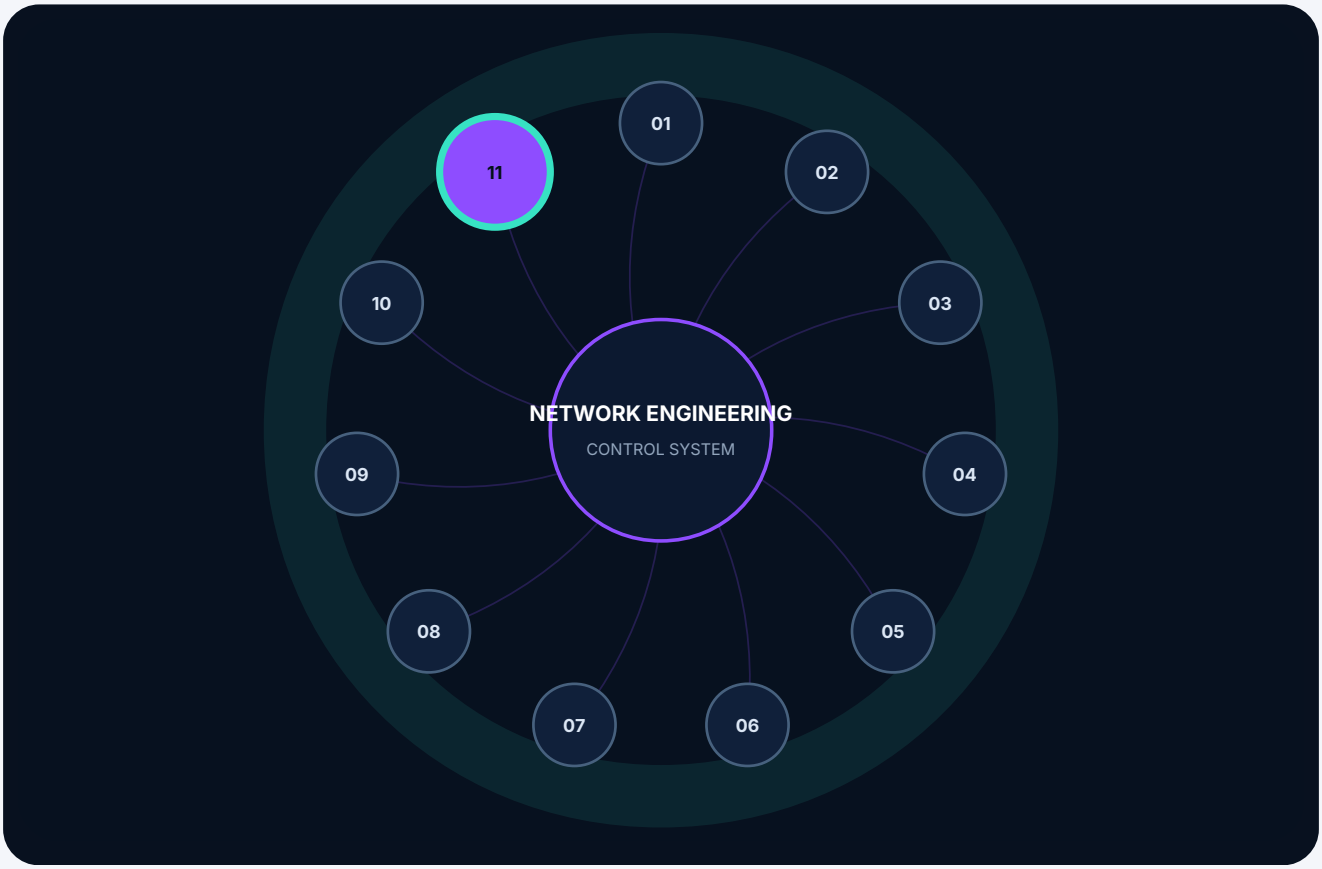

Connected assurance

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

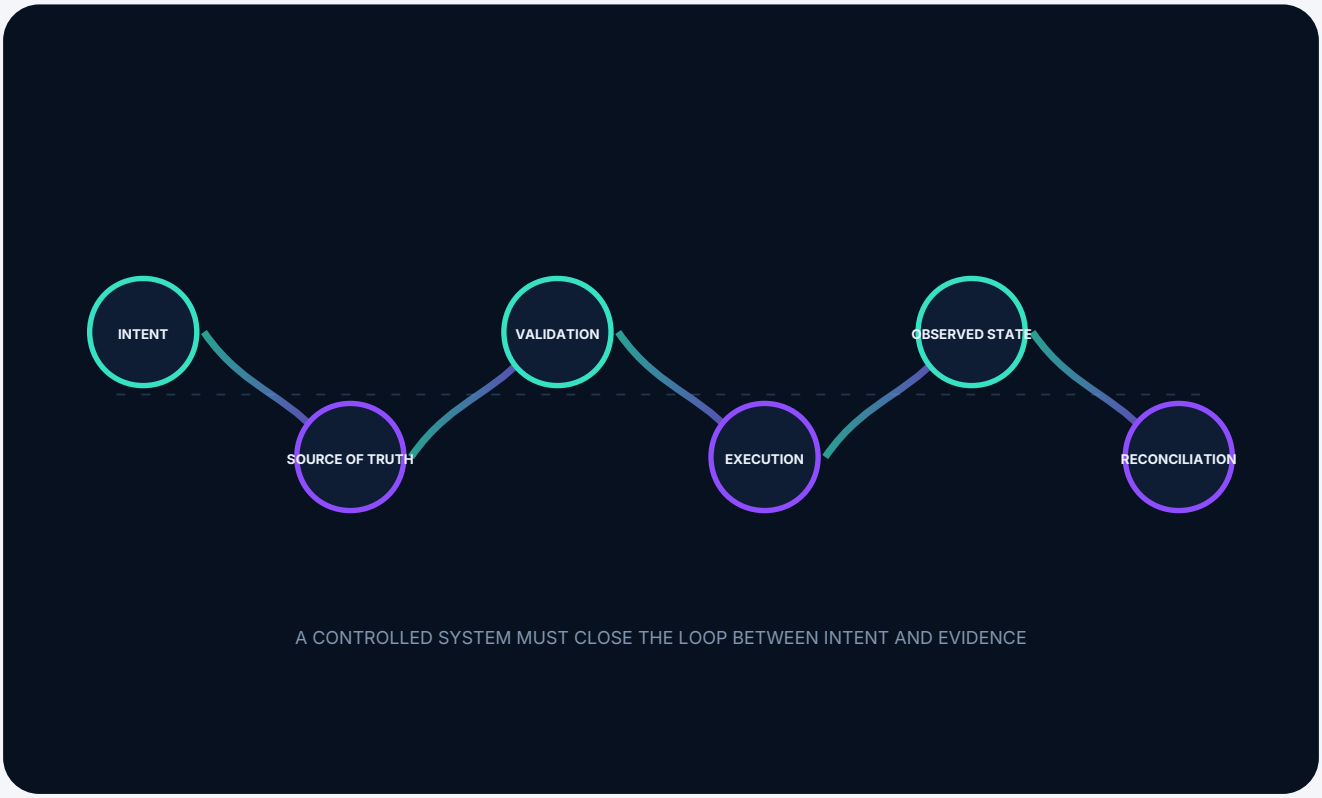

State separation

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-NET-0011

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

The architecture of network access has failed.

Organizations continue to secure their internal networks using flat Layer 2 VLANs and Pre-Shared Keys (PSKs) for Wi-Fi. They grant remote access via legacy VPNs that put unmanaged, potentially compromised devices directly onto the corporate LAN. When a single device is compromised, the attacker instantly gains lateral movement access to the entire flat network, bypassing all perimeter firewalls.

This standard exists because:

1. **PSKs and MAC Auth are Security Theater:** A Wi-Fi PSK is shared, leaks to departing employees, and is easily cracked. MAC authentication is trivially spoofed. Network access MUST be predicated on cryptographic identity (802.1X / EAP-TLS), not shared strings or burned-in hardware addresses.
2. **VPNs are Perimeter Anachronisms:** A VPN grants Layer 3 network access based on a single cryptographic check. Once connected, the user has the keys to the kingdom. Systems MUST replace VPNs with Zero Trust Network Access (ZTNA), which grants Layer 7 application access based on continuous identity and device posture verification.
3. **Flat Networks Breed Ransomware:** A network where an authenticated device can reach any other device is a ransomware acceleration vector. NAC MUST be used to dynamically assign devices to microsegmented VLANs or Software-Defined Perimeters based on their identity, role, and real-time security posture.
4. **Wi-Fi is a Physical Attack Surface:** RF signals do not respect physical walls. An attacker in the parking lot can launch deauthentication attacks, spin up Evil Twins, or exploit WPA3 downgrade vulnerabilities. Wireless engineering MUST include Wireless Intrusion Prevention (WIPS), spectrum analysis, and strict certificate pinning.

This document establishes the **Mythos Access & NAC Standard (MANS)**, a comprehensive, evidence-based methodology for evaluating network access architectures against identity-based isolation, continuous posture validation, and RF resilience.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Network Access Control (NAC) platforms (Cisco ISE, Aruba ClearPass, Forescout)
- 802.1X authentication (EAP-TLS, EAP-PEAP) and RADIUS/TACACS+
- Wireless engineering (Wi-Fi 6E/7, WPA3-Enterprise, OWE, 6GHz spectrum)
- Zero Trust Network Access (ZTNA) and SASE/SSE architectures
- Device posture attestation (TPM, Secure Boot, EDR status)
- Wireless Intrusion Prevention Systems (WIPS) and RF spectrum analysis

1.2 Out of Scope

- General network routing and IPAM (covered in MYTHOS-NET-0010)
- General zero-trust policy architecture (covered in MYTHOS-SEC-0003)
- Endpoint OS hardening (covered in MYTHOS-END-0001)

1.3 Audience

- Network access engineers and WLAN architects
- Security engineers designing NAC and ZTNA pipelines
- Platform engineers managing device fleet onboarding
- Standards bodies seeking a reference access control methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Access Dimensions

Dimension	Definition
802.1X	The IEEE standard for port-based Network Access Control, requiring authentication before a switch port or WLAN allows traffic.
EAP-TLS	Extensible Authentication Protocol - Transport Layer Security. The most secure 802.1X protocol, requiring mutual cryptographic certificate validation by both the client and the authentication server.
ZTNA (Zero Trust Network Access)	A security paradigm replacing VPNs, granting contextual, application-level access based on user identity and device posture, without exposing the corporate L3 network.
Device Posture Attestation	The process of cryptographically verifying the security state (e.g., TPM measurements, EDR running, OS patched) of a device before granting it network access.
Evil Twin	A fraudulent wireless access point that imitates a legitimate one, often using the same SSID, to capture credentials or execute Man-in-the-Middle (MITM) attacks.

2.2 The Access Control Doctrine

> A PSK is a shared liability. A VPN is a wide-open door. A flat VLAN is a ransomware superhighway. If the device is not attested, the network 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; PSKs, flat networks, traditional VPNs, no NAC
1	Basic 802.1X exists, but allows MAC auth fallback and flat VLANs
2	Functional NAC, but lacks continuous posture checks or ZTNA
3	Solid production performance; EAP-TLS enforced, dynamic VLANs, ZTNA, WPA3-Enterprise
4	Strong performance; TPM attestation, WIPS, automated quarantine, 6GHz Wi-Fi 7
5	Best-in-class; sets the standard for mathematically verified, zero-trust network access

Weighting

Category	Weight	Rationale
Identity-Based Access (802.1X/NAC)	20%	PSKs and MAC auth cannot guarantee identity
ZTNA & VPN Replacement	20%	Legacy VPNs grant excessive lateral movement access
Dynamic Microsegmentation	15%	Flat networks allow unchecked ransomware spread
Device Posture & Attestation	15%	A compromised device with valid credentials is a threat
Wireless Security (WPA3/RF)	15%	Wi-Fi is a physical boundary; downgrades and Evil Twins are constant threats
Device Admin (RADIUS/TACACS+)	15%	Local device accounts prevent centralized auditing

Total: 100%

A system MUST score at least 3.0 in Identity-Based Access and ZTNA to be considered for production use.

EVALUATION AND ASSURANCE SYSTEM

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

6 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

Identity-Based Access (802.1X / NAC)

SOURCE WEIGHT 20%

4.1.1
Authentication Rigor

MYTHOS-NET-0011-EVAL-4.1.1

Authentication Rigor

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Identity-Based Access (802.1X / NAC)	4.1.1	20%	Evidence-led

Is network access predicated on cryptographic identity rather than shared secrets?

0
Open Wi-Fi or WPA2-PSK; switches use MAC authentication bypass (MAB)

1
802.1X using PEAP-MSCHAPv2 (vulnerable to rogue RADIUS/AP attacks)

2
EAP-TLS enforced for corporate devices, but IoT devices use PSK

3
EAP-TLS enforced for all entities (humans, devices, IoT); zero MAC auth fallback

4
Continuous authentication: session keys rotated and validated periodically; not just at join time

5
Leading identity: formally verified cryptographic bounds, zero ability to spoof identity, hardware-bound client certificates

ENGINEERING INTERPRETATION

This criterion evaluates whether authentication rigor 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.

4.2 / CAPABILITY FAMILY

ZTNA and VPN Replacement

SOURCE WEIGHT 20%

MYTHOS-NET-0011-EVAL-4.2.1

Remote Access Architecture

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
ZTNA and VPN Replacement	**4.2.1**	**20%**	**Evidence-led**

Are remote users granted network access or application access?

0
Traditional IPsec/SSL VPN placing clients on the flat L3 corporate LAN

1
VPN with split-tunneling, but lateral movement is still possible

2
Basic ZTNA deployed, but coexists with VPN (legacy crutch)

3
100% ZTNA: VPNs removed from the controlled workflow; users connect via identity-aware proxy to specific L7 applications

4
Continuous ZTNA: access is re-evaluated on every request; idle sessions are terminated

5
Leading isolation: formally verified L7 boundaries, zero L3 network exposure, cryptographic proof of application isolation

ENGINEERING INTERPRETATION

This criterion evaluates whether remote access architecture 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

Dynamic Microsegmentation

SOURCE WEIGHT 15%

4.3.1
Lateral Movement Prevention

MYTHOS-NET-0011-EVAL-4.3.1

Lateral Movement Prevention

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Dynamic Microsegmentation	4.3.1	15%	Evidence-led

Can a compromised device communicate with unrelated devices on the same network?

0
Single flat corporate VLAN; any device can ping any other device

1
Basic VLAN segmentation by department (e.g., HR, IT), but broad internal access

2
NAC assigns VLANs dynamically based on identity

3
Software-Defined Perimeter (SDP) or dynamic microsegmentation (Cisco TrustSec/Cilium); devices can only reach explicitly allowed resources

4
Real-time policy updates: if device posture changes, NAC instantly moves it to a quarantine VLAN

5
Leading segmentation: formally verified zero-trust east-west traffic, zero lateral movement possible, mathematical proof of isolation

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

- 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.

4.4 / CAPABILITY FAMILY

Device Posture and Attestation

SOURCE WEIGHT 15%

4.4.1
Health Verification

MYTHOS-NET-0011-EVAL-4.4.1

Health Verification

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Device Posture and Attestation	**4.4.1**	**15%**	**Evidence-led**

Is the security state of the device verified before and during network access?

0
No posture checks; any device with a valid cert is admitted

1
Basic agent checks if antivirus is running

2
OS patch level and EDR status verified at connection time

3
Hardware-backed attestation (TPM 2.0) verifying Secure Boot and OS integrity via RADIUS extensions

4
Continuous posture monitoring; degraded posture instantly triggers NAC quarantine

5
Leading attestation: formally verified hardware roots of trust, zero ability to bypass posture agents, cryptographic proof of system integrity

ENGINEERING INTERPRETATION

This criterion evaluates whether health 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

- attestation evidence, endorsement chain, measurement policy, key-release logs, and verifier records
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test valid, stale, revoked, mismatched, and unavailable attestation paths
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- trusting self-asserted state, stale measurements, and fail-open key release
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- attestation coverage
- verification failure rate
- revocation propagation
- 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

Wireless Security and RF

SOURCE WEIGHT 15%

4.5.1
Airwave Integrity

MYTHOS-NET-0011-EVAL-4.5.1

Airwave Integrity

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Wireless Security and RF	4.5.1	15%	Evidence-led

Is the wireless infrastructure resilient to RF attacks and downgrades?

0
WPA2-PSK; no RF monitoring

1
WPA3 supported, but allows WPA2 transition mode (vulnerable to Dragonblood downgrade)

2
WPA3-Enterprise strictly enforced; 802.11w (PMF) mandatory

3
Wi-Fi 6E/7 (6GHz) utilized for spectrum relief; WIPS detecting and containing Evil Twins

4
Client certificate pinning preventing rogue RADIUS MITM; spectrum AI auto-optimizing channels/power

5
Leading integrity: formally verified cryptographic handshakes, zero downgrade capability, real-time geofencing of RF boundaries

ENGINEERING INTERPRETATION

This criterion evaluates whether airwave integrity 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

- RF survey, channel and power plan, authentication logs, and client telemetry
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test roaming, capacity, interference, authentication, and degraded-band behavior
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- coverage holes, sticky clients, co-channel interference, and insecure fallback
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- coverage compliance
- authentication success
- roaming interruption
- airtime utilization
- 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

Device Admin (RADIUS/ TACACS+)

SOURCE WEIGHT 15%

4.6.1
Infrastructure Access Control

MYTHOS-NET-0011-EVAL-4.6.1

Infrastructure Access Control

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Device Admin (RADIUS/ TACACS+)	4.6.1	15%	Evidence-led

How are engineers authenticated to network devices (switches, routers)?

0
Local accounts on every switch/router; shared passwords

1
Centralized TACACS+/RADIUS, but uses shared passwords

2
Centralized auth with individual named accounts (RADIUS/TACACS+)

3
Just-in-Time (JIT) access: engineers check out a time-boxed session via PAM; commands logged via TACACS+

4
Command-level authorization (TACACS+); dangerous commands (e.g., `reload`, `write erase`) require dual-control

5
Leading governance: formally verified zero local accounts, cryptographic attestation of admin sessions, zero unlogged CLI actions

ENGINEERING INTERPRETATION

This criterion evaluates whether infrastructure access 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.

LIFECYCLE AND ASSURANCE

Part V. The Mythos Access & NAC Suite (MANS)

Traditional network benchmarks measure authentication latency. The MANS evaluates posture enforcement, lateral movement resistance, and RF attack resilience.

5.1 Suite Composition

5.1.1 MANS-Access

- **Rogue Device Test:** 100+ tests attempting to connect a device without a valid client certificate, verifying NAC blocks it at the port level.
- **Posture Degradation Test:** 50+ tests disabling EDR or altering firewall rules on an active device, verifying NAC instantly quarantines it.

5.1.2 MANS-Wireless

- **Evil Twin Simulation:** 100+ tests spinning up a high-power Rogue AP with the same SSID, verifying WIPS detects and contains it, and clients refuse to connect.
- **Downgrade Attempt:** 50+ tests forcing a WPA2 fallback, verifying the client strictly requires WPA3 and drops the connection.

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.