MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

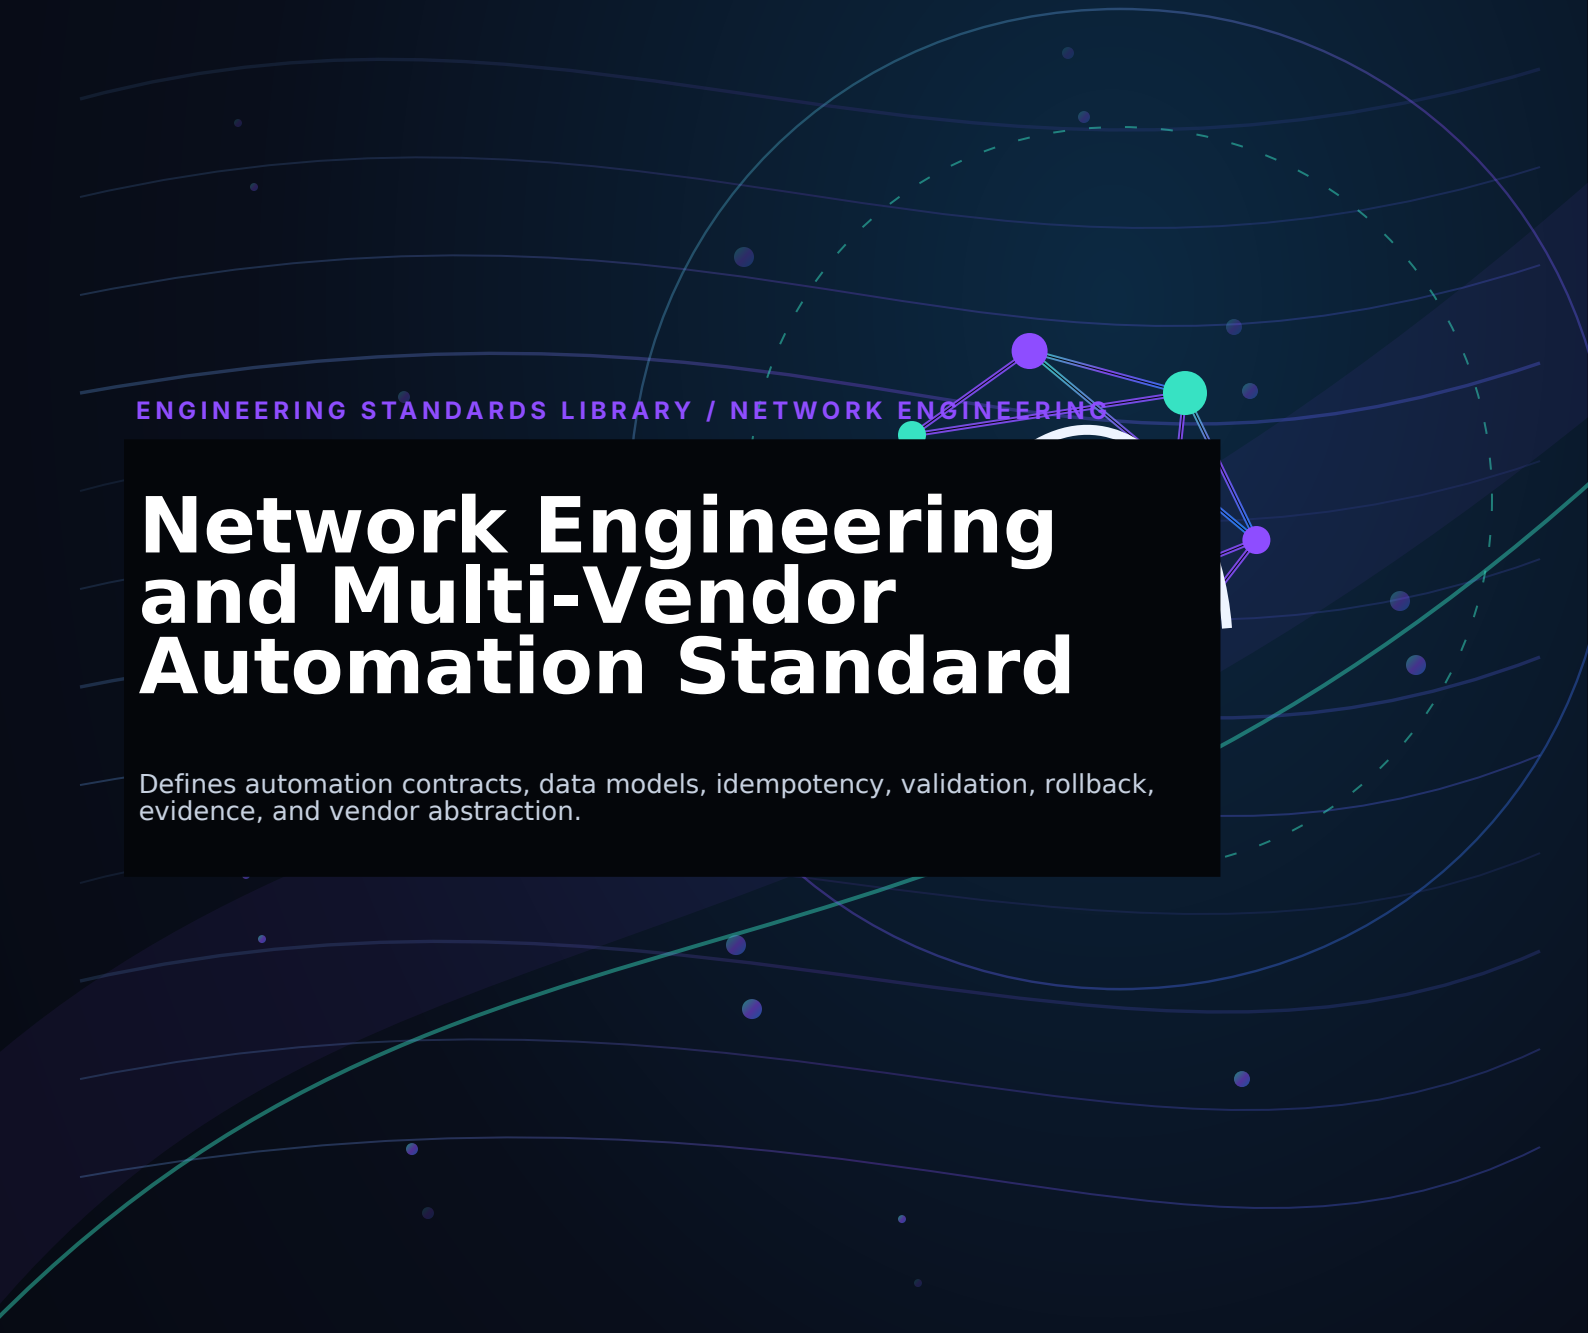

ENGINEERING STANDARDS LIBRARY / NETWORK ENGINEERING

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

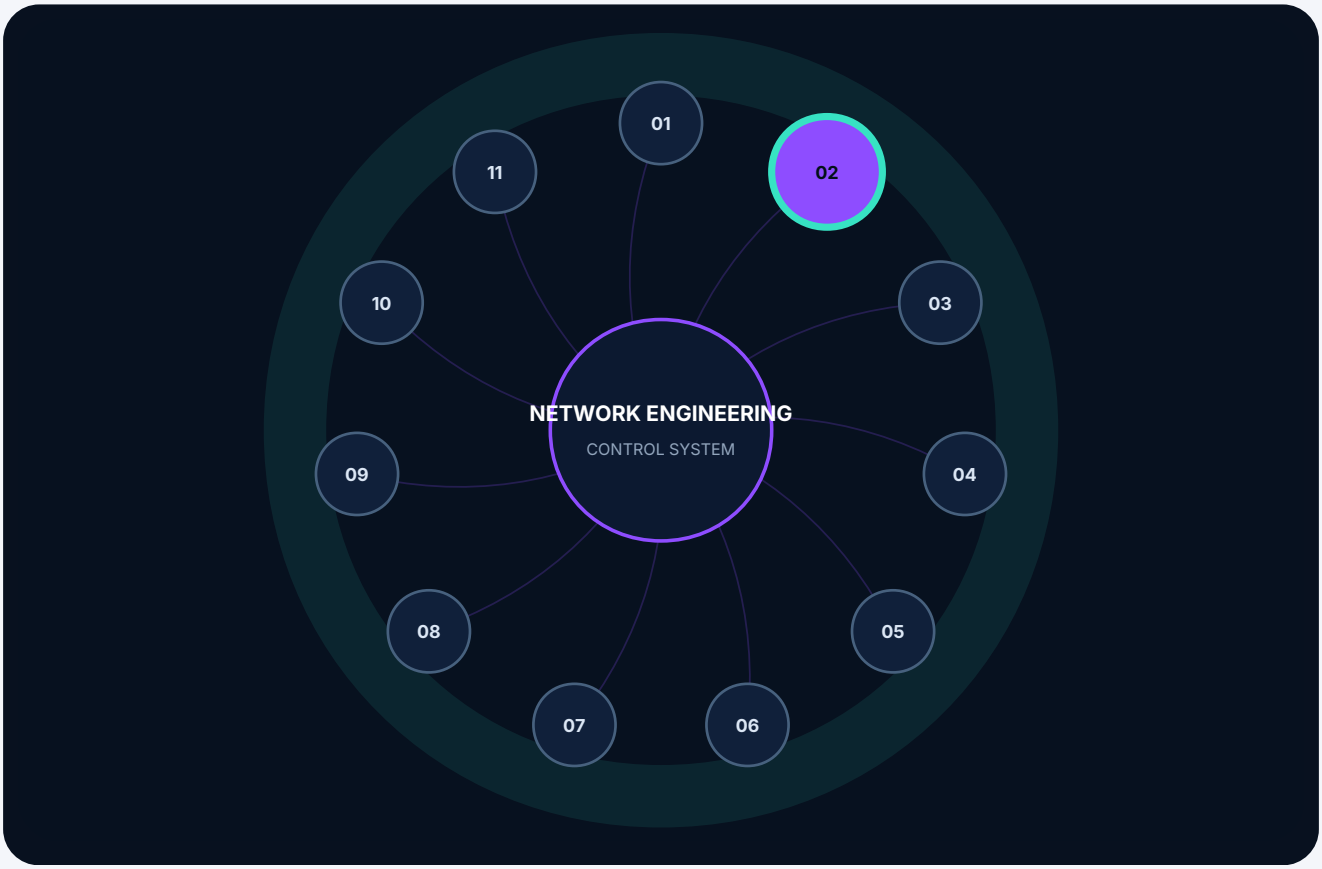

Connected assurance

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

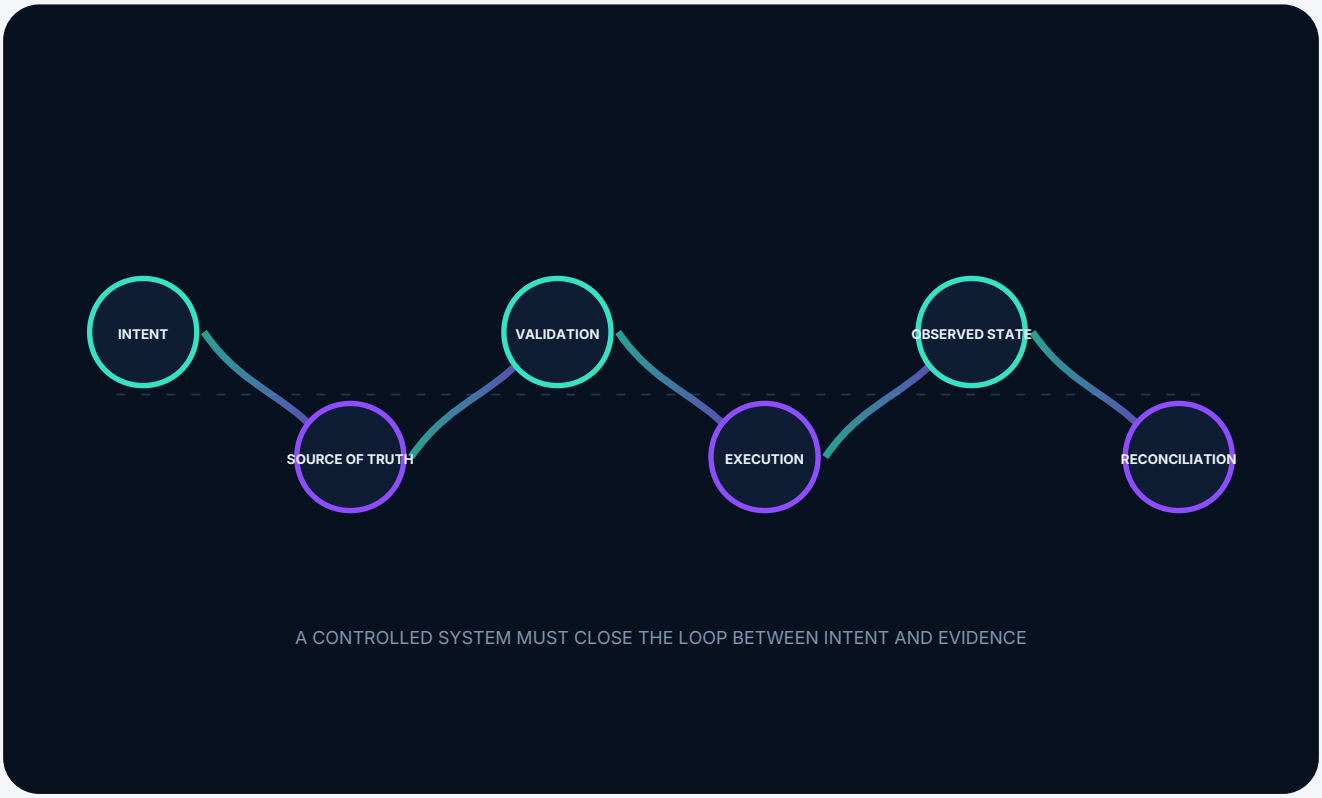

State separation

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.