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

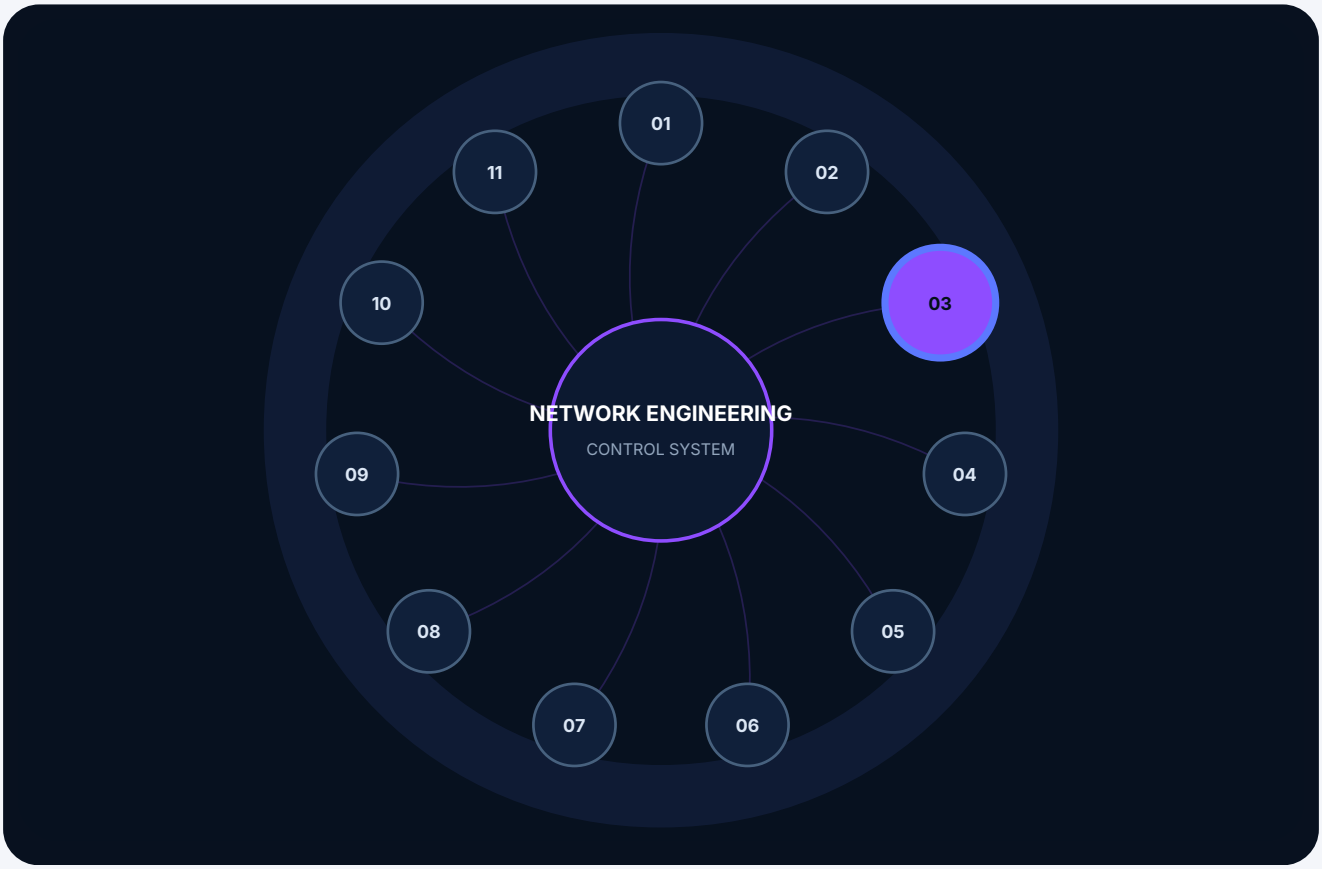

Connected assurance

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

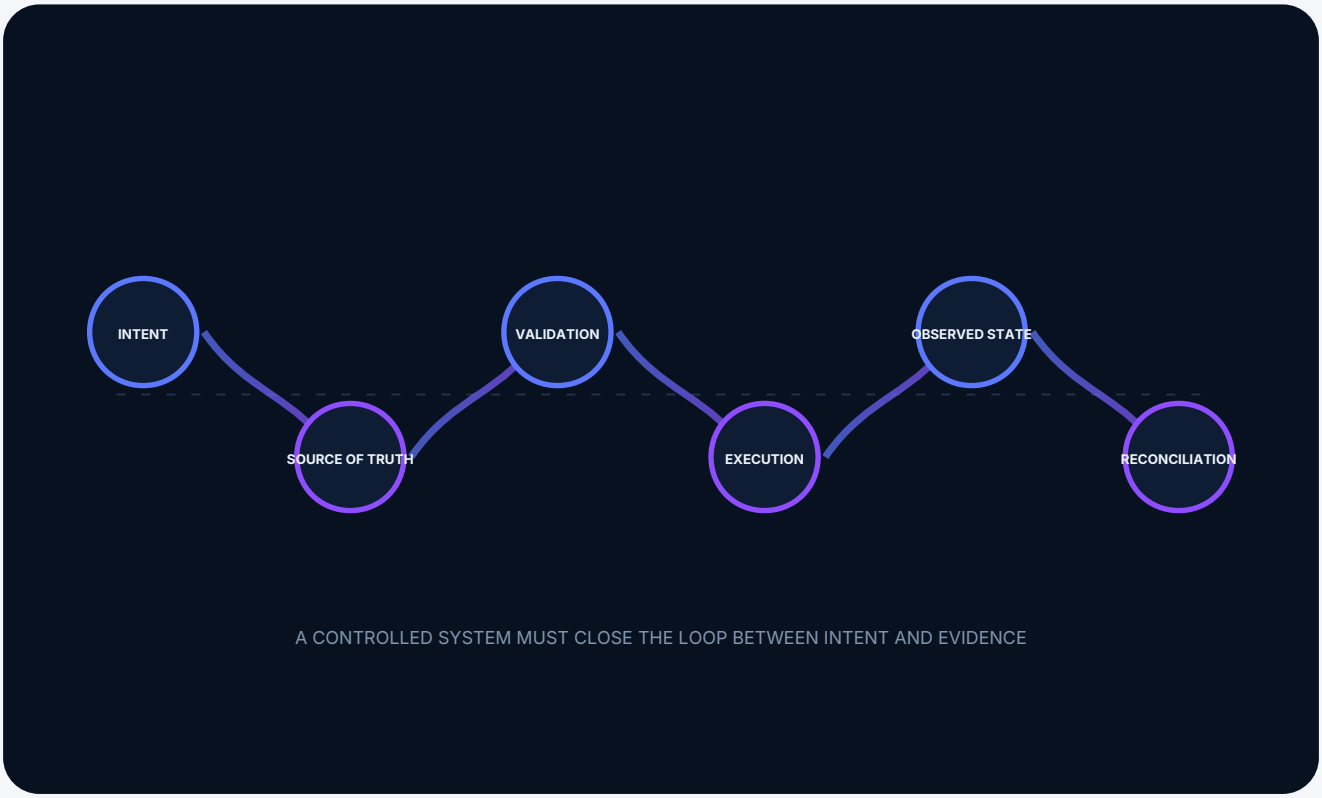

State separation

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.