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

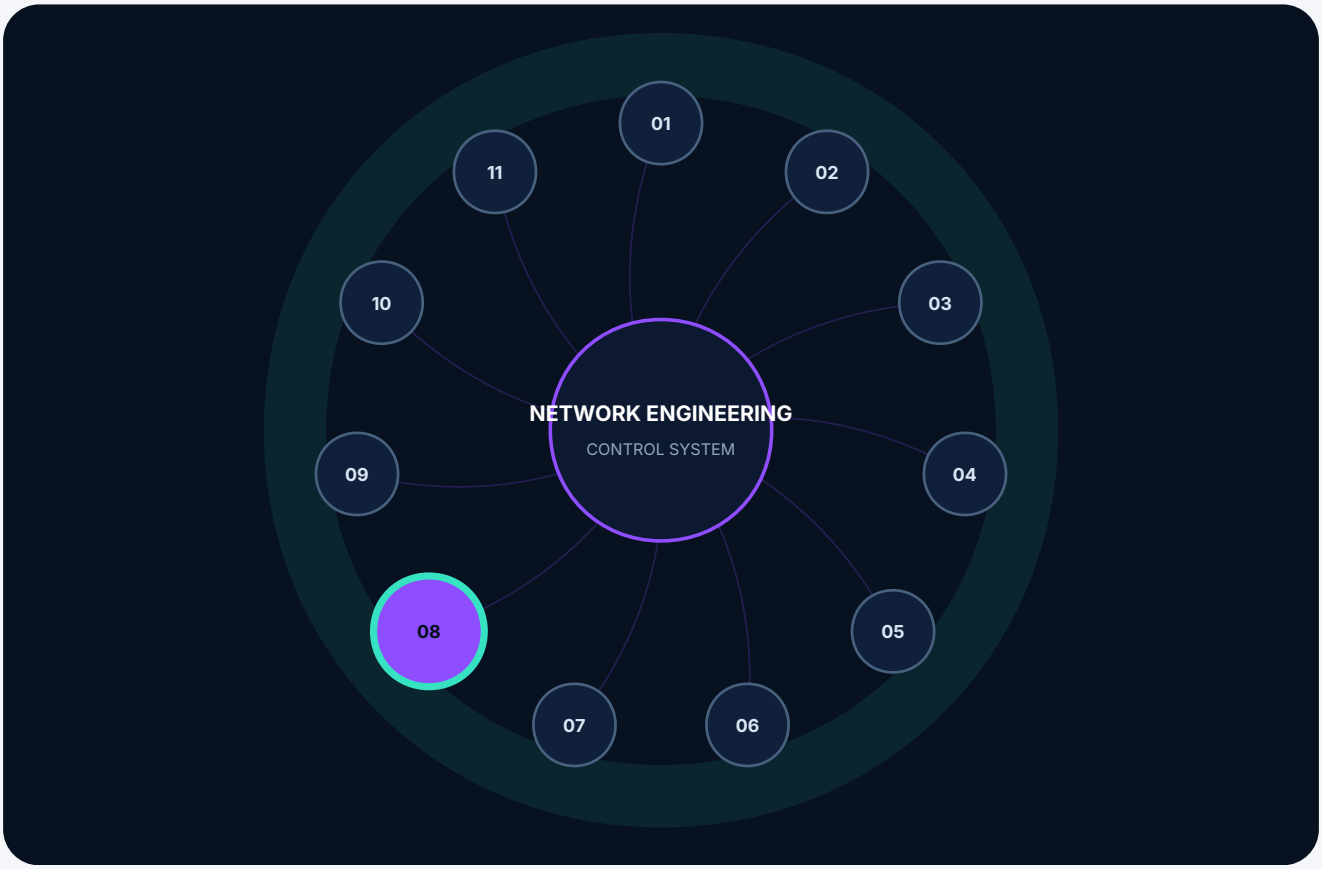

Connected assurance

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

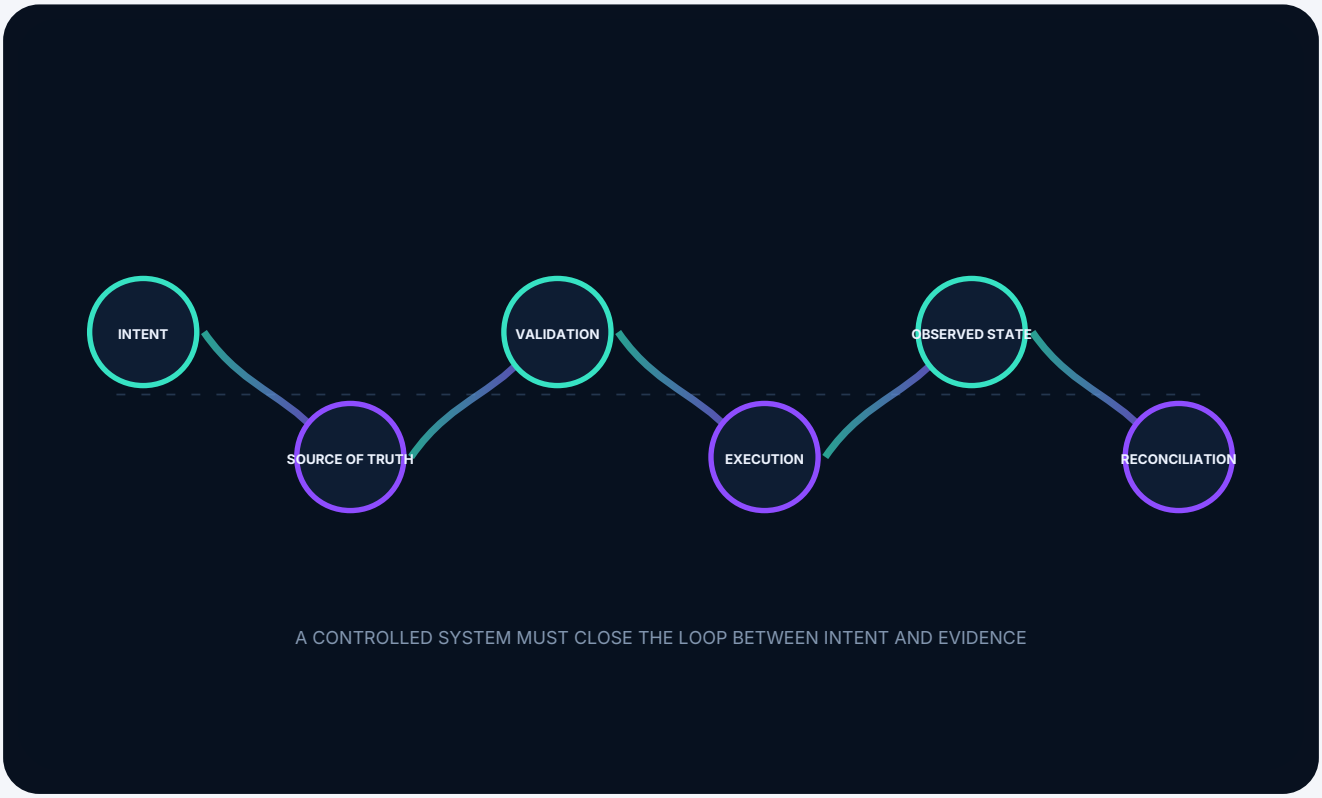

State separation

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.