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

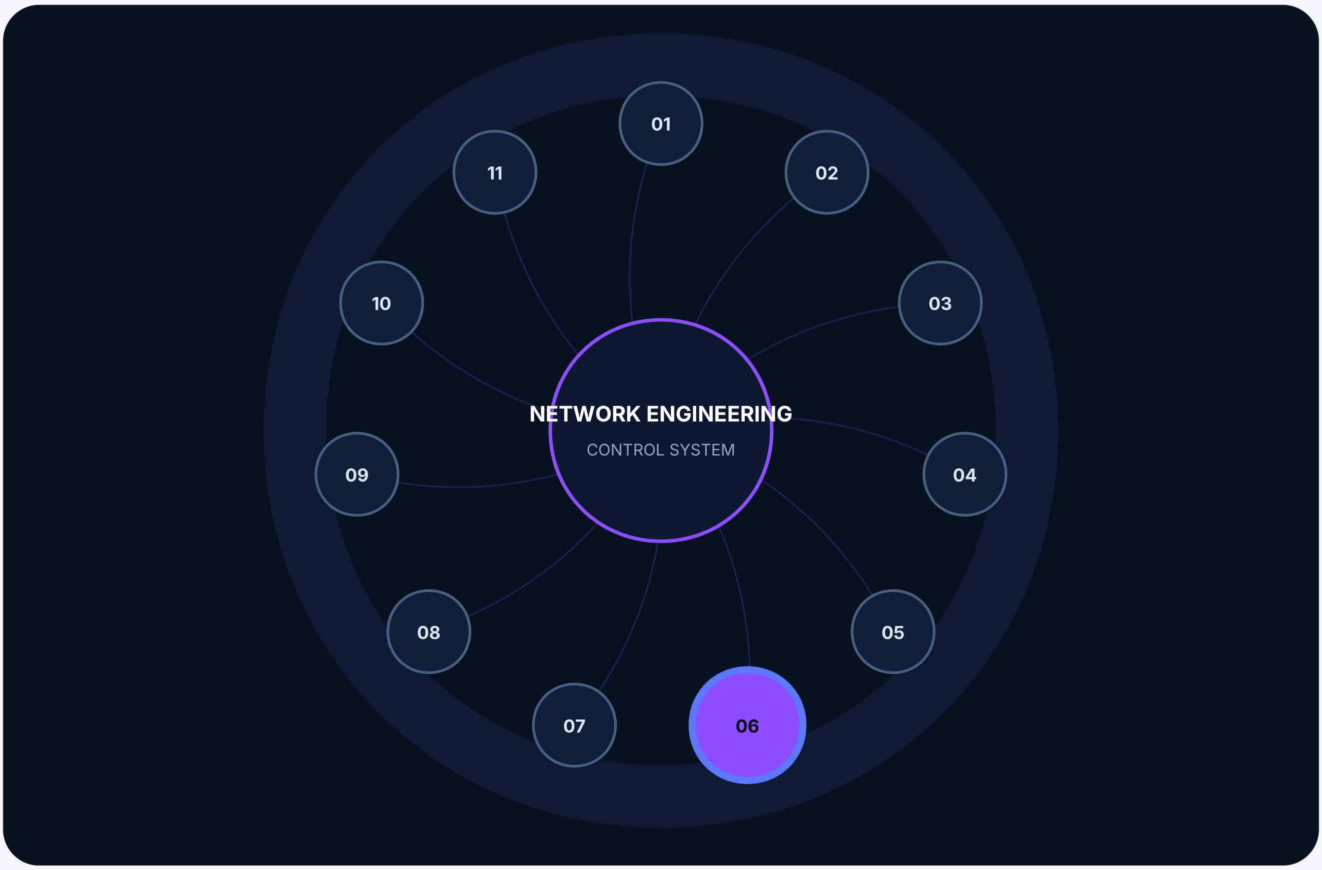

Connected assurance

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

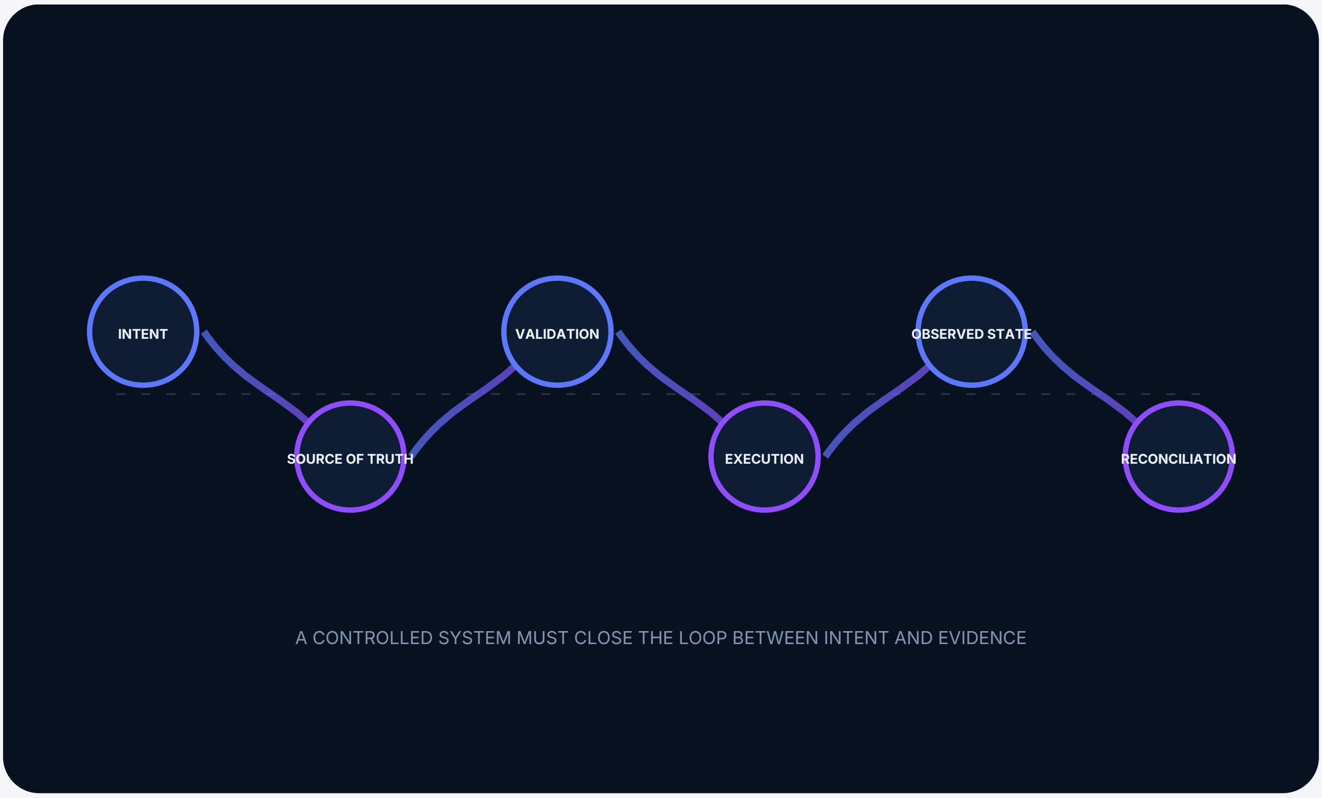

State separation

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.)?

Score	Description
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?

Score	Description
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?

Score	Description
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?

Score	Description
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.

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.