CANONICAL LIVING OVERVIEW GUIDE

Network Source of Truth, Intent, and Drift Reconciliation

A governance and implementation guide for network source of truth, normalized intent, inventory, topology, addressing, services, lifecycle, discovery, reconciliation, drift, approvals, APIs, and evidence-driven automation.

PERMANENT ID
MYTHOS-FG-NET-0017

PUBLICATION CLASS
FIELD GUIDE

VERSION / STATUS
1.0.0-candidate / CANDIDATE

PERMANENT PUBLICATION IDENTITY

Network Source of Truth, Intent, and Drift Reconciliation

Document ID
MYTHOS-FG-NET-0017

Version
1.0.0-candidate

Status
Candidate Focused Field Guide

Review date
2026-10-26

Publication position
This is a living field guide. It is designed for frequent revision while retaining a stable permanent identity. Candidate status means the content is ready for structured use and review but is not represented as an external certification or authority publication.

LEVEL	CLASS	CADENCE	EVIDENCE
L1-L4	FIELD GUIDE	QUARTERLY	SOURCE-BOUND
Foundational through frontier	Practical and educational	Interim updates as needed	Limitations remain visible

Reader orientation

Dimension	Engineering position
Primary domain	MYTHOS-NET - Network Engineering & Architecture
Secondary domain	MYTHOS-DAT; MYTHOS-SWE; MYTHOS-SEC
Audience	Network architects, automation engineers, platform engineers, operations, CMDB owners, and governance teams.
Prerequisites	Network architecture, data modeling, and automation fundamentals.
Publisher / owner	Mythos Systems / Aaron Stovall
Public classification	Public technical guidance

Expected outcomes

Define authority for network objects, fields, lifecycle states, and derived state.

Model physical, logical, service, security, and operational intent with stable identity.

Reconcile discovered and intended state without overwriting truth blindly.

Detect, classify, approve, remediate, and evidence drift safely.

Use source-of-truth data to drive configuration, validation, documentation, and lifecycle.

How to use this guide

Read the guide as a progressive system rather than a disconnected encyclopedia. Foundational material establishes the mental model; practitioner sections convert it into repeatable implementation and operations; advanced sections expose architecture, scale, security, and failure behavior; frontier sections describe technologies whose evidence or maturity is still evolving.

Suggested reading path
New practitioners should follow the chapters in order and complete the laboratories. Experienced engineers may begin with the architecture, failure, validation, or frontier sections, then use the related-document map to deepen a specific topic.

Learning and assurance model

Level	Reader outcome
L1 - Foundational	Terminology, first principles, component roles, packet or state flow, and simple working examples.
L2 - Practitioner	Implementation, configuration, automation, testing, troubleshooting, and operational evidence.
L3 - Advanced	Architecture, scale, concurrency, resilience, threat models, performance, and complex migrations.
L4 - Frontier	Emerging systems, research directions, technical uncertainty, readiness gates, and adoption signals.

Contents

CHAPTER 01

Authority and truth models

Define what is authoritative, observed, derived, proposed, and historical.

Chapter operating position
Define what is authoritative, observed, derived, proposed, and historical.

1.1 Sources and authority boundaries

Sources and authority boundaries must be treated as an engineering mechanism rather than a product checkbox. Assign field-level ownership across network source of truth, IPAM, CMDB, cloud, identity, controllers, and devices. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for sources and authority boundaries.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating sources and authority boundaries as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for sources and authority boundaries.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

1.2 Intended, applied, and operational state

Intended, applied, and operational state must be treated as an engineering mechanism rather than a product checkbox. Use datastore and origin concepts to distinguish desired configuration, active state, and measured behavior. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for intended, applied, and operational state.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Implementation sketch

```
object:
  id: dev_01J2STABLEIDENTITY
  type: network_device
  canonical:
    role: leaf
    site: dc-west-1
    lifecycle: active
  observed:
    hostname: leaf-17
    software: 10.4.2
    observed_at: 2026-07-26T16:00:00Z
    source: gnmi
  drift:
    classification: unauthorized-change
    confidence: high
    disposition: review-before-enforcement
```

Failure patterns

- Treating intended, applied, and operational state as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for intended, applied, and operational state.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

1.3 Stable identity

Stable identity must be treated as an engineering mechanism rather than a product checkbox. Create immutable object IDs independent of hostname, address, serial, site, owner, or role changes. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for stable identity.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating stable identity as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for stable identity.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

1.4 Confidence and provenance

Confidence and provenance must be treated as an engineering mechanism rather than a product checkbox. Record source, method, timestamp, version, confidence, reviewer, and transformation for every important fact. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

Define authority, ownership, scope, dependencies, and acceptance criteria for confidence and provenance.

Model the expected state and traffic path before implementation or troubleshooting.

Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.

Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.

Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.

Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

Treating confidence and provenance as a vendor feature instead of an end-to-end system.

Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for confidence and provenance.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

Chapter 01 laboratory and review

Practical laboratory

Create an authority matrix for devices, interfaces, links, prefixes, VLANs, circuits, software, owners, and services across five systems.

Laboratory evidence package

- Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.
- Versioned configuration or code with explanatory comments and reproducible setup instructions.
- Nominal-path validation plus at least two injected failure or boundary conditions.
- Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.
- A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore cryptographically verifiable state, decentralized authority, semantic knowledge graphs, and local-first truth.

Review gate

Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 02

Network data model

Represent the infrastructure and services that must be designed, operated, and changed.

Chapter operating position
Represent the infrastructure and services that must be designed, operated, and changed.

2.1 Physical inventory

Physical inventory must be treated as an engineering mechanism rather than a product checkbox. Model sites, rooms, racks, devices, modules, ports, cables, circuits, power, optics, and spares. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for physical inventory.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating physical inventory as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for physical inventory.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

2.2 Logical topology

Logical topology must be treated as an engineering mechanism rather than a product checkbox. Model interfaces, LAGs, VLANs, VRFs, prefixes, tunnels, routing sessions, policies, and dependencies. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for logical topology.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Implementation sketch

```
object:
  id: dev_01J2STABLEIDENTITY
  type: network_device
  canonical:
    role: leaf
    site: dc-west-1
    lifecycle: active
  observed:
    hostname: leaf-17
    software: 10.4.2
    observed_at: 2026-07-26T16:00:00Z
    source: gnmi
  drift:
    classification: unauthorized-change
    confidence: high
    disposition: review-before-enforcement
```

Failure patterns

- Treating logical topology as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for logical topology.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

2.3 Service and intent

Service and intent must be treated as an engineering mechanism rather than a product checkbox. Model connectivity, isolation, QoS, security, application paths, customers, owners, objectives, and lifecycle. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for service and intent.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating service and intent as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for service and intent.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

2.4 Extensibility and normalization

Extensibility and normalization must be treated as an engineering mechanism rather than a product checkbox. Use typed custom fields, plugins, schemas, relationships, namespaces, and controlled vendor-specific extensions. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

Define authority, ownership, scope, dependencies, and acceptance criteria for extensibility and normalization.

Model the expected state and traffic path before implementation or troubleshooting.

Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.

Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.

Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.

Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

Treating extensibility and normalization as a vendor feature instead of an end-to-end system.

Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for extensibility and normalization.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

Chapter 02 laboratory and review

Practical laboratory
Model a campus, WAN, EVPN fabric, and cloud connectivity using one normalized schema. Identify where vendor-specific data remains necessary.

Laboratory evidence package

- Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.
- Versioned configuration or code with explanatory comments and reproducible setup instructions.
- Nominal-path validation plus at least two injected failure or boundary conditions.
- Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.
- A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Extend the model to digital twins, photonic layers, AI agents, quantum-safe assets, and dynamic service intent.

Review gate
Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 03

Lifecycle and workflow

Make every object and service progress through explicit states and approvals.

Chapter operating position
Make every object and service progress through explicit states and approvals.

3.1 Object lifecycle

Object lifecycle must be treated as an engineering mechanism rather than a product checkbox. Use planned, reserved, staged, active, maintenance, deprecated, quarantined, failed, retired, and archived states. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for object lifecycle.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating object lifecycle as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for object lifecycle.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

3.2 Change proposals

Change proposals must be treated as an engineering mechanism rather than a product checkbox. Record desired deltas, rationale, risk, dependencies, evidence, approvals, schedule, and rollback. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for change proposals.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Implementation sketch

```
object:
  id: dev_01J2STABLEIDENTITY
  type: network_device
  canonical:
    role: leaf
    site: dc-west-1
    lifecycle: active
  observed:
    hostname: leaf-17
    software: 10.4.2
    observed_at: 2026-07-26T16:00:00Z
    source: gnmi
  drift:
    classification: unauthorized-change
    confidence: high
    disposition: review-before-enforcement
```

Failure patterns

- Treating change proposals as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for change proposals.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

3.3 Ownership and separation of duties

Ownership and separation of duties must be treated as an engineering mechanism rather than a product checkbox. Define requestor, designer, reviewer, approver, executor, verifier, operator, and auditor responsibilities. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for ownership and separation of duties.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating ownership and separation of duties as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for ownership and separation of duties.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

3.4 Temporal state and history

Temporal state and history must be treated as an engineering mechanism rather than a product checkbox. Retain effective time, discovered time, planned time, supersession, snapshots, and historical reconstruction. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

Define authority, ownership, scope, dependencies, and acceptance criteria for temporal state and history.

Model the expected state and traffic path before implementation or troubleshooting.

Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.

Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.

Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.

Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

Treating temporal state and history as a vendor feature instead of an end-to-end system.

Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for temporal state and history.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

Chapter 03 laboratory and review

Practical laboratory
Create a lifecycle workflow for a new site and for decommissioning a device. Include reservations, approvals, implementation, verification, and archive.

Laboratory evidence package

- Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.
- Versioned configuration or code with explanatory comments and reproducible setup instructions.
- Nominal-path validation plus at least two injected failure or boundary conditions.
- Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.
- A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore event-sourced source of truth, bitemporal data, policy engines, and machine-verifiable approvals.

Review gate
Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 04

Discovery and ingestion

Collect existing state without treating discovery as automatically authoritative.

Chapter operating position
Collect existing state without treating discovery as automatically authoritative.

4.1 Discovery methods

Discovery methods must be treated as an engineering mechanism rather than a product checkbox. Use APIs, CLI, SNMP, telemetry, LLDP, routing, ARP or ND, cloud inventory, controllers, and imported records. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for discovery methods.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating discovery methods as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for discovery methods.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

4.2 Normalization and correlation

Normalization and correlation must be treated as an engineering mechanism rather than a product checkbox. Map vendor names, identifiers, addresses, interfaces, serials, topology, and services to canonical objects. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for normalization and correlation.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Implementation sketch

```
object:
  id: dev_01J2STABLEIDENTITY
  type: network_device
  canonical:
    role: leaf
    site: dc-west-1
    lifecycle: active
  observed:
    hostname: leaf-17
    software: 10.4.2
    observed_at: 2026-07-26T16:00:00Z
    source: gnmi
  drift:
    classification: unauthorized-change
    confidence: high
    disposition: review-before-enforcement
```

Failure patterns

- Treating normalization and correlation as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for normalization and correlation.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

4.3 Conflict and duplicate handling

Conflict and duplicate handling must be treated as an engineering mechanism rather than a product checkbox. Detect multiple identities, reused names, stale records, partial observations, and contradictory sources. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for conflict and duplicate handling.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating conflict and duplicate handling as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for conflict and duplicate handling.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

4.4 Ingestion safety

Ingestion safety must be treated as an engineering mechanism rather than a product checkbox. Stage changes, record provenance, require confidence, prevent mass overwrite, and route uncertainty for review. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

Define authority, ownership, scope, dependencies, and acceptance criteria for ingestion safety.

Model the expected state and traffic path before implementation or troubleshooting.

Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.

Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.

Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.

Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

Treating ingestion safety as a vendor feature instead of an end-to-end system.

Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for ingestion safety.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

Chapter 04 laboratory and review

Practical laboratory
Ingest data from three inconsistent sources, correlate objects, preserve conflicts, and produce a review queue rather than overwriting canonical data.

Laboratory evidence package

- Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.
- Versioned configuration or code with explanatory comments and reproducible setup instructions.
- Nominal-path validation plus at least two injected failure or boundary conditions.
- Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.
- A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore graph matching, probabilistic correlation, local discovery agents, and privacy-preserving federated inventory.

Review gate
Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 05

Drift detection and classification

Compare intended, configured, operational, and service state with useful semantics.

Chapter operating position
Compare intended, configured, operational, and service state with useful semantics.

5.1 Structural and configuration drift

Structural and configuration drift must be treated as an engineering mechanism rather than a product checkbox. Detect missing, extra, modified, reordered, defaulted, or unsupported objects and settings. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for structural and configuration drift.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating structural and configuration drift as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for structural and configuration drift.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

5.2 Operational and behavioral drift

Operational and behavioral drift must be treated as an engineering mechanism rather than a product checkbox. Detect route, neighbor, link, client, path, performance, policy, and service divergence despite matching config. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for operational and behavioral drift.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Implementation sketch

```
object:
  id: dev_01J2STABLEIDENTITY
  type: network_device
  canonical:
    role: leaf
    site: dc-west-1
    lifecycle: active
  observed:
    hostname: leaf-17
    software: 10.4.2
    observed_at: 2026-07-26T16:00:00Z
    source: gnmi
  drift:
    classification: unauthorized-change
    confidence: high
    disposition: review-before-enforcement
```

Failure patterns

- Treating operational and behavioral drift as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for operational and behavioral drift.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

5.3 Authorized and unauthorized change

Authorized and unauthorized change must be treated as an engineering mechanism rather than a product checkbox. Use change windows, tickets, actors, approvals, controller events, and signatures to classify drift. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for authorized and unauthorized change.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating authorized and unauthorized change as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for authorized and unauthorized change.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

5.4 Severity and impact

Severity and impact must be treated as an engineering mechanism rather than a product checkbox. Rank by service, security, blast radius, persistence, confidence, and failure or exploit potential. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

Define authority, ownership, scope, dependencies, and acceptance criteria for severity and impact.

Model the expected state and traffic path before implementation or troubleshooting.

Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.

Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.

Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.

Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

Treating severity and impact as a vendor feature instead of an end-to-end system.

Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for severity and impact.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

Chapter 05 laboratory and review

Practical laboratory
Create drift scenarios involving harmless defaults, urgent break-glass change, unauthorized policy, failed link, and stale source data. Classify each.

Laboratory evidence package

- Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.
- Versioned configuration or code with explanatory comments and reproducible setup instructions.
- Nominal-path validation plus at least two injected failure or boundary conditions.
- Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.
- A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Extend drift analysis with reachability proof, attack-path impact, predictive risk, and AI explanations grounded in evidence.

Review gate
Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 06

Reconciliation and remediation

Resolve divergence without assuming the source, device, or observer is always correct.

Chapter operating position
Resolve divergence without assuming the source, device, or observer is always correct.

6.1 Reconciliation decisions

Reconciliation decisions must be treated as an engineering mechanism rather than a product checkbox. Choose adopt observed, enforce intended, merge, quarantine, defer, override, or escalate with explicit evidence. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for reconciliation decisions.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating reconciliation decisions as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for reconciliation decisions.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

6.2 Change generation

Change generation must be treated as an engineering mechanism rather than a product checkbox. Produce minimal deltas, dependency ordering, validation, transaction, canary, rollout, and rollback. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for change generation.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Implementation sketch

```
object:
  id: dev_01J2STABLEIDENTITY
  type: network_device
  canonical:
    role: leaf
    site: dc-west-1
    lifecycle: active
  observed:
    hostname: leaf-17
    software: 10.4.2
    observed_at: 2026-07-26T16:00:00Z
    source: gnmi
  drift:
    classification: unauthorized-change
    confidence: high
    disposition: review-before-enforcement
```

Failure patterns

- Treating change generation as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for change generation.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

6.3 Human authority and exceptions

Human authority and exceptions must be treated as an engineering mechanism rather than a product checkbox. Require approvals by risk, preserve emergency changes, set expiry, and prevent silent permanent exceptions. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for human authority and exceptions.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating human authority and exceptions as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for human authority and exceptions.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

6.4 Closure and proof

Closure and proof must be treated as an engineering mechanism rather than a product checkbox. Re-read state, test reachability and service, update history, resolve drift, and monitor recurrence. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

Define authority, ownership, scope, dependencies, and acceptance criteria for closure and proof.

Model the expected state and traffic path before implementation or troubleshooting.

Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.

Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.

Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.

Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

Treating closure and proof as a vendor feature instead of an end-to-end system.

Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for closure and proof.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

Chapter 06 laboratory and review

Practical laboratory
Build a reconciliation workflow for configuration and operational drift with dry run, approval, remediation, rollback, and post-change proof.

Laboratory evidence package

- Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.
- Versioned configuration or code with explanatory comments and reproducible setup instructions.
- Nominal-path validation plus at least two injected failure or boundary conditions.
- Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.
- A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore autonomous reconciliation constrained by invariants, formal plans, capability grants, and reversible actions.

Review gate
Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 07

APIs, integrations, and platform operation

Expose source-of-truth capabilities safely to humans, automation, controllers, and agents.

Chapter operating position
Expose source-of-truth capabilities safely to humans, automation, controllers, and agents.

7.1 API design

API design must be treated as an engineering mechanism rather than a product checkbox. Use typed objects, filtering, pagination, transactions, webhooks, idempotency, optimistic locking, and versioning. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for api design.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating api design as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for api design.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

7.2 Event and workflow integration

Event and workflow integration must be treated as an engineering mechanism rather than a product checkbox. Publish changes, approvals, lifecycle events, drift, and evidence to automation, ticketing, observability, and Git. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for event and workflow integration.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Implementation sketch

```
object:
  id: dev_01J2STABLEIDENTITY
  type: network_device
  canonical:
    role: leaf
    site: dc-west-1
    lifecycle: active
  observed:
    hostname: leaf-17
    software: 10.4.2
    observed_at: 2026-07-26T16:00:00Z
    source: gnmi
  drift:
    classification: unauthorized-change
    confidence: high
    disposition: review-before-enforcement
```

Failure patterns

- Treating event and workflow integration as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for event and workflow integration.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

7.3 Security and tenancy

Security and tenancy must be treated as an engineering mechanism rather than a product checkbox. Protect credentials, object scopes, fields, secrets, customers, approvals, audit, and administrative actions. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for security and tenancy.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating security and tenancy as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for security and tenancy.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

7.4 Backup and recovery

Backup and recovery must be treated as an engineering mechanism rather than a product checkbox. Protect database, media, plugins, schemas, configuration, secrets, audit, and tested restoration. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

Define authority, ownership, scope, dependencies, and acceptance criteria for backup and recovery.

Model the expected state and traffic path before implementation or troubleshooting.

Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.

Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.

Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.

Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

Treating backup and recovery as a vendor feature instead of an end-to-end system.

Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for backup and recovery.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

Chapter 07 laboratory and review

Practical laboratory
Integrate a source-of-truth event with an automation workflow and ticket. Test duplicate events, stale version, authorization, and recovery.

Laboratory evidence package

- Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.
- Versioned configuration or code with explanatory comments and reproducible setup instructions.
- Nominal-path validation plus at least two injected failure or boundary conditions.
- Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.
- A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore signed events, decentralized replicas, sovereign domains, and agent-accessed knowledge graphs.

Review gate
Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 08

Governance, quality, and adoption

Keep the source of truth useful, current, trusted, and embedded in engineering work.

Chapter operating position
Keep the source of truth useful, current, trusted, and embedded in engineering work.

8.1 Data quality objectives

Data quality objectives must be treated as an engineering mechanism rather than a product checkbox. Measure completeness, validity, uniqueness, consistency, freshness, lineage, ownership, and reconciliation time. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for data quality objectives.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating data quality objectives as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for data quality objectives.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

8.2 Operating model

Operating model must be treated as an engineering mechanism rather than a product checkbox. Define product ownership, domain stewards, platform engineering, review boards, users, support, and training. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for operating model.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Implementation sketch

```
object:
  id: dev_01J2STABLEIDENTITY
  type: network_device
  canonical:
    role: leaf
    site: dc-west-1
    lifecycle: active
  observed:
    hostname: leaf-17
    software: 10.4.2
    observed_at: 2026-07-26T16:00:00Z
    source: gnmi
  drift:
    classification: unauthorized-change
    confidence: high
    disposition: review-before-enforcement
```

Failure patterns

- Treating operating model as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.
- Ignoring reverse paths, external dependencies, identity, time, or partial failure.
- Changing several variables before preserving original evidence and stating a hypothesis.
- Allowing automation or AI to act on stale, ambiguous, or unauthorized data.
- Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for operating model.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

8.3 Adoption and migration

Adoption and migration must be treated as an engineering mechanism rather than a product checkbox. Import legacy systems, establish authority gradually, prioritize high-value workflows, and retire shadow data. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

- Define authority, ownership, scope, dependencies, and acceptance criteria for adoption and migration.
- Model the expected state and traffic path before implementation or troubleshooting.
- Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.
- Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.
- Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.
- Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

- Treating adoption and migration as a vendor feature instead of an end-to-end system.
- Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for adoption and migration.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

8.4 Continuous improvement

Continuous improvement must be treated as an engineering mechanism rather than a product checkbox. Review schemas, workflows, exceptions, incidents, automation outcomes, and user friction. The design should name the authoritative inputs, state transitions, dependencies, limits, ownership, telemetry, expected output, and conditions that make the result unsafe or inconclusive.

The guide traces this subject through the authority, canonical object, observation, comparison, decision, remediation, and history chain. Configuration, controller health, or a green dashboard is never sufficient proof by itself. The engineer must reconcile intended state with protocol or platform state, installed forwarding or policy, real transactions, failure behavior, and the experience of the affected user, device, service, or application.

Implementation begins with a minimal, known-good baseline and explicit invariants. Introduce one capability or dependency at a time, validate positive and negative cases, inject realistic failures, measure recovery, and preserve before-and-after evidence. Automation must be idempotent, bounded, observable, reversible, and able to stop safely when inputs are stale or outcomes diverge.

Troubleshooting should locate the first divergence from the expected transaction or state machine. Inspect both directions and all authority boundaries; account for caching, eventual consistency, tunnels, load sharing, policy, identity, offload, time, and partial failure. Closure requires causal evidence, restored service, rollback proof, residual-risk disclosure, and prevention or monitoring actions.

Engineering practices

Define authority, ownership, scope, dependencies, and acceptance criteria for continuous improvement.

Model the expected state and traffic path before implementation or troubleshooting.

Validate intent, control state, applied policy or forwarding state, and end-to-end service independently.

Test normal, denied, degraded, failed, recovery, upgrade, and rollback conditions.

Retain topology, configuration, packet, telemetry, log, identity, timing, and service evidence.

Automate through typed data, deterministic gates, canaries, bounded concurrency, and reversible changes.

Failure patterns

Treating continuous improvement as a vendor feature instead of an end-to-end system.

Relying on intended configuration or controller state without proving applied and operational behavior.

Ignoring reverse paths, external dependencies, identity, time, or partial failure.

Changing several variables before preserving original evidence and stating a hypothesis.

Allowing automation or AI to act on stale, ambiguous, or unauthorized data.

Declaring success after one probe without negative tests, scale checks, rollback, and recurrence monitoring.

Validation and evidence

Method	Expected evidence
Examine	Requirements, topology, models, policy, configuration, ownership, and dependencies for continuous improvement.
Observe	Control-plane or platform state, applied policy, forwarding state, telemetry, logs, and endpoint behavior.
Test	Expected traffic, prohibited traffic, dependency loss, partial failure, convergence, recovery, and rollback.
Measure	Availability, reachability, latency, loss, convergence, scale, resource use, change success, and user experience.
Reconcile	Intended state against observed platform, network, security, identity, application, and evidence records.

Chapter 08 laboratory and review

Practical laboratory

Create a source-of-truth adoption roadmap with authority milestones, data-quality metrics, workflows, training, and system retirement gates.

Laboratory evidence package

- Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.
- Versioned configuration or code with explanatory comments and reproducible setup instructions.
- Nominal-path validation plus at least two injected failure or boundary conditions.
- Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.
- A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Evaluate self-describing infrastructure, intent marketplaces, continuous verification, and machine-readable organizational authority.

Review gate

Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

Integrated learning roadmap

The guide is intended to be revisited. First establish fluency, then build a small implementation, operate it under failure, automate repeatable work, and finally evaluate emerging techniques against a stable baseline. Progress is demonstrated through evidence rather than reading completion alone.

Stage	Demonstrated capability
Foundation	Explain the system from first principles and trace its critical path without relying on product terminology.
Build	Implement the smallest representative system with explicit contracts, identities, telemetry, and tests.
Operate	Diagnose injected failures, recover safely, and preserve evidence of the causal chain.
Automate	Convert a proven manual workflow into bounded, idempotent, observable automation.
Architect	Design for scale, resilience, security, interoperability, and lifecycle change.
Explore	Evaluate frontier techniques using stated hypotheses, limitations, and adoption gates.

Source authority register

Source policy
Official standards, specifications, and project documentation remain with their issuing organizations. Mythos Systems provides original engineering interpretation, synthesis, implementation guidance, and relationship mapping. Source verification dates do not guarantee that an authority has not changed since review.

01. IETF / RFC Editor - RFC 8342 - Network Management Datastore Architecture
Role: Intended, applied, operational, and origin datastore model.
Verified: 2026-07-26
Official source: https://www.rfc-editor.org/info/rfc8342

02. IETF / RFC Editor - RFC 8969 - Automation with YANG
Role: Operator-oriented service and network automation framework.
Verified: 2026-07-26
Official source: https://www.rfc-editor.org/info/rfc8969

03. IETF / RFC Editor - RFC 7950 - YANG 1.1
Role: YANG data-modeling language authority.
Verified: 2026-07-26
Official source: https://www.rfc-editor.org/info/rfc7950

04. OpenConfig - OpenConfig data models
Role: Vendor-neutral network configuration and operational-state models.
Verified: 2026-07-26
Official source: https://www.openconfig.net/projects/models/

05. OpenConfig - gNMI specification
Role: gRPC configuration retrieval and streaming telemetry protocol.
Verified: 2026-07-26
Official source: https://www.openconfig.net/docs/gnmi/gnmi-specification/

06. NetBox Labs - NetBox Documentation
Role: Official NetBox source-of-truth and data-model documentation.
Verified: 2026-07-26
Official source: https://netboxlabs.com/docs/netbox/

07. Batfish Project - Batfish network configuration analysis
Role: Official network configuration-analysis and differential-verification project.
Verified: 2026-07-26
Official source: https://batfish.org/

08. OpenTelemetry - OpenTelemetry semantic conventions 1.43.0
Role: Current common telemetry naming and semantic conventions.
Verified: 2026-07-26
Official source: https://opentelemetry.io/docs/specs/semconv/

Authority application and refresh triggers

Authority class	Required library action
Protocols and specifications	Recheck interoperability, security considerations, defaults, and migration guidance when a referenced specification changes.
Security and assurance frameworks	Update control mappings, evidence expectations, risk language, and assessment procedures after material framework revision.
Languages, runtimes, and projects	Re-run examples, tests, deprecation checks, and operational recommendations against supported releases.
Benchmarks and evaluations	Re-execute claims when workloads, methods, models, compilers, drivers, hardware, or scoring rules change.
Operational evidence	Incorporate validated incidents, failure tests, reader corrections, and measured implementation outcomes into the next revision.

Limitations, freshness, and revision control

Boundary	Statement
Publication limitation	This guide synthesizes broad engineering practice. It does not replace product documentation, legal advice, safety certification, or environment-specific design review.
Evidence limitation	Not every pattern is benchmarked in every environment. Examples illustrate engineering methods and must be validated against actual versions, workloads, hardware, policy, and failure conditions.
Freshness limitation	Vendor behavior, software libraries, AI models, security threats, and emerging standards can change quickly. Volatile claims require source revalidation before operational adoption.
Adoption limitation	Frontier content is not an implicit adoption recommendation. Adoption requires defined value, qualification gates, controlled proof, rollback, ownership, and residual-risk acceptance.
Status boundary	Candidate Focused Field Guide describes the publication, not certification of a product, implementation, engineer, or organization.

Revision record

Version	Revision statement
1.0.0-candidate	Initial candidate living-guide edition. Published 2026-07-26; scheduled review 2026-10-26.

Search and relationship metadata

Dimension	Engineering position
Primary domain	MYTHOS-NET - Network Engineering & Architecture
Secondary domain	MYTHOS-DAT; MYTHOS-SWE; MYTHOS-SEC
Publication class	Field Guide / Canonical Living Overview Guide
Skill levels	L1 Foundational; L2 Practitioner; L3 Advanced; L4 Frontier
Tags	source of truth, intent, drift, reconciliation, NetBox, data model, inventory, lifecycle, automation, governance
Canonical route	/library/field-guides/mythos-fg-net-0017/

Living publication

Corrections, expanded laboratories, improved diagrams, authority changes, and new engineering evidence should revise this permanent publication rather than create disconnected replacements.