CANONICAL LIVING OVERVIEW GUIDE

Streaming Telemetry, gNMI, OpenTelemetry, and Network Observability

A complete network observability guide covering telemetry architecture, YANG-Push, gNMI, OpenConfig, metrics, logs, events, traces, OpenTelemetry, collectors, schemas, time, cardinality, storage, SLOs, anomaly detection, security, and troubleshooting.

PERMANENT ID
MYTHOS-FG-NET-0019

PUBLICATION CLASS
FIELD GUIDE

VERSION / STATUS
1.0.0-candidate / CANDIDATE

PERMANENT PUBLICATION IDENTITY

Streaming Telemetry, gNMI, OpenTelemetry, and Network Observability

Document ID
MYTHOS-FG-NET-0019

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 observability, SRE, automation, platform, security, and operations engineers.
Prerequisites	Network operations, APIs, metrics, and basic data-platform concepts.
Publisher / owner	Mythos Systems / Aaron Stovall
Public classification	Public technical guidance

Expected outcomes

Design telemetry around operational questions and service objectives rather than raw data volume.

Use gNMI, YANG-Push, OpenConfig, and secure streaming subscriptions correctly.

Normalize network metrics, events, logs, traces, resources, and topology into correlated evidence.

Control time, sampling, cardinality, gaps, retention, cost, and privacy.

Build dashboards, alerts, investigations, and automation that remain explainable and testable.

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

Observability questions and architecture

Start with decisions, users, services, and failure hypotheses.

Chapter operating position
Start with decisions, users, services, and failure hypotheses.

1.1 Monitoring versus observability

Monitoring versus observability must be treated as an engineering mechanism rather than a product checkbox. Use known checks and exploratory evidence together without claiming telemetry automatically explains causality. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 monitoring versus observability.
- 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 monitoring versus observability 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 monitoring versus observability.
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 Service and network objectives

Service and network objectives must be treated as an engineering mechanism rather than a product checkbox. Define availability, reachability, loss, latency, convergence, capacity, security, and change objectives. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 network 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.

Implementation sketch

```
subscription:
  target: leaf-17
  protocol: gnmi
  mode: stream
  paths:
    - /interfaces/interface/state/counters
    - /network-instances/network-instance/protocols/protocol/bgp/neighbors
  sampling:
    counters_ms: 10000
    state: on_change
  security:
    mtls: true
    role: telemetry-reader
  evidence:
    detect_gap: true
    retain_source_timestamp: true
```

Failure patterns

- Treating service and network 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 service and network 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.

1.3 Signals and evidence

Signals and evidence must be treated as an engineering mechanism rather than a product checkbox. Combine metrics, state, events, logs, packets, flows, traces, topology, configuration, and application results. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 signals and evidence.
- 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 signals and evidence 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 signals and evidence.
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 Collection architecture

Collection architecture must be treated as an engineering mechanism rather than a product checkbox. Design targets, collectors, brokers, processing, storage, query, visualization, alerting, automation, and archive. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 collection architecture.

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 collection architecture 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 collection architecture.
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
Define ten operational questions and map each to required signals, collection method, freshness, retention, and validation.

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 architecture to edge analytics, privacy-preserving telemetry, digital twins, and autonomous evidence systems.

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

gNMI, YANG-Push, and subscriptions

Collect modeled configuration and state with explicit paths, modes, time, and security.

Chapter operating position
Collect modeled configuration and state with explicit paths, modes, time, and security.

2.1 gNMI Get, Set, and Subscribe

gNMI Get, Set, and Subscribe must be treated as an engineering mechanism rather than a product checkbox. Use targets, paths, prefixes, origins, encodings, atomicity, timestamps, aliases, and errors. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 gnmi get, set, and subscribe.
- 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 gnmi get, set, and subscribe 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 gnmi get, set, and subscribe.
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 Subscription modes

Subscription modes must be treated as an engineering mechanism rather than a product checkbox. Choose once, poll, stream, sample, on-change, heartbeat, suppress redundant, and updates-only 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 target, model, subscription, collector, processing, storage, query, and operator action 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 subscription modes.
- 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

```
subscription:
  target: leaf-17
  protocol: gnmi
  mode: stream
  paths:
    - /interfaces/interface/state/counters
    - /network-instances/network-instance/protocols/protocol/bgp/neighbors
  sampling:
    counters_ms: 10000
    state: on_change
  security:
    mtls: true
    role: telemetry-reader
  evidence:
    detect_gap: true
    retain_source_timestamp: true
```

Failure patterns

- Treating subscription modes 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 subscription modes.
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 YANG-Push

YANG-Push must be treated as an engineering mechanism rather than a product checkbox. Use dynamic or configured subscriptions, datastore updates, dampening, periodic update, and receiver state. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 yang-push.
- 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 yang-push 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 yang-push.
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 Security and authorization

Security and authorization must be treated as an engineering mechanism rather than a product checkbox. Use TLS, target identity, client identity, access control, accounting, certificate lifecycle, and least privilege. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 authorization.

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 authorization 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 authorization.
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
Create sample and on-change subscriptions for interfaces and BGP, then test disconnect, replay gap, heartbeat, time skew, and authorization.

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-driven network APIs, signed telemetry, confidential collectors, and post-quantum management channels.

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

Models and semantic conventions

Create stable meaning across devices, vendors, collectors, and consumers.

Chapter operating position
Create stable meaning across devices, vendors, collectors, and consumers.

3.1 YANG paths and OpenConfig

YANG paths and OpenConfig must be treated as an engineering mechanism rather than a product checkbox. Use canonical models, keys, config and state containers, vendor deviations, versions, and origin. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 yang paths and openconfig.
- 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 yang paths and openconfig 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 yang paths and openconfig.
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 Resources and identity

Resources and identity must be treated as an engineering mechanism rather than a product checkbox. Represent device, component, interface, site, tenant, service, software, role, and stable identifiers. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 resources and 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.

Implementation sketch

```
subscription:
  target: leaf-17
  protocol: gnmi
  mode: stream
  paths:
    - /interfaces/interface/state/counters
    - /network-instances/network-instance/protocols/protocol/bgp/neighbors
  sampling:
    counters_ms: 10000
    state: on_change
  security:
    mtls: true
    role: telemetry-reader
  evidence:
    detect_gap: true
    retain_source_timestamp: true
```

Failure patterns

- Treating resources and 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 resources and 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.

3.3 Metric and event naming

Metric and event naming must be treated as an engineering mechanism rather than a product checkbox. Use namespaces, units, descriptions, types, temporality, monotonicity, and semantic consistency. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 metric and event naming.
- 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 metric and event naming 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 metric and event naming.
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 OpenTelemetry integration

OpenTelemetry integration must be treated as an engineering mechanism rather than a product checkbox. Map network resources and signals into traces, metrics, logs, context, attributes, and correlation while noting gaps. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 opentelemetry 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.

Failure patterns

Treating opentelemetry 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 opentelemetry 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.

Chapter 03 laboratory and review

Practical laboratory
Define a semantic profile for interfaces, routing neighbors, WLAN clients, and network changes across gNMI and OpenTelemetry.

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 standard network semantic conventions, knowledge graphs, and cross-domain service context.

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

Collection and pipeline engineering

Move high-rate data reliably without losing provenance, ordering, or control.

Chapter operating position
Move high-rate data reliably without losing provenance, ordering, or control.

4.1 Collectors and gateways

Collectors and gateways must be treated as an engineering mechanism rather than a product checkbox. Use direct, regional, edge, proxy, fan-in, load-balanced, and tenant-isolated patterns. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 collectors and gateways.
- 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 collectors and gateways 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 collectors and gateways.
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 Buffering and backpressure

Buffering and backpressure must be treated as an engineering mechanism rather than a product checkbox. Handle burst, disconnect, queue, retry, loss, sampling, prioritization, and stale 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 target, model, subscription, collector, processing, storage, query, and operator action 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 buffering and backpressure.
- 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

```
subscription:
  target: leaf-17
  protocol: gnmi
  mode: stream
  paths:
    - /interfaces/interface/state/counters
    - /network-instances/network-instance/protocols/protocol/bgp/neighbors
  sampling:
    counters_ms: 10000
    state: on_change
  security:
    mtls: true
    role: telemetry-reader
  evidence:
    detect_gap: true
    retain_source_timestamp: true
```

Failure patterns

- Treating buffering and backpressure 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 buffering and backpressure.
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 Processing and enrichment

Processing and enrichment must be treated as an engineering mechanism rather than a product checkbox. Normalize, validate, derive rates, join topology, attach ownership, classify events, redact, and route. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 processing and enrichment.
- 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 processing and enrichment 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 processing and enrichment.
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 Pipeline observability

Pipeline observability must be treated as an engineering mechanism rather than a product checkbox. Measure collection lag, missing targets, decode errors, dropped data, queue depth, storage failure, and cost. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 pipeline observability.

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 pipeline observability 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 pipeline observability.
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
Build a collector pipeline that survives target and backend outages, records gaps, resumes safely, and preserves original timestamps and identity.

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 stream processing at SmartNICs, eBPF collectors, edge AI summarization, and verifiable telemetry chains.

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

Time, sampling, cardinality, and storage

Control the properties that most often make telemetry misleading or unaffordable.

Chapter operating position
Control the properties that most often make telemetry misleading or unaffordable.

5.1 Time synchronization

Time synchronization must be treated as an engineering mechanism rather than a product checkbox. Use NTP, PTP where required, source timestamp, receive timestamp, clock quality, drift, leap behavior, and ordering. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 time synchronization.
- 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 time synchronization 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 time synchronization.
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 Sampling and aggregation

Sampling and aggregation must be treated as an engineering mechanism rather than a product checkbox. Choose periods, on-change, histograms, percentiles, rates, rollups, and retention by question and volatility. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 sampling and aggregation.
- 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

```
subscription:
  target: leaf-17
  protocol: gnmi
  mode: stream
  paths:
    - /interfaces/interface/state/counters
    - /network-instances/network-instance/protocols/protocol/bgp/neighbors
  sampling:
    counters_ms: 10000
    state: on_change
  security:
    mtls: true
    role: telemetry-reader
  evidence:
    detect_gap: true
    retain_source_timestamp: true
```

Failure patterns

- Treating sampling and aggregation 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 sampling and aggregation.
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 Cardinality

Cardinality must be treated as an engineering mechanism rather than a product checkbox. Control labels such as client, prefix, flow, policy, path, tenant, and session without destroying diagnostic value. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 cardinality.
- 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 cardinality 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 cardinality.
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 Storage and retention

Storage and retention must be treated as an engineering mechanism rather than a product checkbox. Choose time-series, logs, object, search, columnar, graph, and archive tiers with integrity and deletion policy. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 storage and retention.

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 storage and retention 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 storage and retention.
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
Estimate data rate and cardinality for interface, BGP, WLAN, flow, and client telemetry. Design sampling, rollups, and retention within a cost budget.

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 adaptive sampling, sketch algorithms, semantic compression, and energy-aware telemetry.

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

Dashboards, alerts, and SLOs

Turn signals into operator decisions without hiding uncertainty or creating alarm fatigue.

Chapter operating position
Turn signals into operator decisions without hiding uncertainty or creating alarm fatigue.

6.1 Service-oriented views

Service-oriented views must be treated as an engineering mechanism rather than a product checkbox. Show user and application path, dependencies, policy, topology, impact, trend, and current incidents. 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 target, model, subscription, collector, processing, storage, query, and operator action 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-oriented views.
- 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-oriented views 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-oriented views.
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 Alert design

Alert design must be treated as an engineering mechanism rather than a product checkbox. Use symptoms, causes, correlation, severity, duration, confidence, suppression, maintenance, and routing. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 alert 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.

Implementation sketch

```
subscription:
  target: leaf-17
  protocol: gnmi
  mode: stream
  paths:
    - /interfaces/interface/state/counters
    - /network-instances/network-instance/protocols/protocol/bgp/neighbors
  sampling:
    counters_ms: 10000
    state: on_change
  security:
    mtls: true
    role: telemetry-reader
  evidence:
    detect_gap: true
    retain_source_timestamp: true
```

Failure patterns

- Treating alert 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 alert 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.

6.3 Network SLOs

Network SLOs must be treated as an engineering mechanism rather than a product checkbox. Define indicators for reachability, latency, loss, convergence, change success, telemetry health, and client experience. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 network slos.
- 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 network slos 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 network slos.
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 Error budgets and action

Error budgets and action must be treated as an engineering mechanism rather than a product checkbox. Use service risk to prioritize reliability, capacity, automation, and release decisions. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 error budgets and action.

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 error budgets and action 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 error budgets and action.
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

Create an SLO and alert set for branch reachability and wireless authentication. Test normal, noisy, partial, and real failure conditions.

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 causal graphs, adaptive thresholds, natural-language investigation, and human-centered operational UX.

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

Troubleshooting and causal evidence

Correlate signals across endpoints, paths, control planes, changes, and services.

Chapter operating position
Correlate signals across endpoints, paths, control planes, changes, and services.

7.1 Golden flow and topology context

Golden flow and topology context must be treated as an engineering mechanism rather than a product checkbox. Trace source, destination, identity, path, policy, route, tunnel, queue, and service 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 target, model, subscription, collector, processing, storage, query, and operator action 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 golden flow and topology context.
- 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 golden flow and topology context 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 golden flow and topology context.
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 Change correlation

Change correlation must be treated as an engineering mechanism rather than a product checkbox. Link configuration, software, policy, controller, topology, and maintenance events to observed impact. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 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

```
subscription:
  target: leaf-17
  protocol: gnmi
  mode: stream
  paths:
    - /interfaces/interface/state/counters
    - /network-instances/network-instance/protocols/protocol/bgp/neighbors
  sampling:
    counters_ms: 10000
    state: on_change
  security:
    mtls: true
    role: telemetry-reader
  evidence:
    detect_gap: true
    retain_source_timestamp: true
```

Failure patterns

- Treating change 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 change 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.

7.3 Packets and flows

Packets and flows must be treated as an engineering mechanism rather than a product checkbox. Use packet capture, flow records, synthetic tests, and endpoint telemetry when aggregate metrics are insufficient. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 packets and flows.
- 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 packets and flows 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 packets and flows.
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 Evidence bundle

Evidence bundle must be treated as an engineering mechanism rather than a product checkbox. Preserve raw data, queries, dashboards, timestamps, topology, changes, interpretation, confidence, and residual unknowns. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 evidence bundle.

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 evidence bundle 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 evidence bundle.
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

Investigate a synthetic outage using only telemetry at first, then add packet and endpoint evidence. Record where each signal changed the conclusion.

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

Develop evidence-citing AI investigation agents with deterministic queries, source links, confidence, and human escalation.

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, security, and automation

Protect sensitive telemetry and use it safely for closed-loop operations.

Chapter operating position
Protect sensitive telemetry and use it safely for closed-loop operations.

8.1 Access and privacy

Access and privacy must be treated as an engineering mechanism rather than a product checkbox. Control topology, client, identity, location, packet, flow, configuration, security, and tenant 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 target, model, subscription, collector, processing, storage, query, and operator action 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 access and privacy.
- 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 access and privacy 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 access and privacy.
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 Integrity and provenance

Integrity and provenance must be treated as an engineering mechanism rather than a product checkbox. Track target identity, collector, transformation, schema, time, loss, signature, storage, and query. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 integrity 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.

Implementation sketch

```
subscription:
  target: leaf-17
  protocol: gnmi
  mode: stream
  paths:
    - /interfaces/interface/state/counters
    - /network-instances/network-instance/protocols/protocol/bgp/neighbors
  sampling:
    counters_ms: 10000
    state: on_change
  security:
    mtls: true
    role: telemetry-reader
  evidence:
    detect_gap: true
    retain_source_timestamp: true
```

Failure patterns

- Treating integrity 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 integrity 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.

8.3 Automation triggers

Automation triggers must be treated as an engineering mechanism rather than a product checkbox. Use thresholds, policies, approvals, simulation, canaries, action limits, rollback, and 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 target, model, subscription, collector, processing, storage, query, and operator action 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 automation triggers.
- 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 automation triggers 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 automation triggers.
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 Platform lifecycle

Platform lifecycle must be treated as an engineering mechanism rather than a product checkbox. Manage schemas, collectors, certificates, agents, backends, dashboards, retention, cost, and migrations. 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 target, model, subscription, collector, processing, storage, query, and operator action 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 platform 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 platform 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 platform 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.

Chapter 08 laboratory and review

Practical laboratory

Design a telemetry-triggered remediation that can drain a link or roll back a change only after validation and approval 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

Explore autonomous network operations, confidential analytics, cross-domain observability, and machine-verifiable assurance cases.

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. 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/

02. OpenConfig - gNMI Authentication and Encryption
Role: Secure channel, authorization, and accounting guidance for gNMI.
Verified: 2026-07-26
Official source: https://www.openconfig.net/docs/gnmi/gnmi-authentication/

03. IETF / RFC Editor - RFC 8639 - Subscription to YANG Notifications
Role: Dynamic and configured telemetry subscriptions.
Verified: 2026-07-26
Official source: https://www.rfc-editor.org/info/rfc8639

04. IETF / RFC Editor - RFC 8641 - Subscription to YANG Datastores
Role: YANG-Push datastore update subscriptions.
Verified: 2026-07-26
Official source: https://www.rfc-editor.org/info/rfc8641

05. 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

06. 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/

07. 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/

08. IETF / RFC Editor - RFC 8341 - Network Configuration Access Control Model
Role: NETCONF and RESTCONF access-control authority.
Verified: 2026-07-26
Official source: https://www.rfc-editor.org/info/rfc8341

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	streaming telemetry, gNMI, OpenConfig, OpenTelemetry, YANG-Push, observability, SLO, metrics, events, automation
Canonical route	/library/field-guides/mythos-fg-net-0019/

Living publication

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