



CANONICAL LIVING OVERVIEW GUIDE

API Design, Typed Contracts, Schemas, and Compatibility

An API-engineering guide covering HTTP and message semantics, resource and operation design, typed contracts, OpenAPI 3.2, AsyncAPI 3.1, JSON Schema, Protocol Buffers, GraphQL, errors, identity, pagination, idempotency, compatibility, versioning, testing, gateways, observability, governance, and AI-consumable tool interfaces.

PERMANENT PUBLICATION IDENTITY

API Design, Typed Contracts, Schemas, and Compatibility

Document ID
MYTHOS-FG-SWE-0003

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

Reader orientation

Dimension	Engineering position
Primary domain	MYTHOS-SWE - Software Engineering & Design
Secondary domain	MYTHOS-DAT; MYTHOS-CLD; MYTHOS-SEC
Audience	API architects, software and platform engineers, integration teams, data engineers, security engineers, technical writers, SDK teams, and agent-tool developers.
Prerequisites	HTTP, distributed applications, schemas, serialization, authentication, testing, and software architecture fundamentals.
Publisher / owner	Mythos Systems / Aaron Stovall
Public classification	Public technical guidance

Expected outcomes

- Design APIs around domain capabilities, resource state, and explicit contracts rather than framework defaults.
- Use typed schemas to define inputs, outputs, errors, events, metadata, and compatibility rules.
- Apply HTTP semantics, including the 2026 QUERY method, safely and intentionally.
- Evolve synchronous and asynchronous contracts without silently breaking consumers.
- Verify APIs through contract tests, generated examples, negative cases, telemetry, and consumer evidence.

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

How to use this guide	3
Contents	3
Chapter 01 - API purpose, consumers, and boundary design	5
Chapter 01 laboratory and review	10
Chapter 02 - HTTP resources, operations, and semantics	11
Chapter 02 laboratory and review	16
Chapter 03 - Typed contracts and schema engineering	17
Chapter 03 laboratory and review	22
Chapter 04 - OpenAPI, AsyncAPI, and contract-first workflows	23
Chapter 04 laboratory and review	28
Chapter 05 - Errors, retries, idempotency, and workflow semantics	29
Chapter 05 laboratory and review	34

Chapter 06 - Security, privacy, and abuse resistance	35
Chapter 06 laboratory and review	40
Chapter 07 - Compatibility, versioning, and migration	41
Chapter 07 laboratory and review	46
Chapter 08 - Testing, observability, governance, and lifecycle	47
Chapter 08 laboratory and review	52
Integrated learning roadmap	53
Source authority register	54
Limitations, freshness, and revision control	55

CHAPTER 01

API purpose, consumers, and boundary design

Define the capability, ownership, audience, and trust boundary before choosing a protocol.

Chapter operating position

Define the capability, ownership, audience, and trust boundary before choosing a protocol.

1.1 Capability and domain alignment

Capability and domain alignment must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Identify the business or platform capability, owning context, authoritative state, invariants, and permitted operations. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for capability and domain alignment.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating capability and domain alignment as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for capability and domain alignment.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

1.2 Consumer and use-case analysis

Consumer and use-case analysis must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Map human, web, mobile, service, batch, partner, device, agent, and administrative consumers with distinct needs. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for consumer and use-case analysis.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Implementation sketch

```
openapi: 3.2.0
paths:
  /contacts:
    query:
      operationId: searchContacts
      requestBody:
        required: true
        content:
          application/json:
            schema: {$ref: '#/components/schemas/ContactQuery'}
      responses:
        '200': {description: Query result}
        '400': {$ref: '#/components/responses/Problem'}
    compatibility:
      unknown_response_fields: tolerated
      removed_enum_values: prohibited
```

Failure patterns

Treating consumer and use-case analysis as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for consumer and use-case analysis.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.



1.3 Trust, exposure, and service boundary

Trust, exposure, and service boundary must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define internal, partner, public, tenant, management, data, network, identity, privacy, and regulatory boundaries. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for trust, exposure, and service boundary.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating trust, exposure, and service boundary as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for trust, exposure, and service boundary.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

1.4 Protocol and interaction selection

Protocol and interaction selection must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Choose request-response, query, streaming, event, command, subscription, file, or batch according to semantics and failure behavior. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for protocol and interaction selection.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating protocol and interaction selection as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for protocol and interaction selection.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

Chapter 01 laboratory and review

Practical laboratory

Create an API boundary and consumer map for one domain capability and justify each chosen interaction style.

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 boundary design to agent tools, machine-readable capabilities, confidential APIs, and post-quantum service identity.

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

HTTP resources, operations, and semantics

Use protocol semantics accurately so clients, intermediaries, and operations can reason about behavior.

Chapter operating position

Use protocol semantics accurately so clients, intermediaries, and operations can reason about behavior.



2.1 Resources, representations, and identifiers

Resources, representations, and identifiers must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Design stable URIs, canonical identity, representations, content negotiation, links, and state transitions. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for resources, representations, and identifiers.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating resources, representations, and identifiers as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for resources, representations, and identifiers.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

2.2 Methods and safety properties

Methods and safety properties must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use GET, HEAD, POST, PUT, PATCH, DELETE, QUERY, and OPTIONS according to safety, idempotency, cacheability, and payload semantics. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for methods and safety properties.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Implementation sketch

```
openapi: 3.2.0
paths:
  /contacts:
    query:
      operationId: searchContacts
      requestBody:
        required: true
        content:
          application/json:
            schema: {$ref: '#/components/schemas/ContactQuery'}
      responses:
        '200': {description: Query result}
        '400': {$ref: '#/components/responses/Problem'}
compatibility:
  unknown_response_fields: tolerated
  removed_enum_values: prohibited
```

Failure patterns

Treating methods and safety properties as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for methods and safety properties.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.



2.3 Status, headers, and conditional requests

Status, headers, and conditional requests must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use status codes, validators, preconditions, ranges, locations, caching, retries, and concurrency control consistently. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for status, headers, and conditional requests.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating status, headers, and conditional requests as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for status, headers, and conditional requests.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.



2.4 Queries, commands, and asynchronous operations

Queries, commands, and asynchronous operations must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Model safe body-bearing queries, long-running jobs, operation resources, callbacks, polling, cancellation, and completion evidence. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for queries, commands, and asynchronous operations.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating queries, commands, and asynchronous operations as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for queries, commands, and asynchronous operations.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

Chapter 02 laboratory and review

Practical laboratory

Design and test a resource API using conditional updates, idempotent creation, RFC 9457 errors, and a safe QUERY operation.

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 HTTP message signatures, content-addressed APIs, capability links, and verifiable operation receipts.

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

Typed contracts and schema engineering

Make data meaning, validation, evolution, and ownership explicit.

Chapter operating position

Make data meaning, validation, evolution, and ownership explicit.



3.1 Schema vocabulary and type systems

Schema vocabulary and type systems must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define scalar, object, collection, union, enumeration, optionality, defaults, constraints, formats, and semantic annotations. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for schema vocabulary and type systems.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating schema vocabulary and type systems as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for schema vocabulary and type systems.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

3.2 JSON Schema and validation

JSON Schema and validation must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use dialect declarations, references, vocabularies, applicators, conditional rules, unevaluated properties, and output formats. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for json schema and validation.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Implementation sketch

```
openapi: 3.2.0
paths:
  /contacts:
    query:
      operationId: searchContacts
      requestBody:
        required: true
        content:
          application/json:
            schema: {$ref: '#/components/schemas/ContactQuery'}
      responses:
        '200': {description: Query result}
        '400': {$ref: '#/components/responses/Problem'}
compatibility:
  unknown_response_fields: tolerated
  removed_enum_values: prohibited
```

Failure patterns

Treating json schema and validation as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for json schema and validation.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.



3.3 Protocol Buffers and binary contracts

Protocol Buffers and binary contracts must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use field numbers, wire types, oneof, presence, reserved fields, packages, services, unknown fields, and code generation. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for protocol buffers and binary contracts.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating protocol buffers and binary contracts as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for protocol buffers and binary contracts.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

3.4 GraphQL and query contracts

GraphQL and query contracts must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define schema, types, fields, arguments, nullability, interfaces, unions, execution, errors, authorization, and cost limits. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for graphql and query contracts.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating graphql and query contracts as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for graphql and query contracts.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

Chapter 03 laboratory and review

Practical laboratory

Represent one domain model in JSON Schema, Protocol Buffers, and GraphQL, then compare compatibility, validation, and consumer ergonomics.

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 semantic schemas, type-level policy, proof-carrying data, and generated contracts for multimodal or agent interactions.

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

OpenAPI, AsyncAPI, and contract-first workflows

Use standard descriptions as governed sources for documentation, code, testing, and change analysis.

Chapter operating position

Use standard descriptions as governed sources for documentation, code, testing, and change analysis.

4.1 OpenAPI 3.2 documents

OpenAPI 3.2 documents must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define paths, QUERY operations, operations, parameters, query strings, request bodies, responses, callbacks, webhooks, security, and reusable components. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for openapi 3.2 documents.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating openapi 3.2 documents as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for openapi 3.2 documents.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

4.2 AsyncAPI 3.1 documents

AsyncAPI 3.1 documents must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define channels, operations, messages, correlation, bindings, servers, security, traits, and message-driven interaction. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for asyncapi 3.1 documents.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Implementation sketch

```
openapi: 3.2.0
paths:
  /contacts:
    query:
      operationId: searchContacts
      requestBody:
        required: true
        content:
          application/json:
            schema: {$ref: '#/components/schemas/ContactQuery'}
      responses:
        '200': {description: Query result}
        '400': {$ref: '#/components/responses/Problem'}
compatibility:
  unknown_response_fields: tolerated
  removed_enum_values: prohibited
```

Failure patterns

Treating asyncapi 3.1 documents as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for asyncapi 3.1 documents.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.



4.3 Contract repository and generation

Contract repository and generation must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Version descriptions, validate, lint, bundle, generate SDKs, mocks, documentation, tests, policies, and catalogs. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for contract repository and generation.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating contract repository and generation as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for contract repository and generation.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.



4.4 Design review and implementation conformance

Design review and implementation conformance must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Compare contract, code, gateway, runtime traffic, examples, errors, and consumer behavior for drift. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for design review and implementation conformance.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating design review and implementation conformance as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for design review and implementation conformance.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

Chapter 04 laboratory and review

Practical laboratory

Create OpenAPI and AsyncAPI descriptions for a capability with request-response and event behavior, then generate tests and detect drift.

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 unified API intermediate representations and agent-readable capability manifests with signed provenance.

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

Errors, retries, idempotency, and workflow semantics

Design failure responses that help consumers act correctly without leaking unsafe detail.

Chapter operating position

Design failure responses that help consumers act correctly without leaking unsafe detail.

5.1 Problem details and error taxonomy

Problem details and error taxonomy must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use stable problem types, status, title, instance, extensions, causes, correlation, and remediation guidance. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for problem details and error taxonomy.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating problem details and error taxonomy as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for problem details and error taxonomy.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

5.2 Retry and backoff semantics

Retry and backoff semantics must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Identify transient, permanent, throttled, conflict, validation, dependency, timeout, and unknown outcomes with retry safety. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for retry and backoff semantics.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Implementation sketch

```
openapi: 3.2.0
paths:
  /contacts:
    query:
      operationId: searchContacts
      requestBody:
        required: true
        content:
          application/json:
            schema: {$ref: '#/components/schemas/ContactQuery'}
      responses:
        '200': {description: Query result}
        '400': {$ref: '#/components/responses/Problem'}
compatibility:
  unknown_response_fields: tolerated
  removed_enum_values: prohibited
```

Failure patterns

Treating retry and backoff semantics as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for retry and backoff semantics.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.



5.3 Idempotency and duplicate control

Idempotency and duplicate control must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use idempotency keys, operation identity, deduplication windows, stored results, side-effect detection, and replay handling. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for idempotency and duplicate control.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating idempotency and duplicate control as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for idempotency and duplicate control.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

5.4 Long-running and compensating workflows

Long-running and compensating workflows must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Expose state, progress, cancellation, partial completion, compensation, recovery, and terminal outcomes. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for long-running and compensating workflows.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating long-running and compensating workflows as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for long-running and compensating workflows.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

Chapter 05 laboratory and review

Practical laboratory

Implement a mock operation that experiences timeout after side effects, duplicate requests, retry, cancellation, and compensation; prove client-safe behavior.

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 verifiable workflow receipts, exactly-once effect protocols, and agent-planned compensation under policy.

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

Security, privacy, and abuse resistance

Protect API authority, data, capacity, and administrative surfaces end to end.

Chapter operating position

Protect API authority, data, capacity, and administrative surfaces end to end.

6.1 Authentication and service identity

Authentication and service identity must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use OAuth, OIDC, mTLS, workload identity, client credentials, tokens, certificates, key rotation, and trust domains. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for authentication and service identity.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating authentication and service identity as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for authentication and service identity.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

6.2 Authorization and tenancy

Authorization and tenancy must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Enforce operation, object, field, tenant, relationship, purpose, environment, and transaction constraints server side. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for authorization and tenancy.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Implementation sketch

```
openapi: 3.2.0
paths:
  /contacts:
    query:
      operationId: searchContacts
      requestBody:
        required: true
        content:
          application/json:
            schema: {$ref: '#/components/schemas/ContactQuery'}
      responses:
        '200': {description: Query result}
        '400': {$ref: '#/components/responses/Problem'}
compatibility:
  unknown_response_fields: tolerated
  removed_enum_values: prohibited
```

Failure patterns

Treating authorization and tenancy as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for authorization and tenancy.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.



6.3 Input, output, and data protection

Input, output, and data protection must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Validate schemas, sizes, encodings, files, URLs, callbacks, secrets, privacy, classification, retention, and egress. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for input, output, and data protection.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating input, output, and data protection as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for input, output, and data protection.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

6.4 Abuse and resource protection

Abuse and resource protection must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Apply quotas, rate limits, concurrency, query complexity, pagination, payload limits, cost controls, bot defenses, and monitoring. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for abuse and resource protection.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating abuse and resource protection as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for abuse and resource protection.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

Chapter 06 laboratory and review

Practical laboratory

Threat-model and test a public API for object authorization, unsafe consumption, SSRF, replay, rate abuse, data leakage, and administrative exposure.

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 controls to agent delegation, tool-call authorization, signed requests, confidential processing, and privacy-enhancing queries.

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

Compatibility, versioning, and migration

Evolve contracts according to observable consumer behavior and explicit compatibility policy.

Chapter operating position

Evolve contracts according to observable consumer behavior and explicit compatibility policy.

7.1 Compatibility dimensions

Compatibility dimensions must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Assess source, binary, wire, schema, semantic, behavioral, performance, operational, and security compatibility. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for compatibility dimensions.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating compatibility dimensions as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for compatibility dimensions.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

7.2 Additive and breaking changes

Additive and breaking changes must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Classify fields, values, validation, defaults, ordering, errors, methods, auth, timing, quotas, and data meaning changes. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for additive and breaking changes.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Implementation sketch

```
openapi: 3.2.0
paths:
  /contacts:
    query:
      operationId: searchContacts
      requestBody:
        required: true
        content:
          application/json:
            schema: {$ref: '#/components/schemas/ContactQuery'}
      responses:
        '200': {description: Query result}
        '400': {$ref: '#/components/responses/Problem'}
compatibility:
  unknown_response_fields: tolerated
  removed_enum_values: prohibited
```

Failure patterns

Treating additive and breaking changes as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for additive and breaking changes.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

7.3 Version strategy and deprecation

Version strategy and deprecation must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Choose media, path, header, schema, message, capability, or continuous evolution with notices, telemetry, deadlines, and support. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for version strategy and deprecation.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating version strategy and deprecation as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for version strategy and deprecation.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

7.4 Consumer-driven migration

Consumer-driven migration must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Inventory consumers, test usage, provide adapters, shadow traffic, dual run, canaries, rollback, and retirement evidence. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for consumer-driven migration.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating consumer-driven migration as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for consumer-driven migration.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

Chapter 07 laboratory and review

Practical laboratory

Process a contract change through compatibility analysis, consumer tests, dual-version rollout, deprecation, and final retirement.

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 automated semantic compatibility analysis, traffic-derived consumer models, and formally checked migrations.

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

Testing, observability, governance, and lifecycle

Operate APIs as products and critical distributed-system boundaries.

Chapter operating position

Operate APIs as products and critical distributed-system boundaries.



8.1 Contract and conformance testing

Contract and conformance testing must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use schema, examples, negative, fuzz, property, consumer, provider, gateway, protocol, and compatibility tests. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for contract and conformance testing.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating contract and conformance testing as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for contract and conformance testing.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

8.2 API telemetry and SLOs

API telemetry and SLOs must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Measure requests, operations, errors, problem types, latency, retries, saturation, quota, auth, dependencies, and consumer impact. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for api telemetry and slo.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Implementation sketch

```
openapi: 3.2.0
paths:
  /contacts:
    query:
      operationId: searchContacts
      requestBody:
        required: true
        content:
          application/json:
            schema: {$ref: '#/components/schemas/ContactQuery'}
      responses:
        '200': {description: Query result}
        '400': {$ref: '#/components/responses/Problem'}
compatibility:
  unknown_response_fields: tolerated
  removed_enum_values: prohibited
```

Failure patterns

Treating api telemetry and slos as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for api telemetry and slos.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

8.3 Catalog, ownership, and governance

Catalog, ownership, and governance must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Maintain identity, owner, domain, lifecycle, documentation, risk, dependencies, versions, consumers, and review status. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for catalog, ownership, and governance.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating catalog, ownership, and governance as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for catalog, ownership, and governance.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.



8.4 SDKs, documentation, and developer experience

SDKs, documentation, and developer experience must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Provide generated and curated examples, authentication guidance, errors, limits, changelogs, migration, support, and sandboxes. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the consumer, operation, contract, identity, request, validation, state change, response, event, and compatibility outcome chain. A diagram, passing test, successful request, exploit result, benchmark, or green dashboard is not sufficient proof by itself. Intended behavior must be reconciled with implemented state, negative and boundary tests, degraded conditions, raw evidence, recovery, and the resulting user, service, or risk outcome.

Implementation begins with a minimal known-good baseline and explicit invariants. Introduce one variable or dependency at a time, preserve before-state evidence, test expected and prohibited behavior, inject realistic failure, and retain a reversible path. Automation and AI assistance must stop safely when authority, freshness, identity, state, or expected outcomes are ambiguous.

Troubleshooting and assurance should locate the first divergence from the expected contract or state machine. Account for time, concurrency, caching, eventual consistency, data and schema drift, partial failure, environment differences, hidden coupling, resource saturation, and missing telemetry. Completion requires reproducible evidence, a causal explanation, validated restoration, residual limitations, and prevention or monitoring actions.

Engineering practices

Define the objective, owner, authority, scope, assumptions, dependencies, and acceptance criteria for sdks, documentation, and developer experience.

Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.

Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.

Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.

Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.

Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

Treating sdks, documentation, and developer experience as a document or tool activity disconnected from actual system behavior.

Relying on intended design or configuration without proving implementation and runtime outcomes.

Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.

Changing several variables before preserving the original state and stating a falsifiable hypothesis.

Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.

Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

Method	Expected evidence
Examine	Requirements, architecture, source, contracts, data, versions, ownership, and assumptions for sdks, documentation, and developer experience.
Observe	Runtime state, control flow, telemetry, artifacts, resource behavior, errors, test results, and user-visible outcomes.
Test	Expected, prohibited, boundary, malformed, concurrent, degraded, failed, recovery, and rollback cases.
Measure	Correctness, latency, throughput, resource use, reliability, coverage, risk, maintainability, and operational impact.
Reconcile	Intended requirements and design against code, configuration, runtime behavior, tests, evidence, and residual limitations.

Chapter 08 laboratory and review

Practical laboratory

Build a contract pipeline that lints, tests, generates a client, runs compatibility checks, deploys a mock, and publishes catalog metadata.

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 API digital twins and governed agent-tool registries that continuously validate contracts against runtime behavior.

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. OpenAPI Initiative - OpenAPI Specification 3.2.0

Role: Current language-agnostic description standard for HTTP APIs.

Verified: 2026-07-26

Official source: <https://spec.openapis.org/oas/v3.2.0.html>

02. AsyncAPI Initiative - AsyncAPI Specification 3.1.0

Role: Current machine-readable description standard for message-driven APIs.

Verified: 2026-07-26

Official source: <https://www.asyncapi.com/docs/reference/specification/latest>

03. JSON Schema - JSON Schema Draft 2020-12

Role: Current released JSON Schema dialect and validation vocabulary.

Verified: 2026-07-26

Official source: <https://json-schema.org/specification>

04. IETF / RFC Editor - RFC 9110 - HTTP Semantics

Role: HTTP method, status, representation, caching, and semantic requirements.

Verified: 2026-07-26

Official source: <https://www.rfc-editor.org/info/rfc9110>

05. IETF / RFC Editor - RFC 9457 - Problem Details for HTTP APIs

Role: Machine-readable HTTP API error representation.

Verified: 2026-07-26

Official source: <https://www.rfc-editor.org/info/rfc9457>

06. IETF / RFC Editor - RFC 10008 - The HTTP QUERY Method

Role: Safe and idempotent request method carrying query content, published in 2026.

Verified: 2026-07-26

Official source: <https://www.rfc-editor.org/info/rfc10008>

07. Protocol Buffers - Protocol Buffers Language Guide

Role: Typed message schemas, field evolution, and binary compatibility.

Verified: 2026-07-26

Official source: <https://protobuf.dev/programming-guides/proto3/>

08. GraphQL Foundation - GraphQL Specification

Role: Typed query and execution semantics for GraphQL APIs.

Verified: 2026-07-26

Official source: <https://spec.graphql.org/>

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-SWE - Software Engineering & Design
Secondary domain	MYTHOS-DAT; MYTHOS-CLD; MYTHOS-SEC
Publication class	Field Guide / Canonical Living Overview Guide
Skill levels	L1 Foundational; L2 Practitioner; L3 Advanced; L4 Frontier
Tags	API design, OpenAPI 3.2, AsyncAPI 3.1, JSON Schema, Protocol Buffers, GraphQL, typed contracts, compatibility, versioning, HTTP QUERY
Canonical route	/library/field-guides/mythos-fg-swe-0003/

Living publication

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