



CANONICAL LIVING OVERVIEW GUIDE

Software Requirements, Specifications, and Acceptance Criteria

A requirements-engineering guide covering stakeholder needs, system and software requirements, product discovery, domain language, functional and quality requirements, constraints, interfaces, use cases, scenarios, normative specifications, acceptance criteria, traceability, change control, validation, and

PERMANENT PUBLICATION IDENTITY

Software Requirements, Specifications, and Acceptance Criteria

Document ID
MYTHOS-FG-SWE-0001

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-SEC; MYTHOS-UX
Audience	Product managers, business analysts, software and systems engineers, architects, test and quality engineers, security engineers, technical writers, and AI-assisted development teams.
Prerequisites	Basic software lifecycle, product, architecture, testing, and stakeholder-collaboration concepts.
Publisher / owner	Mythos Systems / Aaron Stovall
Public classification	Public technical guidance

Expected outcomes

- Transform ambiguous goals into bounded, testable, traceable software requirements.
- Separate stakeholder needs, system requirements, software requirements, constraints, assumptions, and design decisions.
- Write normative specifications and acceptance criteria that human and AI implementers interpret consistently.
- Validate requirements for correctness, completeness, feasibility, consistency, priority, risk, and verifiability.
- Maintain lineage from business objective through architecture, implementation, tests, release evidence, and change.

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 - Requirements engineering and decision context	5
Chapter 01 laboratory and review	10
Chapter 02 - Elicitation, discovery, and domain understanding	11
Chapter 02 laboratory and review	16
Chapter 03 - Requirement types and quality attributes	17
Chapter 03 laboratory and review	22
Chapter 04 - Specification forms and normative language	23
Chapter 04 laboratory and review	28
Chapter 05 - Acceptance criteria and executable examples	29
Chapter 05 laboratory and review	34

Chapter 06 - Validation, review, and conflict resolution	35
Chapter 06 laboratory and review	40
Chapter 07 - Traceability, change, and requirements lifecycle	41
Chapter 07 laboratory and review	46
Chapter 08 - Specification-driven and AI-assisted development	47
Chapter 08 laboratory and review	52
Integrated learning roadmap	53
Source authority register	54
Limitations, freshness, and revision control	55

CHAPTER 01

Requirements engineering and decision context

Establish why the software exists, which decisions it supports, and who owns acceptance.

Chapter operating position

Establish why the software exists, which decisions it supports, and who owns acceptance.



1.1 Mission, outcomes, and stakeholder needs

Mission, outcomes, and stakeholder needs must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define the problem, desired outcome, affected people, business capability, value, harm, success measures, and non-goals. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 mission, outcomes, and stakeholder needs.

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 mission, outcomes, and stakeholder needs 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 mission, outcomes, and stakeholder needs.
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 Stakeholders, users, and decision rights

Stakeholders, users, and decision rights must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Identify sponsors, product owners, users, operators, security, privacy, legal, support, vendors, and acceptance authority. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 stakeholders, users, and decision rights.

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

```
requirement:
  id: REQ-PAY-042
  statement: >
    The service MUST reject a payment approval when the request tenant
    differs from the tenant owning the payment object.
  rationale: prevent cross-tenant authorization failure
  acceptance:
    - same-tenant-authorized-user: allow
    - different-tenant-user: deny-and-audit
    - missing-tenant-context: deny
  traces_to: [ARCH-AUTH-007, TEST-AUTH-042]
```

Failure patterns

Treating stakeholders, users, and decision rights 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 stakeholders, users, and decision rights.
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 System boundary and environment

System boundary and environment must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define the product, adjacent systems, people, devices, networks, data, third parties, deployment environments, and lifecycle. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 system boundary and environment.

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 system boundary and environment 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 system boundary and environment.
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 Assumptions, constraints, and dependencies

Assumptions, constraints, and dependencies must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Record legal, technical, organizational, schedule, budget, platform, supplier, safety, and interoperability 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 assumptions, constraints, and dependencies.

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 assumptions, constraints, and dependencies 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 assumptions, constraints, and dependencies.
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 a decision-context and stakeholder map for a representative application, then identify conflicting outcomes and unresolved authority.

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 requirements engineering to autonomous agents, cyber-physical systems, sovereign deployments, and model-governed software.

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

Elicitation, discovery, and domain understanding

Gather evidence from work, systems, users, and constraints rather than relying on requested features alone.

Chapter operating position

Gather evidence from work, systems, users, and constraints rather than relying on requested features alone.



2.1 Interviews, workshops, and observation

Interviews, workshops, and observation must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use structured questions, event walkthroughs, task observation, contextual inquiry, and disagreement capture. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 interviews, workshops, and observation.

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 interviews, workshops, and observation 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 interviews, workshops, and observation.
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 Existing-system and data analysis

Existing-system and data analysis must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Inspect workflows, interfaces, logs, incidents, defects, reports, policies, schemas, metrics, and shadow processes. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 existing-system and data 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

```
requirement:
  id: REQ-PAY-042
  statement: >
    The service MUST reject a payment approval when the request tenant
    differs from the tenant owning the payment object.
  rationale: prevent cross-tenant authorization failure
  acceptance:
    - same-tenant-authorized-user: allow
    - different-tenant-user: deny-and-audit
    - missing-tenant-context: deny
  traces_to: [ARCH-AUTH-007, TEST-AUTH-042]
```

Failure patterns

Treating existing-system and data 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 existing-system and data 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.



2.3 Domain language and conceptual models

Domain language and conceptual models must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define terms, entities, states, events, rules, invariants, calculations, and ownership using a shared ubiquitous language. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 domain language and conceptual models.

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 domain language and conceptual models 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 domain language and conceptual models.
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 Prototypes, experiments, and discovery tests

Prototypes, experiments, and discovery tests must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use mockups, spikes, simulations, data probes, and technical experiments to reduce uncertainty before commitment. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 prototypes, experiments, and discovery tests.

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 prototypes, experiments, and discovery tests 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 prototypes, experiments, and discovery tests.
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

Run a requirements workshop from a synthetic case, produce a glossary and domain model, and resolve at least three ambiguous terms.

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 transcript-grounded discovery agents, process mining, executable domain models, and privacy-preserving user research.

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

Requirement types and quality attributes

Express behavior, quality, constraints, and operational needs at the correct level.

Chapter operating position

Express behavior, quality, constraints, and operational needs at the correct level.



3.1 Functional and behavioral requirements

Functional and behavioral requirements must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Specify triggers, actors, preconditions, actions, state transitions, outputs, exceptions, and business rules. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 functional and behavioral requirements.

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 functional and behavioral requirements 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 functional and behavioral requirements.
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 Quality requirements

Quality requirements must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define performance, reliability, security, usability, accessibility, compatibility, maintainability, portability, and safety with measures. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 quality requirements.

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

```
requirement:
  id: REQ-PAY-042
  statement: >
    The service MUST reject a payment approval when the request tenant
    differs from the tenant owning the payment object.
  rationale: prevent cross-tenant authorization failure
  acceptance:
    - same-tenant-authorized-user: allow
    - different-tenant-user: deny-and-audit
    - missing-tenant-context: deny
  traces_to: [ARCH-AUTH-007, TEST-AUTH-042]
```

Failure patterns

Treating quality requirements 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 quality requirements.
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 Data and information requirements

Data and information requirements must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Specify entities, meaning, quality, lineage, privacy, retention, residency, synchronization, backup, and recovery. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 data and information requirements.

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 data and information requirements 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 data and information requirements.
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 Operational and lifecycle requirements

Operational and lifecycle requirements must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Cover deployment, configuration, observability, support, migration, update, rollback, incident, decommission, and 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 operational and lifecycle requirements.

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 operational and lifecycle requirements 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 operational and lifecycle requirements.
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

Rewrite vague requirements such as fast, secure, intuitive, and highly available into measurable quality scenarios and acceptance thresholds.

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 quality requirements to energy, carbon, model behavior, agent autonomy, provenance, and cryptographic agility.

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

Specification forms and normative language

Choose representations that preserve meaning for readers, tools, tests, and implementation.

Chapter operating position

Choose representations that preserve meaning for readers, tools, tests, and implementation.



4.1 Natural-language requirements

Natural-language requirements must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use one obligation per statement, defined subjects, active voice, explicit conditions, measurable outcomes, and controlled terms. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 natural-language requirements.

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 natural-language requirements 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 natural-language requirements.
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 Use cases, scenarios, and state models

Use cases, scenarios, and state models must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Represent normal, alternate, exception, failure, recovery, concurrency, and lifecycle 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 use cases, scenarios, and state models.

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

```
requirement:
  id: REQ-PAY-042
  statement: >
    The service MUST reject a payment approval when the request tenant
    differs from the tenant owning the payment object.
  rationale: prevent cross-tenant authorization failure
  acceptance:
    - same-tenant-authorized-user: allow
    - different-tenant-user: deny-and-audit
    - missing-tenant-context: deny
  traces_to: [ARCH-AUTH-007, TEST-AUTH-042]
```

Failure patterns

Treating use cases, scenarios, and state models 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 use cases, scenarios, and state models.
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 Schemas, tables, decision models, and examples

Schemas, tables, decision models, and examples must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use structured data, truth tables, decision tables, examples, counterexamples, invariants, and formulas. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 schemas, tables, decision models, and examples.

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 schemas, tables, decision models, and examples 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 schemas, tables, decision models, and examples.
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 Normative key words and interpretation

Normative key words and interpretation must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use MUST, SHOULD, MAY, prohibited behavior, rationale, exceptions, and precedence consistently and intentionally. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 normative key words and interpretation.

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 normative key words and interpretation 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 normative key words and interpretation.
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 the same capability as normative prose, a state model, a decision table, and examples; compare ambiguity and testability.

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 machine-readable specifications, formal constraints, natural-language compilers, and proof-linked requirements.

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

Acceptance criteria and executable examples

Define what evidence proves a requirement is satisfied before implementation begins.

Chapter operating position

Define what evidence proves a requirement is satisfied before implementation begins.



5.1 Given-when-then and scenario criteria

Given-when-then and scenario criteria must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Specify initial state, action, expected output, side effects, timing, identity, and evidence without encoding accidental design. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 given-when-then and scenario criteria.

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 given-when-then and scenario criteria 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 given-when-then and scenario criteria.
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 Boundary, negative, and abuse criteria

Boundary, negative, and abuse criteria must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Cover invalid input, unauthorized action, concurrency, rate, scale, dependency loss, partial failure, and recovery. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 boundary, negative, and abuse criteria.

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

```
requirement:
  id: REQ-PAY-042
  statement: >
    The service MUST reject a payment approval when the request tenant
    differs from the tenant owning the payment object.
  rationale: prevent cross-tenant authorization failure
  acceptance:
    - same-tenant-authorized-user: allow
    - different-tenant-user: deny-and-audit
    - missing-tenant-context: deny
  traces_to: [ARCH-AUTH-007, TEST-AUTH-042]
```

Failure patterns

Treating boundary, negative, and abuse criteria 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 boundary, negative, and abuse criteria.
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 Quality and service acceptance

Quality and service acceptance must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define latency distributions, throughput, availability, error budget, accessibility, security, data quality, and operational readiness. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 quality and service acceptance.

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 quality and service acceptance 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 quality and service acceptance.
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 Release and completion evidence

Release and completion evidence must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Require tests, observability, runbooks, rollback, migration, documentation, vulnerability status, owner approval, and residual limitations. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 release and completion evidence.

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 release and completion evidence 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 release and completion evidence.
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

Develop acceptance criteria for a high-risk workflow with positive, negative, boundary, security, failure, recovery, and operational cases.

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 generated executable examples, property-based acceptance, digital-twin validation, and continuous conformance monitoring.

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

Validation, review, and conflict resolution

Detect requirement defects before they become architecture, code, or test defects.

Chapter operating position

Detect requirement defects before they become architecture, code, or test defects.



6.1 Quality review criteria

Quality review criteria must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Assess necessity, correctness, clarity, singularity, completeness, consistency, feasibility, priority, traceability, and verifiability. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 quality review criteria.

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 quality review criteria 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 quality review criteria.
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 Conflict and trade-off analysis

Conflict and trade-off analysis must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Resolve competing user, security, performance, privacy, schedule, cost, accessibility, and operational 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 conflict and trade-off 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

```
requirement:
  id: REQ-PAY-042
  statement: >
    The service MUST reject a payment approval when the request tenant
    differs from the tenant owning the payment object.
  rationale: prevent cross-tenant authorization failure
  acceptance:
    - same-tenant-authorized-user: allow
    - different-tenant-user: deny-and-audit
    - missing-tenant-context: deny
  traces_to: [ARCH-AUTH-007, TEST-AUTH-042]
```

Failure patterns

Treating conflict and trade-off 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 conflict and trade-off 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.



6.3 Risk and uncertainty

Risk and uncertainty must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Record unknowns, assumptions, confidence, experiments, decision deadlines, reversibility, and contingency. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 risk and uncertainty.

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 risk and uncertainty 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 risk and uncertainty.
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 Stakeholder validation and sign-off

Stakeholder validation and sign-off must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use walkthroughs, prototypes, examples, simulations, test designs, and explicit acceptance or objection records. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 stakeholder validation and sign-off.

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 stakeholder validation and sign-off 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 stakeholder validation and sign-off.
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

Peer-review a flawed specification, classify every defect, negotiate conflicts, and produce a revised acceptance record.

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 consistency checking, formal satisfiability analysis, and multi-agent adversarial specification review.

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

Traceability, change, and requirements lifecycle

Preserve intent as software, organizations, evidence, and environments evolve.

Chapter operating position

Preserve intent as software, organizations, evidence, and environments evolve.

7.1 Requirement identity and hierarchy

Requirement identity and hierarchy must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use stable IDs, parent-child relationships, allocation, rationale, source, owner, priority, status, and version. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 requirement identity and hierarchy.

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 requirement identity and hierarchy 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 requirement identity and hierarchy.
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 Bidirectional traceability

Bidirectional traceability must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Connect objectives, requirements, architecture, code, configuration, tests, controls, releases, incidents, and 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 bidirectional traceability.

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

```
requirement:
  id: REQ-PAY-042
  statement: >
    The service MUST reject a payment approval when the request tenant
    differs from the tenant owning the payment object.
  rationale: prevent cross-tenant authorization failure
  acceptance:
    - same-tenant-authorized-user: allow
    - different-tenant-user: deny-and-audit
    - missing-tenant-context: deny
  traces_to: [ARCH-AUTH-007, TEST-AUTH-042]
```

Failure patterns

Treating bidirectional traceability 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 bidirectional traceability.
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 Change impact and baselines

Change impact and baselines must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Evaluate affected interfaces, data, users, dependencies, tests, operations, suppliers, migration, risk, and documentation. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 change impact and baselines.

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 change impact and baselines 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 change impact and baselines.
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 Supersession and retirement

Supersession and retirement must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Record replacement, merge, split, deprecation, withdrawal, historical reason, route, and retained 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 supersession and retirement.

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 supersession and retirement 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 supersession and retirement.
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

Create a traceability graph and process a change request that affects architecture, API, data, tests, operations, and security.

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 a living specification graph that automatically identifies drift while preserving human product authority.

Review gate

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

CHAPTER 08

Specification-driven and AI-assisted development

Make the specification a controlled build input rather than a forgotten planning artifact.

Chapter operating position

Make the specification a controlled build input rather than a forgotten planning artifact.



8.1 Repository and documentation architecture

Repository and documentation architecture must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Place requirements, decisions, interfaces, schemas, examples, tests, prompts, and evidence under version control. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 repository and documentation architecture.

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 repository and documentation architecture 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 repository and documentation architecture.
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 Work decomposition and implementation contracts

Work decomposition and implementation contracts must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Partition work by bounded scope, prerequisites, files, interfaces, invariants, tests, review, 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 work decomposition and implementation 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.

Implementation sketch

```
requirement:
  id: REQ-PAY-042
  statement: >
    The service MUST reject a payment approval when the request tenant
    differs from the tenant owning the payment object.
  rationale: prevent cross-tenant authorization failure
  acceptance:
    - same-tenant-authorized-user: allow
    - different-tenant-user: deny-and-audit
    - missing-tenant-context: deny
  traces_to: [ARCH-AUTH-007, TEST-AUTH-042]
```

Failure patterns

Treating work decomposition and implementation 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 work decomposition and implementation 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.



8.3 AI implementer instructions and context

AI implementer instructions and context must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Provide authoritative sources, exclusions, style, constraints, tools, checkpoints, budgets, and prohibited assumptions. 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 ai implementer instructions and context.

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 ai implementer instructions and context 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 ai implementer instructions and context.
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 Verified completion and learning

Verified completion and learning must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Compare implementation with requirements, execute acceptance, inspect diffs, record deviations, and revise the specification when reality 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 objective, stakeholder need, requirement, specification, acceptance criterion, implementation, test, and release evidence 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 verified completion and learning.

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 verified completion and learning 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 verified completion and learning.
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

Use a specification to direct a human or coding agent through one feature, requiring plan, implementation, tests, review, and evidence-backed completion.

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 executable specifications, agent-readable contracts, formal plan checking, and autonomous development constrained by verified acceptance.

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. ISO / IEC / IEEE - ISO/IEC/IEEE 29148:2018 - Requirements Engineering

Role: Published requirements-engineering processes and information items.

Verified: 2026-07-26

Official source: <https://www.iso.org/standard/72089.html>

02. ISO / IEC - ISO/IEC 25010:2023 - Product Quality Model

Role: Current product-quality characteristics for systems and software.

Verified: 2026-07-26

Official source: <https://www.iso.org/standard/78176.html>

03. IETF / RFC Editor - RFC 2119 - Key Words for Use in RFCs to Indicate Requirement Levels

Role: Normative requirement-language meanings.

Verified: 2026-07-26

Official source: <https://www.rfc-editor.org/info/rfc2119>

04. IETF / RFC Editor - RFC 8174 - Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words

Role: Current clarification for normative key-word interpretation.

Verified: 2026-07-26

Official source: <https://www.rfc-editor.org/info/rfc8174>

05. NIST - SP 800-218 - Secure Software Development Framework 1.1

Role: Secure development practices integrated into software lifecycle and acceptance.

Verified: 2026-07-26

Official source: <https://csrc.nist.gov/pubs/sp/800/218/final>

06. ISO / IEC / IEEE - ISO/IEC/IEEE 29119-1:2022 - General Concepts

Role: General software-testing concepts and terminology.

Verified: 2026-07-26

Official source: <https://www.iso.org/standard/81291.html>

07. ISO / IEC / IEEE - ISO/IEC/IEEE 29119-2:2021 - Test Processes

Role: Organizational, management, and dynamic test processes.

Verified: 2026-07-26

Official source: <https://www.iso.org/standard/79428.html>

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-SEC; MYTHOS-UX
Publication class	Field Guide / Canonical Living Overview Guide
Skill levels	L1 Foundational; L2 Practitioner; L3 Advanced; L4 Frontier
Tags	requirements engineering, software specification, acceptance criteria, normative language, traceability, quality attributes, use cases, specification driven development, AI coding agents, product discovery
Canonical route	/library/field-guides/mythos-fg-swe-0001/

Living publication

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