CANONICAL LIVING OVERVIEW GUIDE

Software Architecture, Modularity, and Domain-Driven Design

A software-architecture guide covering architectural drivers, views and viewpoints, modularity, coupling and cohesion, domain-driven design, bounded contexts, dependency direction, monoliths, services, events, plugins, data ownership, architecture decisions, evolutionary design, governance, fitness functions, and

PERMANENT ID
MYTHOS-FG-SWE-0002

PUBLICATION CLASS
FIELD GUIDE

VERSION / STATUS
1.0.0-candidate / CANDIDATE

PERMANENT PUBLICATION IDENTITY

Software Architecture, Modularity, and Domain-Driven Design

Document ID
MYTHOS-FG-SWE-0002

Version
1.0.0-candidate

Status
Candidate Focused Field Guide

Review date
2026-10-26

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

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

Reader orientation

Dimension	Engineering position
Primary domain	MYTHOS-SWE - Software Engineering & Design
Secondary domain	MYTHOS-CLD; MYTHOS-DAT; MYTHOS-SEC
Audience	Software architects, staff and principal engineers, technical leads, platform engineers, domain experts, and AI-assisted development teams.
Prerequisites	Requirements, software design, APIs, data, testing, deployment, and basic distributed-systems concepts.
Publisher / owner	Mythos Systems / Aaron Stovall
Public classification	Public technical guidance

Expected outcomes

Translate architectural drivers and quality requirements into explicit structures, boundaries, and decisions.

Design modules with high cohesion, controlled coupling, stable contracts, and independent change where justified.

Use strategic and tactical domain-driven design to align software boundaries with domain meaning and ownership.

Choose monolith, modular monolith, services, event-driven, plugin, and hybrid structures according to evidence.

Evolve and validate architecture through decisions, dependency rules, fitness functions, telemetry, and refactoring.

How to use this guide

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

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

Learning and assurance model

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

Contents

CHAPTER 01

Architecture drivers, stakeholders, and descriptions

Define the concerns and evidence an architecture must address.

Chapter operating position
Define the concerns and evidence an architecture must address.

1.1 Architectural drivers

Architectural drivers must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Identify mission, requirements, quality scenarios, constraints, risk, scale, data, security, team structure, 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 architectural drivers.
- 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 architectural drivers 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 architectural drivers.
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 and concerns

Stakeholders and concerns must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Map users, developers, operators, security, data, support, executives, vendors, regulators, and their architectural questions. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 and concerns.
- 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

```
module:
  name: payment-domain
  owns: [payment, approval-policy, payment-events]
  may_depend_on: [shared-kernel]
  must_not_depend_on: [database-driver, web-framework]
  contracts:
    inbound: [ApprovePayment]
    outbound: [PaymentApproved, PaymentRejected]
  fitness_functions:
    - dependency-direction
    - no-cross-context-database-write
    - public-contract-compatibility
```

Failure patterns

- Treating stakeholders and concerns 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 and concerns.
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 Views, viewpoints, and model kinds

Views, viewpoints, and model kinds must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use context, container, component, code, deployment, data, runtime, security, and operational views 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 views, viewpoints, and model kinds.

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 views, viewpoints, and model kinds 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 views, viewpoints, and model kinds.
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 Architecture description and consistency

Architecture description and consistency must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Maintain elements, relationships, rationale, correspondence, assumptions, decisions, and cross-view consistency. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 architecture description and consistency.
- 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 architecture description and consistency 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 architecture description and consistency.
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
Produce a multi-view architecture description for a representative system and reconcile contradictions between runtime, data, and deployment views.

Laboratory evidence package

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

Frontier extension

Extend architecture descriptions to machine-readable graphs, digital twins, assurance cases, and agent-consumable context.

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

CHAPTER 02

Modularity, cohesion, coupling, and dependency direction

Create boundaries that localize change and make system behavior understandable.

Chapter operating position
Create boundaries that localize change and make system behavior understandable.

2.1 Module responsibilities and cohesion

Module responsibilities and cohesion must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Group behavior and data around a focused reason to change, domain concept, capability, lifecycle, or policy. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 module responsibilities and cohesion.
- 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 module responsibilities and cohesion 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 module responsibilities and cohesion.
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 Forms of coupling

Forms of coupling must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Analyze source, build, runtime, data, temporal, deployment, organizational, semantic, and operational coupling. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 forms of coupling.
- 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

```
module:
  name: payment-domain
  owns: [payment, approval-policy, payment-events]
  may_depend_on: [shared-kernel]
  must_not_depend_on: [database-driver, web-framework]
  contracts:
    inbound: [ApprovePayment]
    outbound: [PaymentApproved, PaymentRejected]
  fitness_functions:
    - dependency-direction
    - no-cross-context-database-write
    - public-contract-compatibility
```

Failure patterns

- Treating forms of coupling 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 forms of coupling.
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 Dependency inversion and stable abstractions

Dependency inversion and stable abstractions must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Direct dependencies toward durable policy and contracts rather than volatile mechanisms and frameworks. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 dependency inversion and stable abstractions.

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 dependency inversion and stable abstractions 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 dependency inversion and stable abstractions.
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 Encapsulation and information hiding

Encapsulation and information hiding must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Protect invariants, state, implementation choices, concurrency, errors, and data representation behind narrow interfaces. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 encapsulation and information hiding.
- 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 encapsulation and information hiding 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 encapsulation and information hiding.
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
Refactor a tightly coupled design into modules, measure dependency changes, and demonstrate that one volatile mechanism can be replaced independently.

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 language-enforced architecture, capability modules, effect systems, and formally checked dependency constraints.

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

Strategic domain-driven design

Align architecture with domain language, business capabilities, and organizational ownership.

Chapter operating position
Align architecture with domain language, business capabilities, and organizational ownership.

3.1 Domains, subdomains, and core differentiation

Domains, subdomains, and core differentiation must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Separate core, supporting, and generic capabilities according to business value, complexity, and investment. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 domains, subdomains, and core differentiation.
- 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 domains, subdomains, and core differentiation 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 domains, subdomains, and core differentiation.
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 Bounded contexts

Bounded contexts must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define the boundary within which a model and language are consistent, owned, versioned, and independently evolved. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 bounded contexts.
- 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

```
module:
  name: payment-domain
  owns: [payment, approval-policy, payment-events]
  may_depend_on: [shared-kernel]
  must_not_depend_on: [database-driver, web-framework]
  contracts:
    inbound: [ApprovePayment]
    outbound: [PaymentApproved, PaymentRejected]
  fitness_functions:
    - dependency-direction
    - no-cross-context-database-write
    - public-contract-compatibility
```

Failure patterns

- Treating bounded contexts 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 bounded contexts.
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 Context mapping and relationships

Context mapping and relationships must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use partnership, shared kernel, customer-supplier, conformist, anticorruption layer, open host, and published language patterns. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 context mapping and relationships.

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 context mapping and relationships 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 context mapping and relationships.
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 Team and governance alignment

Team and governance alignment must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Assign ownership, decision authority, interfaces, service levels, roadmaps, and conflict resolution around contexts. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 team and governance alignment.
- Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.
- Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.
- Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.
- Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.
- Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

- Treating team and governance alignment as a document or tool activity disconnected from actual system behavior.
- Relying on intended design or configuration without proving implementation and runtime outcomes.
- Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.
- Changing several variables before preserving the original state and stating a falsifiable hypothesis.
- Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.
- Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

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

Chapter 03 laboratory and review

Practical laboratory
Create a domain map and bounded-context model for a complex enterprise workflow, including two conflicting meanings of the same term.

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 AI-assisted domain discovery, semantic graphs, executable context maps, and dynamic organizational topology.

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

Tactical domain modeling and invariants

Represent domain behavior without allowing frameworks or databases to become the model.

Chapter operating position
Represent domain behavior without allowing frameworks or databases to become the model.

4.1 Entities, value objects, and identity

Entities, value objects, and identity must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Separate lifecycle identity from descriptive values, equality, immutability, validation, and temporal history. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 entities, value objects, and identity.
- Model expected states, boundaries, inputs, outputs, failure modes, and evidence before implementation.
- Validate intended design, implemented behavior, runtime state, and user or mission outcome independently.
- Test positive, negative, boundary, degraded, failed, recovery, scale, and rollback conditions.
- Preserve source, version, configuration, data, telemetry, artifacts, timestamps, decisions, and limitations.
- Automate through typed contracts, least privilege, deterministic gates, bounded concurrency, canaries, and reversibility.

Failure patterns

- Treating entities, value objects, and identity as a document or tool activity disconnected from actual system behavior.
- Relying on intended design or configuration without proving implementation and runtime outcomes.
- Ignoring failure, concurrency, data, identity, environment, compatibility, and operational boundaries.
- Changing several variables before preserving the original state and stating a falsifiable hypothesis.
- Allowing generated code, tests, findings, or optimizations to proceed without authoritative inputs and review.
- Declaring success after one positive example without negative tests, reproducibility, rollback, and recurrence checks.

Validation and evidence

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

4.2 Aggregates and consistency boundaries

Aggregates and consistency boundaries must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Protect invariants, transaction scope, references, concurrency, event publication, and loading 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 aggregates and consistency boundaries.
- 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

```
module:
  name: payment-domain
  owns: [payment, approval-policy, payment-events]
  may_depend_on: [shared-kernel]
  must_not_depend_on: [database-driver, web-framework]
  contracts:
    inbound: [ApprovePayment]
    outbound: [PaymentApproved, PaymentRejected]
  fitness_functions:
    - dependency-direction
    - no-cross-context-database-write
    - public-contract-compatibility
```

Failure patterns

- Treating aggregates and consistency boundaries 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 aggregates and consistency boundaries.
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 Domain services, policies, and specifications

Domain services, policies, and specifications must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Represent operations and rules that do not naturally belong to one entity or value object. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 services, policies, and specifications.

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 services, policies, and specifications 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 services, policies, and specifications.
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 Domain events and process managers

Domain events and process managers must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Capture facts, reactions, long-running workflows, idempotency, compensation, and ownership across boundaries. The design should identify authoritative inputs, states, boundaries, assumptions, dependencies, limits, ownership, telemetry, failure behavior, and acceptance criteria.

This guide traces the subject through the driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 events and process managers.
- 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 events and process managers 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 events and process managers.
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
Model a domain with entities, values, aggregates, policies, and events; test invariants under concurrency and partial failure.

Laboratory evidence package

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

Frontier extension

Explore event-sourced domain models, temporal logic, verified invariants, and agent-operated business processes.

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

Architectural styles and deployment boundaries

Choose structural patterns according to change, scale, failure, and ownership needs.

Chapter operating position

Choose structural patterns according to change, scale, failure, and ownership needs.

5.1 Layered, ports-and-adapters, and clean structures

Layered, ports-and-adapters, and clean structures must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Separate domain policy, application orchestration, interfaces, infrastructure, and framework concerns. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 layered, ports-and-adapters, and clean structures.
- 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 layered, ports-and-adapters, and clean structures 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 layered, ports-and-adapters, and clean structures.
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 Modular monoliths

Modular monoliths must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use strong internal boundaries, one deployable, transactional simplicity, explicit contracts, and extraction 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 modular monoliths.
- 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

```
module:
  name: payment-domain
  owns: [payment, approval-policy, payment-events]
  may_depend_on: [shared-kernel]
  must_not_depend_on: [database-driver, web-framework]
  contracts:
    inbound: [ApprovePayment]
    outbound: [PaymentApproved, PaymentRejected]
  fitness_functions:
    - dependency-direction
    - no-cross-context-database-write
    - public-contract-compatibility
```

Failure patterns

- Treating modular monoliths 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 modular monoliths.
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 Services and microservices

Services and microservices must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Evaluate independent deployment, data ownership, failure, latency, observability, platform cost, and organizational autonomy. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 services and microservices.

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 services and microservices 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 services and microservices.
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 Event-driven, pipeline, actor, and plugin architectures

Event-driven, pipeline, actor, and plugin architectures must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use asynchronous flows, stages, mailboxes, extensions, capability isolation, and contract governance deliberately. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 event-driven, pipeline, actor, and plugin architectures.
- 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 event-driven, pipeline, actor, and plugin architectures 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 event-driven, pipeline, actor, and plugin architectures.
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
Compare three architecture styles for one set of drivers and build a small modular implementation demonstrating the selected trade-offs.

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 WebAssembly components, local-first architectures, serverless actors, and dynamically governed plugin fabrics.

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

Data, integration, and boundary architecture

Prevent shared state and integration shortcuts from dissolving architectural boundaries.

Chapter operating position
Prevent shared state and integration shortcuts from dissolving architectural boundaries.

6.1 Data ownership and schemas

Data ownership and schemas must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Assign authoritative write ownership, read models, identifiers, history, privacy, migration, and cross-context access. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 ownership and schemas.
- 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 ownership and schemas 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 ownership and schemas.
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 Synchronous integration

Synchronous integration must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Design APIs, timeouts, errors, compatibility, retries, authorization, and failure isolation across calls. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 synchronous integration.
- 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

```
module:
  name: payment-domain
  owns: [payment, approval-policy, payment-events]
  may_depend_on: [shared-kernel]
  must_not_depend_on: [database-driver, web-framework]
  contracts:
    inbound: [ApprovePayment]
    outbound: [PaymentApproved, PaymentRejected]
  fitness_functions:
    - dependency-direction
    - no-cross-context-database-write
    - public-contract-compatibility
```

Failure patterns

- Treating synchronous integration 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 synchronous integration.
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 Asynchronous integration

Asynchronous integration must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Design events, commands, schemas, delivery, ordering, idempotency, replay, dead letters, and consumer autonomy. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 asynchronous integration.

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 asynchronous integration 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 asynchronous integration.
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 Anticorruption and translation layers

Anticorruption and translation layers must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Protect domain meaning when integrating legacy systems, vendors, shared models, or inconsistent terminology. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 anticorruption and translation layers.
- 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 anticorruption and translation layers 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 anticorruption and translation layers.
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
Design an integration between three bounded contexts using APIs and events, including schema evolution, failure, replay, and data ownership.

Laboratory evidence package

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

Frontier extension

Explore semantic interoperability, data contracts, federated schemas, and AI-generated translation with verified meaning.

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

Architecture decisions, governance, and fitness functions

Make architectural intent visible, testable, and revisable.

Chapter operating position
Make architectural intent visible, testable, and revisable.

7.1 Architecture decision records

Architecture decision records must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Record context, decision, options, forces, consequences, status, owners, evidence, and supersession. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 architecture decision records.
- 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 architecture decision records 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 architecture decision records.
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 Principles, standards, and paved roads

Principles, standards, and paved roads must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Define mandatory invariants, preferred patterns, reusable platforms, exceptions, and evidence without centralizing every decision. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 principles, standards, and paved roads.
- 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

```
module:
  name: payment-domain
  owns: [payment, approval-policy, payment-events]
  may_depend_on: [shared-kernel]
  must_not_depend_on: [database-driver, web-framework]
  contracts:
    inbound: [ApprovePayment]
    outbound: [PaymentApproved, PaymentRejected]
  fitness_functions:
    - dependency-direction
    - no-cross-context-database-write
    - public-contract-compatibility
```

Failure patterns

- Treating principles, standards, and paved roads 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 principles, standards, and paved roads.
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 Fitness functions and automated checks

Fitness functions and automated checks must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Test dependency direction, API compatibility, data ownership, security, performance, resilience, and deployment 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 fitness functions and automated checks.

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 fitness functions and automated checks 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 fitness functions and automated checks.
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 Architecture review and assurance

Architecture review and assurance must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use risk-based review, prototypes, tests, telemetry, incidents, cost, and operational evidence instead of diagrams alone. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 architecture review and assurance.
- 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 architecture review and assurance 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 architecture review and assurance.
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 ADRs and automated architecture tests for a modular system, then process an exception and later supersede one decision.

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 continuously verified architecture graphs and governed architecture agents that propose but cannot silently change constraints.

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

Evolution, refactoring, and architecture recovery

Change architecture incrementally while preserving service and evidence.

Chapter operating position
Change architecture incrementally while preserving service and evidence.

8.1 Architecture discovery and recovery

Architecture discovery and recovery must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Infer actual components, dependencies, data flows, runtime calls, deployments, ownership, and drift from code and telemetry. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 architecture discovery and recovery.
- 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 architecture discovery and recovery 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 architecture discovery and recovery.
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 Evolutionary refactoring

Evolutionary refactoring must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use seams, strangler patterns, branch by abstraction, parallel change, migration adapters, canaries, and rollback. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 evolutionary refactoring.
- 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

```
module:
  name: payment-domain
  owns: [payment, approval-policy, payment-events]
  may_depend_on: [shared-kernel]
  must_not_depend_on: [database-driver, web-framework]
  contracts:
    inbound: [ApprovePayment]
    outbound: [PaymentApproved, PaymentRejected]
  fitness_functions:
    - dependency-direction
    - no-cross-context-database-write
    - public-contract-compatibility
```

Failure patterns

- Treating evolutionary refactoring 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 evolutionary refactoring.
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 Technical debt and architecture risk

Technical debt and architecture risk must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Classify structural, dependency, data, operational, security, test, and knowledge debt by impact and options. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 technical debt and architecture risk.

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 technical debt and architecture risk 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 technical debt and architecture risk.
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 Architecture observability and learning

Architecture observability and learning must be treated as an observable engineering mechanism rather than a slogan, framework checkbox, or tool output. Use incidents, performance, changes, defects, team friction, cost, and user outcomes to revise architecture. 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 driver, domain, boundary, component, dependency, runtime interaction, deployment, decision, and fitness 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 architecture observability 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 architecture observability 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 architecture observability 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
Recover an architecture from an unfamiliar codebase, identify drift, and plan an incremental boundary restoration with tests and rollback.

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 architecture digital twins that combine repository graphs, runtime traces, data lineage, policy, and change simulation.

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 42010:2022 - Architecture Description
Role: Architecture-description concepts, viewpoints, model kinds, and required relationships.
Verified: 2026-07-26
Official source: https://www.iso.org/standard/74393.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. Object Management Group - Unified Modeling Language 2.5.1
Role: Standard modeling notation for structure, behavior, interaction, and deployment views.
Verified: 2026-07-26
Official source: https://www.omg.org/spec/UML/2.5.1/PDF

04. Domain Language - Domain-Driven Design Reference
Role: Authoritative summary of strategic and tactical DDD patterns.
Verified: 2026-07-26
Official source: https://www.domainlanguage.com/ddd/reference/

05. C4 Model - C4 Model for Visualising Software Architecture
Role: Practical hierarchical software-architecture visualization model.
Verified: 2026-07-26
Official source: https://c4model.com/

06. 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

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-CLD; MYTHOS-DAT; MYTHOS-SEC
Publication class	Field Guide / Canonical Living Overview Guide
Skill levels	L1 Foundational; L2 Practitioner; L3 Advanced; L4 Frontier
Tags	software architecture, modularity, domain driven design, bounded contexts, coupling, cohesion, architecture decisions, modular monolith, microservices, fitness functions
Canonical route	/library/field-guides/mythos-fg-swe-0002/

Living publication

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