

# **R.E.T.N.A PROTOCOL SPECIFICATION**

**Real-Time Execution Trust and Notarization Architecture**

**Decision Governance & Dispatch Trace Standard  
for Agentic and Model-Driven Systems**

---

Version: **v0.1**

Status: **Draft Specification**

Intended Status: **Industry Governance Protocol**

Publication Authority:  
**Value Intelligence Solutions Inc.**

Primary Author:  
**Gary Gayle**

Publication Date:  
**March 2026**

Document Identifier:  
**RETNA-SPEC-0.1**

---

**Governance Must Precede Execution**

---

**Value Intelligence Solutions Inc.**  
[www.valueintelligence.io](http://www.valueintelligence.io)

## Document Metadata

Document Title

R.E.T.N.A Protocol Specification

Version

v0.1

Document Identifier

RETNA-SPEC-0.1

Publication Status

Draft Specification

Publication Authority

Value Intelligence Solutions Inc.

Author

Gary Gayle

Publication Date

March 2026

Intended Audience

AI system architects, regulatory bodies, autonomous system developers, governance infrastructure providers.

---

## Intellectual Property Notice

This document describes the R.E.T.N.A Protocol architecture and governance framework.

All trademarks and intellectual property associated with the R.E.T.N.A Protocol are the property of Value Intelligence Solutions Inc.

## License

Distribution of this document for informational and educational purposes is permitted.

Implementation of the R.E.T.N.A Protocol may be subject to licensing terms defined by Value Intelligence Solutions Inc.

## Table of Contents

Document Metadata .....	1
Intellectual Property Notice .....	1
License .....	1
Table of Contents .....	2
Abstract (Informative).....	13
0. Notice and Intent (Normative).....	14
1. Design Principles (Normative).....	15
1.1 Execution Requires Authorization .....	15
1.2 Governance Is Enforceable .....	16
1.3 Model-Agnostic by Default.....	16
1.4 Agent-Agnostic by Default .....	16
1.5 Outcome-Centric, Not Architecture-Centric.....	17
1.6 Governed Consequence over Transport Semantics.....	17
1.7 Receipts over Narratives .....	18
1.8 Deterministic Authorization.....	18
2. Goals and Non-Goals (Normative) .....	18
2.1 Goals .....	18
2.2 Non-Goals .....	19
3. Terminology and Definitions (Normative).....	19
4. Protocol Overview (Normative).....	25

4.0 Normative Flow Summary.....	27
4.1 Capture Context .....	31
4.2 Normalize Evidence.....	31
4.3 Evaluate Policies and Constraints.....	32
4.4 Emit Governance Outcome.....	32
4.5 Resolve Governed State Transition .....	33
4.6 Emit Decision Receipt.....	33
4.7 Persist Receipt and Execute Authorized Decision .....	34
5. System Invariants (Normative).....	35
I1 — Governance Boundary Integrity .....	35
I2 — Mandatory Receipt Emission.....	36
I3 — Trace Completeness.....	36
I4 — Evidence Referencing.....	37
I5 — Policy Disclosure (Structured).....	37
I6 — Outcome Determinism .....	37
I7 — Tamper Awareness.....	38
I8 — Separation of Duties .....	38
I9 — Authority Integrity.....	39
I10 — Admissibility Requirement.....	39
I11 — Formation Integrity.....	40
I12 — Governed State Transition Integrity.....	40
5.1 Invariant Enforcement Scope .....	41

5.2 Invariant Violations .....	42
6. Actors and Roles (Normative) .....	42
6.1 Required Actors .....	42
6.1.1 Producer .....	42
6.1.2 Governor (R.E.T.N.A Implementation).....	43
6.1.3 Executor .....	44
6.1.4 Receipt Store.....	45
6.2 Optional Actors.....	46
6.2.1 Policy Authority.....	46
6.2.2 Evidence Store.....	46
6.2.3 Attestor.....	47
6.3 Role Composition and Separation .....	48
6.4 Role Identification and Traceability .....	48
7. Decision Classification (Normative).....	49
7.1 Classification Requirements .....	49
7.2 Risk Class (Required).....	50
7.2.1 Baseline Risk Classes.....	50
7.2.2 Risk Assignment Rules.....	50
7.3 Action Type (Required).....	51
7.3.1 Baseline Action Types .....	51
7.4 Decision Mode (Required).....	52
7.4.1 Decision Modes.....	52

7.4.2 Decision Mode Constraints .....	52
7.4.3 Formation and Transition Sensitivity.....	52
7.5 Classification Integrity and Disclosure.....	53
7.6 Relationship to Policy Evaluation .....	54
8. The Decision Receipt (Normative Core) .....	54
8.1 Decision Receipt Requirements.....	55
8.2 Receipt Lifecycle .....	56
8.3 Canonical Receipt Structure (v0.1).....	57
8.3.1 Receipt Header.....	57
8.3.2 Identity Block .....	58
8.3.3 Authority Context Block.....	58
8.3.4 Decision Envelope.....	59
8.3.5 Formation Context Block .....	60
8.3.6 State Context.....	61
8.3.7 Participants .....	61
8.3.8 Evidence Block.....	62
8.3.9 Policy Evaluation Block .....	63
8.3.10 Admissibility Block.....	63
8.3.11 Outcome Block.....	64
8.3.12 Integrity and Audit Block.....	65
8.3.13 Canonical Receipt Serialization .....	66
8.3.14 Canonical Receipt Example (Informative) .....	66

8.4 Receipt Immutability and Tamper Awareness .....	69
8.5 Minimal Receipt (RETNA-A) .....	70
9. Canonical Reason Codes (Normative) .....	71
9.1 Purpose of Reason Codes .....	72
9.2 Outcome Reason Codes (Baseline Set) .....	73
9.2.1 Policy Evaluation Codes .....	73
9.2.2 Evidence and Confidence Codes .....	73
9.2.3 Risk and Safety Codes .....	74
9.2.4 System and Dependency Codes.....	74
9.2.5 Governance Control Codes.....	75
9.2.6 Authority Integrity Codes.....	75
9.2.7 Admissibility and Formation Codes .....	76
9.2.8 Governed State Transition Codes .....	76
9.2.9 Continuity and State Drift Codes — Reserved Canonical Set.....	77
9.3 Constraint Severity Codes.....	79
9.4 Reason Code Usage Rules.....	79
9.5 Vendor Extensions .....	80
9.6 Relationship to Audit and Compliance .....	80
10. Protocol Flows (Normative) .....	81
10.1 Standard Flow — Governed Execution (ALLOW) .....	82
10.2 Deny Flow (DENY) .....	85
10.3 Escalation Flow (ESCALATE).....	85

10.4 Defer Flow (DEFER) .....	86
10.5 Degrade Flow (DEGRADE) .....	87
10.6 Retry and Continuity Rules.....	88
10.7 Failure Handling.....	88
10.8 Flow Determinism and Observability .....	89
11. Policy Interface (Normative) .....	89
11.1 Policy Interface Objectives.....	90
11.2 Policy Bundle Concept .....	90
11.2.1 Policy Bundle Requirements .....	91
11.3 Constraints .....	91
11.3.1 Constraint Interface .....	92
11.4 Policy Evaluation Contract.....	92
11.5 Determinism and Reproducibility .....	93
11.6 Policy Disclosure in Decision Receipts.....	93
11.7 Policy Authority and Trust.....	94
11.8 Policy Scope and Applicability .....	94
11.9 Extensibility .....	95
12. Evidence Handling (Normative).....	96
12.1 Evidence Principles.....	96
12.2 Evidence Representation.....	97
12.3 Evidence Item Requirements.....	98
12.4 Evidence Integrity and Verification.....	98

12.4.1 Integrity Guarantees by Profile ..... 99

12.4.2 Verification Failures ..... 100

12.5 Evidence Lifecycle and Retention ..... 100

12.6 Evidence Types (Baseline)..... 100

12.7 Evidence Minimization and Privacy..... 101

12.8 Evidence and Reproducibility..... 102

13. Trust, Security, and Privacy (Normative) ..... 102

13.1 Authentication and Authorization ..... 103

    13.1.1 Actor Authentication..... 103

    13.1.2 Authorization Scope..... 103

13.2 Least Privilege and Separation of Duties ..... 104

13.3 Tamper Resistance and Integrity ..... 105

    13.3.1 Receipt Integrity ..... 105

    13.3.2 Evidence Integrity..... 105

13.4 Privacy Controls ..... 106

    13.4.1 Data Minimization..... 106

    13.4.2 Access Classification ..... 106

    13.4.3 Retention Classification ..... 106

13.5 Human Override and Accountability..... 107

13.6 Threat Model (Baseline) ..... 107

13.7 Fail-Closed Behavior..... 108

14. Interoperability and Extensibility (Normative)..... 109

14.1 Interoperability Principles.....	109
14.2 Canonical Artifact Interoperability.....	110
14.2.1 Decision Receipt as the Interoperable Unit.....	110
14.2.2 Serialization and Canonicalization.....	110
14.2.3 Receipt Validation and Canonical Verification (Normative).....	111
14.3 Vendor Extensions.....	111
14.3.1 Permitted Extensions .....	111
14.3.2 Extension Constraints .....	112
14.3.3 Payload Kind Registry .....	112
14.3.4 Registry-Governed Extension Points.....	113
14.4 Capability Advertisement.....	113
14.5 Compliance Profiles (Normative).....	114
14.5.1 Baseline Compliance Profile Requirements .....	114
14.5.2 RETNA-A — Foundational Governance Profile.....	115
14.5.3 RETNA-B — Governed Execution Profile .....	116
14.5.4 RETNA-C — Critical Consequence Profile.....	116
14.6 Cross-System Validation.....	116
14.7 Backward and Forward Compatibility .....	117
14.7.1 Backward Compatibility.....	117
14.7.2 Forward Compatibility.....	117
14.8 Domain Profiles (Optional).....	117
15. Conformance and Testability (Normative) .....	119

15.1 Conformance Claims.....	119
15.2 Mandatory Conformance Tests .....	120
15.2.1 Governance Boundary Integrity Test .....	120
15.2.2 Receipt Emission Test.....	120
15.2.3 Receipt Completeness Test.....	121
15.2.4 Policy Trigger Test .....	121
15.2.5 Deterministic Outcome Test .....	122
15.2.6 Receipt Validator Conformance Test .....	122
15.2.8 Authority Integrity Validation Test.....	123
15.2.9 Admissibility Enforcement Validation Test .....	124
15.2.10 Formation Integrity Validation Test.....	125
15.2.11 Governed State Transition Blocking Test.....	125
15.2.12 Receipt Persistence Refusal Test.....	126
15.2.13 Runtime Governance Validation Doctrine.....	126
15.3 Profile-Specific Tests .....	127
15.3.1 Profile B (RETNA-B) Tests .....	127
15.3.2 Profile C (RETNA-C) Tests.....	127
15.4 Negative Testing Requirements .....	128
15.5 Audit Replay Capability.....	129
15.6 Independent Verification.....	129
16. Reference Implementation Guidance (Non-Normative).....	130
16.1 Architectural Placement .....	130

16.1.1 Runtime Governance Enforcement (Illustrative).....	131
16.2 Minimal Reference Components .....	133
16.3 Decision Proposal Interface (Illustrative) .....	134
16.4 Receipt Emission Strategy.....	134
16.5 Receipt Storage Patterns.....	135
16.6 Validation and Tooling.....	135
16.6.1 Reference Validator Tooling.....	136
16.7 HomeSphere AI as a Reference Implementation (Informative) .....	136
16.8 Common Implementation Pitfalls.....	136
17. Licensing and Implementation Rights (Normative) .....	137
18. Versioning and Evolution (Normative) .....	138
18.1 Version Identifier Format.....	138
18.2 Minor Version Rules .....	139
18.3 Major Version Rules.....	139
18.4 Deprecation Policy .....	140
18.5 Receipt Compatibility Across Versions .....	140
18.6 Governance of the Protocol Itself.....	141
18.7 Stability Guarantees.....	141
19. Adoption Path (Non-Normative, Strategic) .....	142
19.1 Adoption Philosophy.....	142
19.2 Stage 0 — Observability-Only Adoption .....	143
19.3 Stage 1 — Governance Boundary Introduction .....	143

19.4 Stage 2 — Selective Enforcement.....	144
19.5 Stage 3 — Full Compliance Profiles .....	144
19.6 Tooling and Ecosystem Enablement.....	145
19.7 Industry Standardization Path .....	145
19.8 Strategic Positioning.....	145
19.9 HomeSphere AI as a Catalyst (Informative).....	146
Closing Statement (Non-Normative) .....	146
Appendix A — Glossary.....	147
Appendix B — Payload Kind Registry (Informative Reference).....	164
Appendix C — Canonical Reason Code Registry (Informative Reference) .....	164

## Abstract (Informative)

Artificial intelligence systems are increasingly deployed as continuous decision engines rather than static prompt-response tools. As these systems begin acting upon real-world state—including inventory management, purchasing, access control, safety systems, logistics, healthcare workflows, and financial operations—trust, traceability, policy compliance, and enforceable execution integrity become prerequisites for safe and scalable adoption.

The **R.E.T.N.A Protocol** defines a vendor-neutral governance framework for autonomous and model-driven systems operating across digital, physical, and financial environments. The protocol introduces a mandatory **Governance Boundary** separating **Decision Construction** from **Decision Execution**, ensuring that actions capable of altering external state are evaluated against explicit policy constraints prior to execution commitment.

Through structured policy evaluation, deterministic authorization logic, formation-aware governance controls, and verifiable **Decision Receipts**, the protocol establishes a standardized method for recording the provenance, evidence inputs, policy outcomes, authorization results, and governed state transition context associated with governed actions.

The protocol governs not only authorization visibility, but enforceable execution integrity.

R.E.T.N.A implementations **MAY** refuse decision formation, invalidate inadmissible state transitions, prevent execution under unresolved governance conditions, and require receipt-backed authorization prior to governed consequence.

The protocol distinguishes transport from consequence, forwarding from authorization, and output generation from governed state transition commitment.

Under the protocol, systems are able to:

- capture **state deltas** representing what changed, or was authorized to change, within the environment
- record the **participating actors involved in decision construction, governance evaluation, and execution**
- reference the **evidence artifacts** used to construct and evaluate the decision
- evaluate **applicable policies, constraints, authority conditions, and governance classifications** prior to execution
- emit a **Decision Receipt** recording governance outcomes, authorization state, policy evaluation results, and supporting evidence
- prevent execution when governance integrity conditions cannot be satisfied

By transforming decisions into **auditable governance artifacts** rather than opaque model outputs, the R.E.T.N.A Protocol enables:

- accountable autonomous behavior
- interoperable governance enforcement across heterogeneous systems

- portable compliance evidence suitable for regulatory and enterprise environments
- deterministic authorization for AI-driven actions
- governed state transition enforcement prior to consequential execution
- durable auditability and reproducible governance verification

The protocol is designed to operate independently of model architecture, enabling adoption across heterogeneous agent frameworks, machine learning systems, robotics platforms, and distributed AI services.

R.E.T.N.A establishes **Decision Governance** as a first-class layer of the AI infrastructure stack, comparable in importance to security (TLS), identity (OAuth), and observability (OpenTelemetry), but focused explicitly on governing **decisions, authorization integrity, and execution consequence** rather than transport events, identities, or data flows.

---

## 0. Notice and Intent (Normative)

This document specifies **R.E.T.N.A Protocol v0.1** (“the Protocol”), a vendor-neutral, model-agnostic decision governance protocol for artificial intelligence systems that construct decisions and may execute actions capable of altering digital, financial, or physical state.

The Protocol defines normative requirements for systems that implement governed decision infrastructure, including the evaluation, authorization, and traceability of decisions prior to execution.

The key words **“MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL”** in this document are to be interpreted as described in RFC 2119 and RFC 8174 when, and only when, they appear in all capital letters.

The Protocol defines the following core architectural elements:

- a **mandatory Governance Boundary** separating Decision Construction from Decision Execution
- a **decision classification framework** based on Payload Kind that determines the governance context for policy evaluation and routing behavior
- a **policy-driven authorization mechanism** capable of allowing, denying, escalating, deferring, or degrading decisions
- a **uniform, auditable Decision Receipt artifact** produced for governed decisions
- a **traceability and provenance model** that produces human-auditable and machine-verifiable records of governed decisions

- **canonical Reason Codes** that communicate authorization outcomes and policy evaluation results
- **Edge Policy enforcement**, defining which actors and systems are authorized to issue or execute governed decisions

The Protocol is designed to be implemented by independent vendors, system architects, and platform providers while enabling interoperability, auditability, and regulatory alignment across heterogeneous AI ecosystems.

This Protocol governs **decisions**, not **models**.

The Protocol evaluates **commitments to act**, rather than the internal reasoning processes used to construct those commitments. Model architecture, inference strategy, and implementation details remain outside the scope of this specification.

Implementations claiming compliance with the R.E.T.N.A Protocol **MUST** implement the governance workflow, Decision Receipt construction, and policy evaluation semantics defined in this specification.

Any system that executes governed decisions without constructing a valid Decision Receipt prior to execution **MUST NOT** be considered R.E.T.N.A-compliant.

---

## 1. Design Principles (Normative)

The R.E.T.N.A Protocol is governed by two foundational axioms:

**Governance Must Precede Execution.**  
**This Protocol governs decisions, not models.**

All protocol design decisions and governance mechanisms derive from these principles.

### 1.1 Execution Requires Authorization

Execution without authorization **MUST be impossible**.

Any action capable of altering external state—including digital systems, financial resources, or physical environments—**MUST** pass through the **Governance Boundary** and receive an explicit authorization outcome before execution.

Systems implementing the Protocol **MUST** enforce this requirement at the execution interface, ensuring that unauthorized actions cannot bypass policy evaluation, **Edge Policy enforcement**, governance mechanisms, or **Decision Receipt** generation.

This requirement establishes governance as a mandatory control layer rather than an optional advisory system.

## 1.2 Governance Is Enforceable

The Protocol **MUST** support blocking, modifying, escalating, deferring, degrading, or preventing governed decisions prior to execution. Passive logging alone is insufficient.

Governance systems that merely record decisions after execution cannot prevent unsafe, unauthorized, or non-compliant outcomes. Implementations **MUST** therefore ensure that governance mechanisms operate as an enforceable authorization layer capable of actively altering, constraining, invalidating, or preventing actions before consequential execution occurs.

Governance **MUST NOT** be limited to post-output evaluation or execution-time filtering.

Implementations **SHOULD** preserve sufficient formation integrity to explain how candidate decisions were shaped prior to admissibility evaluation and Governance Boundary enforcement.

Authorization outcomes **MUST** be communicated through standardized **Reason Codes**.

## 1.3 Model-Agnostic by Default

The Protocol **MUST NOT** assume any specific large language model, embedding model, vendor API, or inference architecture.

Autonomous systems may employ heterogeneous reasoning engines including machine learning models, rule engines, heuristics, or hybrid approaches. The Protocol therefore governs the **decision artifact**, rather than the internal reasoning mechanism that produced it.

## 1.4 Agent-Agnostic by Default

The Protocol **MUST** support single-agent, multi-agent, tool-using, and hybrid systems without modification.

The governance architecture **MUST** operate consistently regardless of whether decisions originate from:

- a single autonomous agent
- a coordinated multi-agent workflow
- a system combining automated and human participants

## 1.5 Outcome-Centric, Not Architecture-Centric

The Protocol evaluates governed **decisions and consequential actions**, not internal reasoning strategies, prompts, or chains-of-thought.

Governance focuses on the **proposed outcome** and its **potential state consequence** rather than the intermediate reasoning steps that produced it. This approach preserves model interoperability while ensuring that authorization, admissibility evaluation, and execution governance remain deterministic and auditable.

Decision evaluation **MUST** consider the **Payload Kind** associated with the decision request.

Governed consequence **MUST NOT** be inferred from transport success, message forwarding, routing completion, or output generation.

## 1.6 Governed Consequence over Transport Semantics

Governed actions capable of operational, legal, financial, physical, or material consequence **MUST** resolve into explicit governed state transitions prior to execution commitment.

Execution authorization **MUST** be derived from **governance evaluation outcomes** rather than transport behavior, delivery success, orchestration completion, or model output generation.

Systems implementing the Protocol **MUST** distinguish:

- message forwarding from execution authorization
- routing completion from admissibility satisfaction
- output generation from governed consequence commitment
- transport observability from state transition authorization

The successful transmission, routing, or delivery of a decision artifact **MUST NOT** independently authorize consequential execution.

This requirement establishes **governed state transition enforcement** as distinct from communication-layer success semantics.

## 1.7 Receipts over Narratives

Governance requires structured, attestable artifacts, not prose explanations or free-form justifications.

Decisions governed by the Protocol **MUST** produce structured **Decision Receipts** (historically referred to as "**Dispatch Receipts**") that record evidence inputs, policy evaluations, authorization outcomes, and associated **Reason Codes**.

These artifacts enable deterministic auditing, forensic replay, and regulatory verification.

## 1.8 Deterministic Authorization

Authorization outcomes **MUST** be deterministic given identical policy configuration, evidence inputs, and evaluation context.

Implementations **MUST** ensure that identical decision requests evaluated under the same policy state produce the same authorization outcome and **Reason Codes**, enabling reproducibility, auditability, and regulatory verification.

---

## 2. Goals and Non-Goals (Normative)

### 2.1 Goals

The Protocol **MUST** enable the following capabilities:

- **Decision gating prior to execution**, producing one of the following authorization outcomes:  
ALLOW, DENY, ESCALATE, DEFER, or DEGRADE.
- **Traceable provenance of the decision pathway**, including the actors, evidence artifacts, and policy evaluations that contributed to the decision outcome.
- **Policy and constraint enforcement independent of model selection**, ensuring that governance operates consistently across heterogeneous reasoning engines.
- **Evidence capture**, including relevant **state deltas**, signals, and contextual inputs used during decision construction and policy evaluation.

- **Interoperable Decision Receipts** that can be exchanged, validated, and audited across independent systems and vendors.
- **Compliance profiles** that support deployment across baseline, enterprise, and regulated environments.

Implementations **SHOULD** support standardized classification of decision requests using **Payload Kind**, enabling consistent policy evaluation and routing behavior.

---

## 2.2 Non-Goals

The Protocol explicitly **does NOT define**:

- how to build or train machine learning models, including large language models (LLMs)
- how to implement retrieval-augmented generation (RAG) or other reasoning pipelines
- how to orchestrate autonomous agents or agent workflows
- how systems route or select models during decision construction
- any required user interface (UI) or user experience (UX) implementation
- a specific policy language, rule engine, or domain-specific language (DSL)

The Protocol governs **decision authorization and traceability**, not the internal mechanisms used to construct decisions.

R.E.T.N.A governs **decisions**; it does **not generate them**.

---

## 3. Terminology and Definitions (Normative)

For the purposes of this Protocol, the following definitions apply.

### Action Type

A classification of the type of action proposed by a decision, such as purchase, notification, device actuation, or access control modification.

### Actor

An entity participating in a governed decision process, including systems, agents, services, or human participants.

### **Admissibility Outcome**

The governance determination establishing whether a candidate decision satisfies the conditions required for governed consequence.

### **Anomaly Signal**

A derived signal indicating unusual or potentially unsafe conditions that may influence policy evaluation or risk classification.

### **Attestation**

Optional cryptographic proof verifying the integrity and authenticity of a **Decision Receipt**.

### **Authority Provenance**

The identifiable basis from which execution authority originates, including source, scope, delegation posture, and verification status where applicable.

### **Authorization Boundary**

The enforcement interface at which a proposed decision must obtain authorization before execution. In this Protocol, the Authorization Boundary is implemented as the **Governance Boundary**.

### **Compliance Profile**

A defined implementation posture specifying the minimum governance, persistence, admissibility, and execution requirements for protocol conformance.

### **Canonical Reason Code**

A standardized machine-readable code representing the outcome or justification of a governance decision.

### **Constraint**

An individual enforceable rule evaluated during policy evaluation (e.g., “spend ≤ \$50”, “require confirmation”, “deny hazardous action”).

### **Confidence Threshold**

A policy-defined level of certainty required for a decision to proceed without escalation.

### **Consequential Execution**

Execution capable of producing Governed Consequence within a target system or environment and therefore subject to governance enforcement

### **Context Capture**

The process of collecting relevant state, signals, and environmental information required for decision evaluation.

### **Decision Class**

A classification of decisions based on risk or operational domain.

### **Decision Construction**

The process of forming a candidate action or plan from inputs, signals, models, or rules.

### **Decision Execution**

The act of applying a decision to the environment, such as purchasing, unlocking, deleting, notifying, or actuating a system.

### **Decision Mode**

A configuration defining the operational posture of a system (e.g., fully autonomous, human-in-the-loop, or advisory).

### **Decision Receipt**

A structured artifact containing provenance, evidence references, policy evaluations, authorization outcomes, and associated reason codes for a governed decision.

## **Dispatch Trace**

A graph-like representation of the participating actors and the path taken to construct and authorize a governed decision.

## **Edge Policy**

A governance rule defining which actors or systems are authorized to issue or execute specific categories of decisions.

## **Evidence**

Inputs used in decision construction or policy evaluation, including state deltas, sensor signals, user commands, retrieved documents, and derived signals.

## **Evidence Reference**

A pointer or identifier referencing stored evidence used during decision construction or governance evaluation.

## **Execution Consequence**

The resulting state, effect, or external outcome produced following authorized consequential execution.

## **Execution Refusal**

The prevention of governed execution when required governance conditions remain unresolved, invalid, contradictory, inadmissible, or unverifiable.

## **Executor**

The actor responsible for carrying out an authorized decision.

## **Formation Context**

Governance-relevant context describing how a candidate decision was shaped prior to admissibility evaluation and Governance Boundary enforcement.

## **Formation Integrity**

The preservation and inspectability of the upstream shaping context through which a candidate decision was constructed prior to admissibility evaluation and execution.

## **Governance Boundary**

The mandatory interface between **Decision Construction** and **Decision Execution** at which the Protocol evaluates the decision.

## **Governance Boundary Integrity**

The condition under which governed execution remains blocked until required governance conditions are resolved and execution eligibility is established.

## **Governed Consequence**

An operational, legal, financial, physical, security, environmental, or materially consequential outcome requiring governance authorization prior to execution.

## **Governance Outcome**

The final authorization result produced by governance evaluation (e.g., ALLOW, DENY, ESCALATE, DEFER, DEGRADE).

## **Governed State Transition**

A governance-resolved transition between prior state and resulting state that has satisfied authority, formation, admissibility, evidence, and execution integrity requirements prior to governed consequence.

## **Governor**

The actor responsible for evaluating policies and producing governance outcomes.

## **Interoperability**

The ability of independent systems to exchange and validate Decision Receipts and governance artifacts.

## **Payload Kind**

A classification identifying the semantic category of a decision request, used to route policy evaluation and governance behavior.

## **Policy**

A set of enforceable rules evaluated at the Governance Boundary.

## **Policy Authority**

The system or entity responsible for defining and distributing governance policies.

## **Policy Bundle**

A collection of policy rules, constraints, and configuration parameters evaluated during decision governance.

## **Policy Constraint**

A specific rule within a policy that restricts or conditions execution behavior.

## **Policy Evaluation**

The process of applying policies and constraints to a candidate decision.

## **Privileged Execution**

Execution of a decision that alters external state and therefore requires explicit authorization.

## **Producer**

The actor responsible for constructing a candidate decision.

## **Receipt Persistence**

The durable preservation of Decision Receipts as immutable governance artifacts sufficient for auditability and execution verification.

## **Receipt Store**

A system responsible for persisting Decision Receipts for audit and verification.

## Reason Code

A standardized machine-readable code representing the result of policy evaluation.

## Risk Class

A classification representing the potential impact of a decision, ranging from informational to safety-critical

---

## 4. Protocol Overview (Normative)

The R.E.T.N.A Protocol defines a mandatory governance workflow applied to all **governed decision classes** prior to Decision Execution.

The Protocol governs not only authorization visibility, but enforceable execution integrity across systems capable of operational, legal, financial, physical, security, or material consequence.

All candidate decisions **MUST** pass through the **Governance Boundary** prior to execution commitment. Implementations **MAY** optimize or parallelize internal processing steps, but the logical governance order defined by this Protocol **MUST** be preserved.

The canonical governance chain defined by the R.E.T.N.A Protocol is:

**authority → formation → evidence → admissibility → Governance Boundary → governed state transition → receipt persistence → execution consequence**

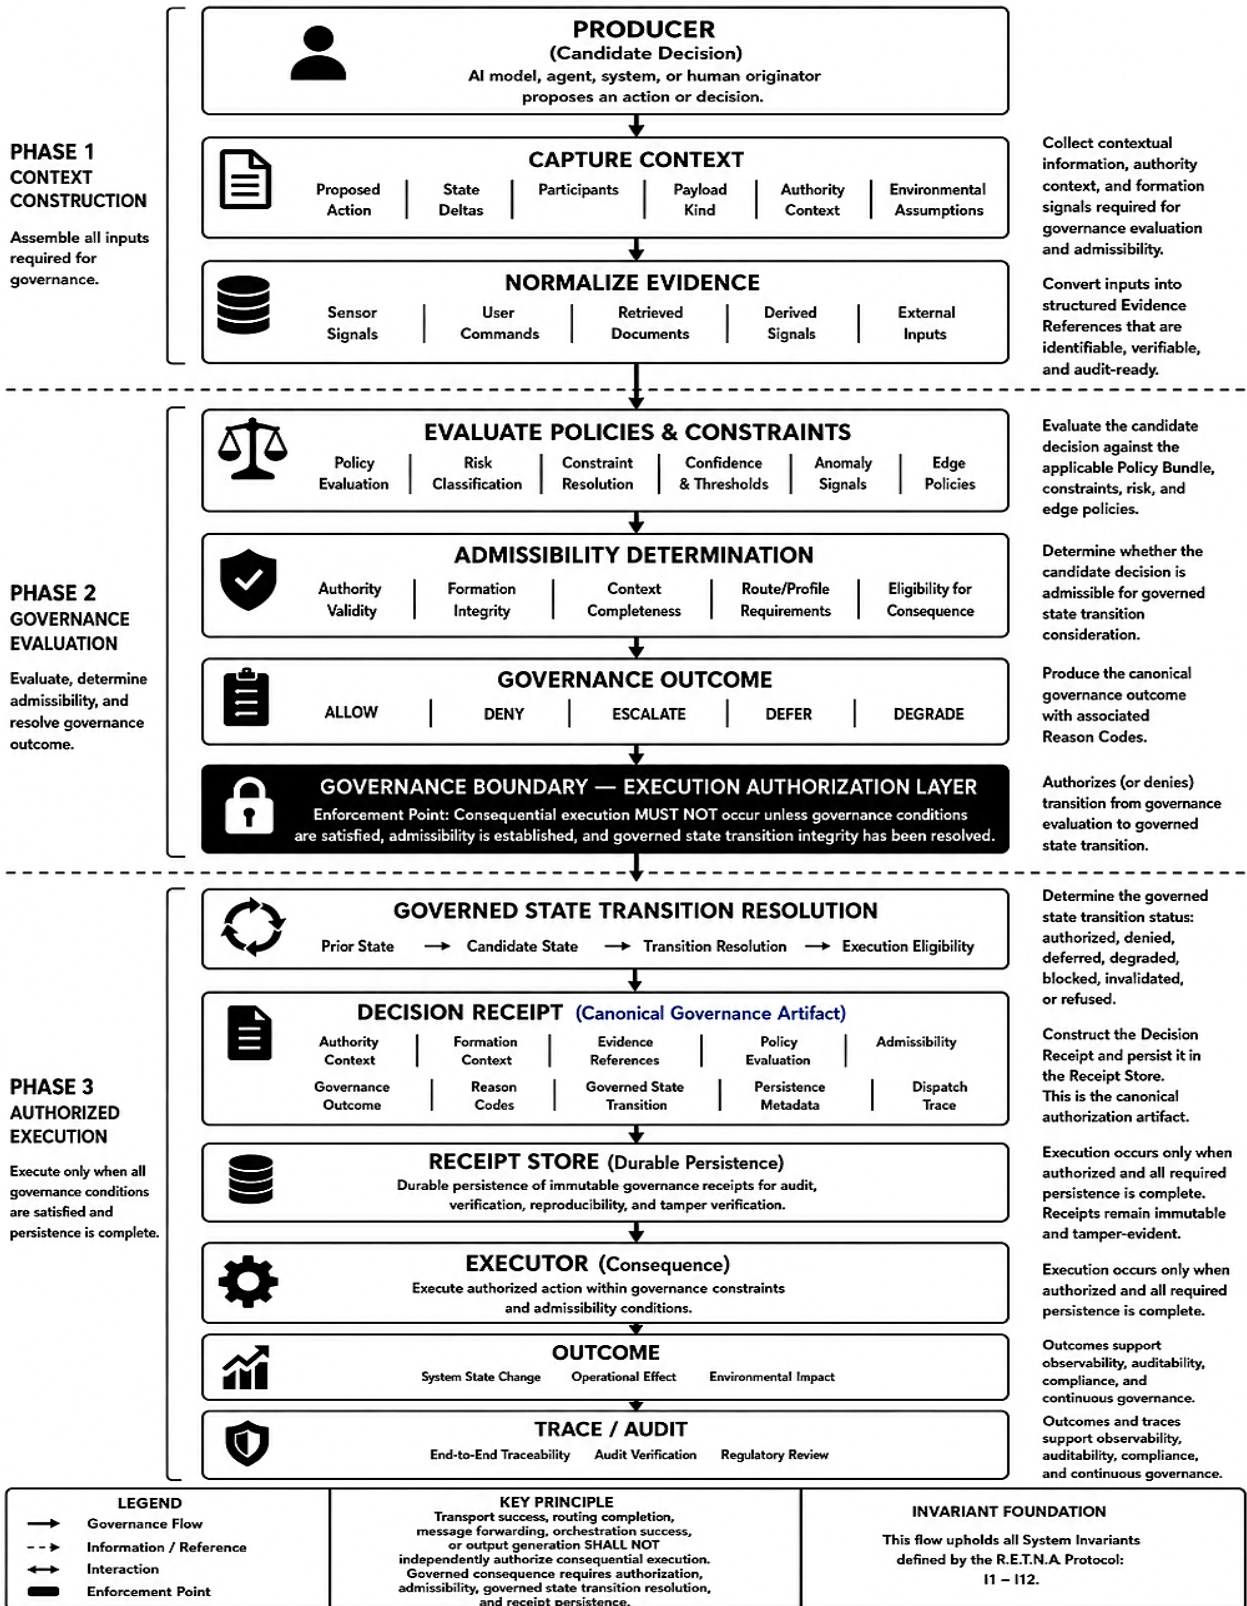
This sequence establishes that governed execution is not derived from transport success, output generation, orchestration completion, or routing continuity.

Governed consequence **MUST** instead result from explicit governance resolution occurring prior to execution commitment.

The Protocol therefore distinguishes:

- output generation from execution authorization
- forwarding from governed consequence
- transport completion from admissibility satisfaction
- orchestration continuity from governed state transition integrity

Implementations **MAY** refuse decision formation, invalidate inadmissible state transitions, prevent execution under unresolved governance conditions, or fail closed where **governance integrity** cannot be established.



**Figure 1. Governed Decision Flow in the R.E.T.N.A Protocol**

This figure illustrates the canonical governance workflow enforced by the R.E.T.N.A Protocol, from context construction and admissibility determination through Governance Boundary enforcement, Decision Receipt formation, governed state transition resolution, and authorized execution. Successful transport, routing, forwarding, or output generation does not independently authorize consequential execution under the Protocol.

---

## 4.0 Normative Flow Summary

The R.E.T.N.A Protocol governs authorization integrity and execution consequence through the following ordered workflow.

All compliant implementations **MUST** preserve this logical governance order even when internal processing steps are parallelized, distributed, cached, retried, or optimized.

The normative processing sequence defined by the Protocol is as follows:

### 1. Authority Resolution

The system resolves the authority context under which the candidate decision is being constructed.

Authority resolution **MAY** include:

- actor authorization validation
- Edge Policy evaluation
- privilege determination
- delegation verification
- execution scope validation
- governance profile applicability

Implementations **MUST NOT** authorize governed execution under unresolved, unverifiable, or invalid authority conditions.

---

### 2. Formation Construction

A Producer constructs a candidate decision request together with the contextual information required for governance evaluation.

Formation context **SHOULD** preserve sufficient information to explain how the candidate decision was shaped prior to admissibility evaluation and Governance Boundary enforcement.

Formation context **MAY** include:

- participating actors

- models involved
  - tools invoked
  - orchestration path
  - environmental assumptions
  - telemetry anchors
  - upstream governance influences
- 

### 3. Evidence Normalization

Captured inputs **MUST** be normalized into structured Evidence References suitable for policy evaluation, traceability, auditability, and reproducibility.

Evidence **MAY** include:

- sensor signals
- user commands
- retrieved documents
- derived signals
- external system inputs
- state snapshots
- anomaly signals

Each evidence artifact **MUST** be identifiable through a stable reference sufficient for verification and audit.

---

### 4. Admissibility Evaluation

The Governor evaluates the candidate decision against the applicable Policy Bundle, classification requirements, constraints, authority conditions, risk posture, and governance rules.

Admissibility evaluation determines whether the proposed decision satisfies the conditions required for governed execution.

Evaluation **MAY** include:

- Policy Constraints
- Risk Class classification
- Confidence Thresholds
- anomaly conditions
- applicable Edge Policy rules
- authority validation
- evidence sufficiency

- formation integrity verification
- 

## 5. Governance Outcome Resolution

Following admissibility evaluation, the Governor **MUST** produce a Governance Outcome.

The Protocol defines the following canonical outcomes:

- **ALLOW** — the decision is authorized for governed execution
- **DENY** — the decision is prohibited and **MUST NOT** execute
- **ESCALATE** — the decision requires higher-authority or human review
- **DEFER** — the decision is postponed pending additional evidence, timing conditions, or governance resolution
- **DEGRADE** — the decision **MAY** proceed only under reduced autonomy or constrained execution conditions

The selected outcome **MUST** be accompanied by one or more Reason Codes describing the basis for the governance determination.

---

## 6. Governed State Transition Resolution

For any governed action capable of operational, legal, financial, physical, security, or material consequence, the system **MUST** resolve a governed state transition prior to execution commitment.

Governed state transition resolution **MUST** distinguish:

- prior state
- candidate state
- resulting state
- governance outcome
- authorization basis
- admissibility status
- evidence references
- execution permissions
- applicable Reason Codes

Successful forwarding, routing, output generation, or transport completion **SHALL NOT** independently constitute governed execution success.

Execution **MUST NOT** proceed where governed state transition integrity cannot be established.

---

## 7. Decision Receipt Construction and Persistence

The system **MUST** construct a Decision Receipt as defined by this Protocol.

All governed decisions **MUST** be represented by a corresponding Decision Receipt prior to execution.

The Decision Receipt **MUST** contain:

- the candidate decision request
- Evidence References
- policy evaluation results
- authority and governance context
- governed state transition metadata
- the Governance Outcome
- associated Reason Codes
- the Dispatch Trace describing participating actors

The Decision Receipt **MUST** be persisted within a Receipt Store prior to consequential execution.

If receipt persistence fails, governed execution **MUST NOT** proceed.

---

## 8. Authorized Execution Consequence

If the Governance Outcome authorizes execution, the Executor **MAY** perform the approved action in accordance with the authorization context recorded in the associated Decision Receipt.

The Executor **MUST NOT** perform privileged actions unless:

- a valid Decision Receipt exists
- governance authorization has been resolved
- governed state transition integrity has been satisfied
- execution authority remains valid at execution time

Execution results **MAY** produce additional evidence, state deltas, environmental changes, telemetry updates, or follow-on decisions subject to subsequent governance evaluation.

---

## Phase 1 — Context Construction

### 4.1 Capture Context

The **Producer** collects the candidate decision request together with the contextual information required for governance evaluation, admissibility determination, and governed state transition resolution.

Captured context **MUST** include:

- the proposed action or decision request
- relevant **state deltas** describing proposed or expected consequence
- participating **Actors** involved in decision construction
- the **Payload Kind** associated with the request
- **authority context** sufficient for governance evaluation
- **environmental assumptions** relevant to execution integrity

Implementations **SHOULD** preserve sufficient formation integrity to explain how the candidate decision was shaped prior to Governance Boundary enforcement.

This context forms the initial governance input evaluated by the Protocol.

**Invariant Mapping:** I3, I11

---

### 4.2 Normalize Evidence

Inputs collected during context capture **MUST** be converted into structured **Evidence References** suitable for policy evaluation, traceability, and audit.

Evidence **MAY** include:

- sensor signals
- user commands
- retrieved documents
- derived signals
- external system inputs

Each evidence artifact **MUST** be identifiable through a stable reference to enable verification and audit.

**Invariant Mapping:** I4, I6

## Phase 2 — Governance Evaluation

### 4.3 Evaluate Policies and Constraints

The **Governor** evaluates the candidate decision against the applicable **Policy Bundle**, including:

- **Policy Constraints**
- **Risk Class** classification
- **Confidence Thresholds**
- **Anomaly Signals**
- applicable **Edge Policy** rules

Governance evaluation determines whether the candidate decision is admissible for governed consequence, requires escalation, must be degraded, deferred, denied, or may proceed toward governed execution authorization.

**Invariant Mapping:** I2, I9, I10

---

### 4.4 Emit Governance Outcome

Following governance evaluation and admissibility determination, the Governor **MUST** produce a **Governance Outcome**.

The Protocol defines the following canonical outcomes:

- **ALLOW** — the decision is authorized for execution
- **DENY** — the decision is prohibited and **MUST NOT** execute
- **ESCALATE** — the decision requires human or higher-authority review
- **DEFER** — the decision is postponed pending additional evidence or time conditions
- **DEGRADE** — the decision may proceed only under reduced autonomy or safer fallback behavior

The selected outcome **MUST** be accompanied by one or more **Reason Codes** describing the basis for the decision.

**Invariant Mapping:** I2, I10

---

### Phase 3 — Governed Execution Authorization

## 4.5 Resolve Governed State Transition

Following Governance Outcome determination, the system **MUST** resolve whether the candidate decision is eligible to produce governed consequence.

Governed state transition resolution **determines whether the proposed action may transition from candidate state into consequential execution under the applicable authority, admissibility, policy, persistence, and execution constraints.**

Governed state transition resolution **MAY** result in:

- Authorized
- Denied
- Deferred
- Degraded
- Blocked
- Invalidated
- Refused

Successful forwarding, routing, orchestration completion, transport success, or output generation **MUST NOT** independently constitute governed execution authorization.

Governed consequence **MUST** result only from explicit governed state transition resolution.

**Invariant Mapping:** I10, I12

---

## 4.6 Emit Decision Receipt

Following governed state transition resolution, the system **MUST** construct a **Decision Receipt** as defined by this Protocol.

The Decision Receipt is the canonical governance artifact defined by the R.E.T.N.A Protocol.

All governed decisions **MUST** be represented by a corresponding Decision Receipt prior to execution.

The Decision Receipt **MUST** contain the following fields:

- the candidate decision request
- **authority context**, where materially applicable

- **formation context**, where required
- **Evidence References**
- policy evaluation results
- admissibility determination
- **Governance Outcome**
- associated **Reason Codes**
- **governed state transition metadata**, where consequential execution is involved
- the **Dispatch Trace** describing materially participating actors
- **receipt persistence** and **integrity metadata**

The Decision Receipt **MUST** be persisted within a **Receipt Store** and **MAY** be returned to requesting systems for audit, verification, or interoperability.

The Decision Receipt serves as the **canonical governance authorization artifact** enabling governed execution by the Executor.

Governed consequence **MUST NOT** occur unless a valid Decision Receipt has been emitted, persisted, and associated with the governed state transition being authorized.

---

## 4.7 Persist Receipt and Execute Authorized Decision

Where governed consequence is permitted, the associated Decision Receipt **MUST** be persisted prior to consequential execution unless otherwise permitted by the applicable compliance profile.

If the Governance Outcome permits execution (for example **ALLOW** or **DEGRADE**), and required authority, formation integrity, admissibility, governed state transition integrity, and persistence conditions remain satisfied, the authorized decision **MAY** proceed toward governed execution by the Executor.

The Executor is responsible for carrying out the approved action within the target system or environment.

Execution **MUST** conform to:

- the **authorization context recorded in the Decision Receipt**
- applicable governance constraints
- admissibility conditions
- **governed state transition** requirements
- reduced-autonomy restrictions where applicable

The Executor **MUST NOT** perform privileged actions unless:

- a valid Decision Receipt exists
- the Governance Outcome authorizes execution
- governed state transition integrity has been resolved
- required receipt persistence has completed successfully

Successful forwarding, routing, orchestration completion, output generation, or transport success **MUST NOT** independently authorize consequential execution.

Execution results **MAY** produce additional evidence, state deltas, telemetry changes, environmental consequence, or follow-on decisions subject to subsequent governance evaluation under this Protocol.

**Invariant Mapping:** I1, I2, I10, I12

---

## 5. System Invariants (Normative)

The following **System Invariants** define mandatory properties that any implementation claiming conformance with the R.E.T.N.A Protocol **MUST** uphold.

These invariants are not implementation guidelines. They represent **protocol-level guarantees** that ensure enforceability, auditability, and interoperability across systems, vendors, and deployment environments.

Violation of any invariant **invalidates claims of R.E.T.N.A compliance** for the affected decision classes.

---

## I1 — Governance Boundary Integrity

For any decision belonging to a governed decision class, **no execution MAY occur unless the decision has passed through the R.E.T.N.A Governance Boundary.**

- The **Governance Boundary MUST be positioned between Decision Construction and Decision Execution.**
- The **Executor MUST NOT accept or act upon a decision unless authorized by a valid Governance Outcome.**
- Any attempt to bypass the Governance Boundary **MUST result in denial of execution and SHOULD be recorded using a canonical Reason Code indicating a policy violation.**

**Rationale:**

Without an enforced boundary, governance degenerates into advisory logging and cannot provide trust guarantees.

---

## I2 — Mandatory Receipt Emission

Every governed decision **MUST produce a Decision Receipt**, regardless of governance outcome.

This includes decisions that are:

- denied
- deferred
- escalated
- degraded
- or never executed

The absence of a Decision Receipt for a governed decision **constitutes a protocol violation**.

**Rationale:**

Accountability requires visibility into both actions taken and actions prevented.

---

## I3 — Trace Completeness

Decision Receipts **MUST identify all materially participating entities** involved in the decision pathway.

At minimum, the trace **MUST include**:

- **Producers** (agents, orchestrators, planners)
- **Governors** (R.E.T.N.A evaluators)
- **Executors** (if execution was attempted)
- **Models and tools materially involved in Decision Construction**

Entities **MUST** be referenced via stable identifiers sufficient to support correlation and audit.

The set of participating entities **MAY** be represented through the **Dispatch Trace**.

**Rationale:**

Partial traces undermine provenance and frustrate post-incident analysis.

---

## I4 — Evidence Referencing

Decision Receipts **MUST contain either:**

- structured evidence summaries, or
- references to externally stored **Evidence Artifacts**

Where references are used, they **MUST** be sufficiently specific to allow retrieval and verification under the applicable compliance profile.

Receipts **MUST NOT require embedding of raw sensitive data where references suffice.**

**Rationale:**

Governance must be auditable without violating privacy or creating unnecessary data exposure.

---

## I5 — Policy Disclosure (Structured)

Decision Receipts **MUST record:**

- which **Policy Bundle(s)** were applied
- the **version(s)** of those bundles
- which **Policy Constraints** were evaluated
- which constraints **materially influenced the Governance Outcome**

Policy disclosure **MUST be structured and machine-readable.**

**Rationale:**

Trust requires not only knowing what happened, but which rules governed it.

---

## I6 — Outcome Determinism

Given identical **Payload Kind**, inputs, evidence references, policy bundles, and configuration state, governance outcomes **SHOULD be reproducible within declared bounds of nondeterminism.**

If nondeterminism is present, it **MUST be explicitly documented and traceable** through Decision Receipts.

Associated **Reason Codes MUST remain consistent** with the reproduced outcome.

**Rationale:**

Determinism (or declared nondeterminism) is required for debugging, auditing, and regulatory review.

---

## I7 — Tamper Awareness

Decision Receipts **MUST include integrity mechanisms appropriate to the declared compliance profile.**

The Protocol defines the following integrity expectations:

- **RETNA-A:** basic integrity awareness (e.g., append-only logging recommended)
- **RETNA-B:** tamper-evident mechanisms **REQUIRED**
- **RETNA-C:** cryptographic attestation or strongly equivalent integrity guarantees **REQUIRED**

Any detected integrity failure **MUST be recorded and flagged.**

**Rationale:**

Receipts that can be altered without detection cannot serve as trust artifacts.

---

## I8 — Separation of Duties

Decision Construction and Decision Governance **SHOULD be separable functions**, even if implemented within the same runtime or service.

Implementations **MUST NOT require governance logic to be embedded within model inference or reasoning processes.**

**Rationale:**

Separation of duties enables independent validation, reduces conflicts of interest, and supports multi-vendor interoperability.

---

## I9 — Authority Integrity

A governed decision **MUST NOT** proceed toward privileged execution unless the authority basis for forming, authorizing, routing, or forwarding the candidate decision is explicit, inspectable, and bound to the associated governance artifact.

Authority integrity **MUST** preserve sufficient information to determine:

- which actor or system asserted execution authority
- the scope under which authority was granted
- **applicable Edge Policy constraints**
- whether delegated authority remained **valid at evaluation time**
- whether execution authorization remained **valid at execution commitment time**

Authority conditions that are unresolved, unverifiable, contradictory, expired, or inadmissible **MUST** prevent governed execution.

Where authority verification materially influences governance evaluation, the resulting authorization basis **MUST** be represented within the Decision Receipt.

### **Rationale:**

Governed execution requires verifiable authority provenance rather than implicit trust derived from transport position, orchestration topology, or system adjacency.

---

## I10 — Admissibility Requirement

A governed decision **MUST NOT** execute unless its admissibility outcome is explicitly evaluated, recorded, and represented through canonical governance metadata.

Admissibility evaluation **MUST** determine whether the candidate decision satisfies the governance conditions required for consequential execution under the applicable Policy Bundle, Risk Class, Decision Mode, authority conditions, and evidence requirements.

Admissibility determination **MAY** include evaluation of:

- policy constraints
- authority validity
- evidence sufficiency
- formation integrity
- execution scope
- risk posture
- Edge Policy restrictions
- anomaly conditions

Where admissibility cannot be established, governed execution **MUST** fail closed.

Admissibility outcomes **MUST** be traceable through the associated Decision Receipt and corresponding Reason Codes.

**Rationale:**

Execution integrity requires explicit admissibility determination prior to consequential state transition authorization.

---

## I11 — Formation Integrity

A governed decision **SHOULD** preserve the formation context of the candidate decision when formation is materially relevant, route-required, policy-sensitive, or necessary to explain how the candidate decision was shaped prior to Governance Boundary enforcement.

Where formation metadata is required, a governed decision **MUST NOT** proceed without formation context sufficient to identify:

- the upstream shaping path
- materially participating actors
- relevant orchestration influences
- applied tools or models where applicable
- governing constraints influencing proposal formation
- the basis upon which the candidate decision was constructed

Formation integrity **MAY** be represented through Dispatch Trace structures, formation metadata blocks, orchestration lineage references, or equivalent governance artifacts.

Formation preservation requirements **MAY** vary according to Compliance Profile, Risk Class, Payload Kind, or domain-specific governance policy.

**Rationale:**

Governance cannot fully evaluate consequential execution without preserving sufficient visibility into how materially relevant candidate decisions were shaped prior to admissibility evaluation.

---

## I12 — Governed State Transition Integrity

A governed action **MUST NOT** be treated as complete merely because it was forwarded, routed, delivered, acknowledged, persisted, emitted as output, or successfully transported between systems.

Governed actions capable of operational, legal, financial, physical, security, or material consequence **MUST** resolve into an **explicit governed state transition prior to execution commitment**.

Governed state transition integrity **MUST** preserve sufficient information to identify:

- prior state
- candidate state
- resulting state
- governance outcome
- authority basis
- formation context
- admissibility outcome
- evidence references
- receipt reference
- receipt persistence status
- execution authorization conditions
- applicable canonical Reason Codes

Successful forwarding, orchestration completion, output generation, transport success, or message delivery **SHALL NOT** independently constitute governed execution success.

Governed consequence **MUST** instead derive from explicit governance resolution resulting in an admissible and authorized governed state transition.

Execution **MUST NOT** proceed where governed state transition integrity cannot be established.

**Rationale:**

The Protocol governs consequential state transition authorization rather than transport semantics, orchestration continuity, or output propagation behavior.

---

## 5.1 Invariant Enforcement Scope

System Invariants apply to:

- all decisions within governed decision classes
- all compliance profiles (**RETNA-A, RETNA-B, RETNA-C**)
- all deployment environments (cloud, edge, embedded, hybrid)

Implementations **MAY apply stricter guarantees**, but **MUST NOT weaken or bypass these invariants**.

## 5.2 Invariant Violations

If any invariant is violated:

- the affected decision **MUST be treated as non-compliant**
- the system **MUST NOT claim R.E.T.N.A compliance** for that decision
- remediation **SHOULD be documented and auditable**

---

### Transition Note (Non-Normative)

With the invariants defined, the following sections specify how systems fulfill these guarantees through roles, artifacts, and protocol flows.

System Invariants are presented before Actors and Roles intentionally, because the invariants define the guarantees that the subsequent roles and protocol mechanisms must satisfy.

---

## 6. Actors and Roles (Normative)

This section defines the actors recognized by the **R.E.T.N.A Protocol** and the responsibilities assigned to each.

Actors represent **functional roles**, not required deployment units. A single system component **MAY implement multiple roles**, provided that all applicable **System Invariants** are upheld.

For **governed decision classes**, any interaction that may result in execution **MUST pass through the Governance Boundary**.

---

### 6.1 Required Actors

The following actors are **REQUIRED** for any system claiming **R.E.T.N.A compliance**.

#### 6.1.1 Producer

The **Producer** is any component that constructs and proposes a decision for potential execution.

Examples include:

- agent orchestrators
- planners
- workflow engines
- model routers
- task schedulers

## Responsibilities

The Producer **MUST**:

- construct a **Decision Proposal** containing sufficient context, evidence references, and intent
- submit the Decision Proposal to the **Governance Boundary**
- await an explicit **Governance Outcome** prior to execution

## Constraints

A Producer **MUST NOT**:

- execute a governed decision directly
- bypass or suppress governance evaluation
- forge or modify **Decision Receipts**

**Invariant Mapping:** I1, I2, I3, I8

---

### 6.1.2 Governor (R.E.T.N.A Implementation)

The **Governor** is the component that implements **R.E.T.N.A governance logic** and enforces the **Governance Boundary**.

The Governor **is the authoritative decision authority for governed decision classes**.

The Governor **MAY** be implemented as:

- a standalone service
- an embedded library
- a policy enforcement point within an execution gateway

## Responsibilities

The Governor **MUST**:

- receive **Decision Proposals** from Producers
- evaluate proposals against policies, constraints, and evidence
- determine a deterministic **Governance Outcome**
- emit a **Decision Receipt** for every evaluated proposal
- communicate authorization (or denial) to the **Executor**

### Constraints

The Governor **MUST**:

- be authoritative over execution authorization
- record all evaluated decisions, including denied or deferred proposals
- apply policies consistently with the declared **compliance profile**

**Invariant Mapping:** I1, I2, I5, I6, I7

---

### 6.1.3 Executor

The **Executor** is any component capable of applying a decision to the environment.

Examples include:

- payment processors
- IoT actuators
- database mutation services
- notification dispatchers
- access control systems

### Responsibilities

The Executor **MUST**:

- verify governance authorization prior to execution
- validate the authenticity and integrity of the **Decision Receipt** where applicable
- execute only decisions explicitly authorized by the **Governor**
- report execution status where applicable

## Constraints

An Executor **MUST NOT**:

- execute a governed decision without authorization
- treat missing authorization as recoverable
- fabricate or alter **Decision Receipts**

Lack of authorization **MUST be treated as a hard failure.**

**Invariant Mapping:** I1, I2, I7

---

## 6.1.4 Receipt Store

The **Receipt Store** is a durable storage system responsible for persisting **Decision Receipts**.

Examples include:

- databases
- append-only logs
- audit ledgers
- compliance data stores

## Responsibilities

The Receipt Store **MUST**:

- persist all **Decision Receipts** emitted by the Governor
- preserve receipt integrity according to the declared **compliance profile**
- support retrieval and export for audit and investigation

## Constraints

The Receipt Store **MUST NOT permit undetectable modification of stored receipts.**

Receipt storage **MUST provide tamper-aware durability consistent with the declared compliance profile.**

Receipt deletion or alteration **MUST be governed by policy.**

**Invariant Mapping:** I2, I7

---

## 6.2 Optional Actors

The following actors are **OPTIONAL but RECOMMENDED** for enterprise and regulated deployments.

---

### 6.2.1 Policy Authority

The **Policy Authority** defines, versions, and publishes **Policy Bundles** evaluated by the Governor.

#### Responsibilities

The Policy Authority **MUST**:

- publish identifiable, versioned policy bundles
- ensure policy immutability once published
- provide policy provenance and ownership metadata

#### Constraints

Policy bundles **SHOULD be independently auditable**.

Policy updates **SHOULD follow change-management practices**.

**Invariant Mapping:** I5, I6

---

### 6.2.2 Evidence Store

The **Evidence Store** holds raw or derived **Evidence Artifacts** referenced by **Decision Receipts**.

Examples include:

- document stores
- sensor data repositories
- feature stores

- snapshot archives

### Responsibilities

The Evidence Store **MUST**:

- store evidence artifacts referenced by receipts
- provide integrity references (hashes or immutable identifiers)
- enforce access and retention policies

### Constraints

Evidence Stores **SHOULD minimize exposure of sensitive data**.

Evidence retrieval **MUST be auditable where required**.

**Invariant Mapping:** I4, I7

---

## 6.2.3 Attestor

The **Attestor** provides cryptographic or strongly equivalent integrity guarantees for **Decision Receipts**.

### Responsibilities

The Attestor **MUST**:

- sign or otherwise attest to receipt integrity
- support verification by third parties where required

### Constraints

Attestation mechanisms **MUST align with the declared compliance profile**.

Key management **MUST follow established security best practices**.

**Invariant Mapping:** I7

---

## 6.3 Role Composition and Separation

A single component **MAY implement multiple roles** (for example, **Governor + Receipt Store**) provided that all applicable **System Invariants** are preserved.

Implementations **SHOULD maintain logical separation between:**

- **Decision Construction (Producer)**
- **Decision Governance (Governor)**
- **Decision Execution (Executor)**

### Rationale

Logical separation:

- enables independent validation
  - reduces conflicts of interest
  - supports multi-vendor interoperability
- 

## 6.4 Role Identification and Traceability

All actors referenced in a **Decision Receipt MUST be identifiable via stable identifiers** sufficient to support:

- correlation across receipts
- audit and investigation
- version and provenance tracking

Identifiers **MAY include:**

- service identifiers
  - component names
  - version identifiers
  - deployment environment identifiers
-

## 7. Decision Classification (Normative)

This section defines the mandatory **classification schema** applied to all governed decisions under the **R.E.T.N.A Protocol**.

Decision Classification establishes the structural context in which governance evaluation occurs. Before policies are evaluated, **the Governor MUST determine the decision's Risk Class, Action Type, and Decision Mode**. These classification attributes define the severity of potential outcomes, the operational authority required for execution, and the escalation pathways available to the system.

Decision classification additionally establishes the governance sensitivity associated with **consequential state transitions, authority requirements, admissibility enforcement, and formation integrity obligations**.

Classification therefore informs not only policy strictness, but also the level of governance preservation required prior to execution commitment.

Decision classification establishes:

- the **risk envelope** of a proposed action
- the **execution modality** permitted
- the **policy strictness** required at the Governance Boundary

All **R.E.T.N.A-compliant implementations MUST classify each governed decision prior to policy evaluation**.

Classification **MUST remain stable during policy evaluation** and **MUST be recorded in the Decision Receipt**.

---

### 7.1 Classification Requirements

For every governed decision, the **Governor MUST assign the following classification fields**:

1. **Risk Class** (required)
2. **Action Type** (required)
3. **Decision Mode** (required)

These fields are **normative inputs to policy evaluation**.

All classification fields **MUST be included in the Decision Receipt**.

## 7.2 Risk Class (Required)

The **Risk Class** expresses the **maximum plausible impact** of executing the proposed decision, assuming worst-case conditions.

Risk classification **MUST be conservative** and **MUST NOT be downgraded to avoid policy enforcement**.

Risk classification **SHOULD** additionally consider:

- irreversible consequence potential
- authority sensitivity
- admissibility strictness requirements
- governed state transition impact
- execution recoverability limitations

Risk classification **MUST consider the worst-case consequences of successful execution**.

### 7.2.1 Baseline Risk Classes

A governed decision **MUST be assigned exactly one** of the following risk classes.

Risk Class	Identifier	Description
Informational	R0_INFO	No direct external effect; advisory or observational only
Low Impact	R1_LOW	Reversible, non-financial, minimal consequence
Financial	R2_FINANCIAL	Commits funds, inventory, or contractual resources
Physical	R3_PHYSICAL	Actuates physical devices or controls environments
Safety-Critical	R4_SAFETY	Potential harm to persons, property, or regulated domains

### 7.2.2 Risk Assignment Rules

- Risk Class **MUST reflect potential impact**, not intent.
- Risk Class **MUST be assigned prior to policy evaluation**.
- Risk escalation **MAY occur during evaluation** if additional evidence indicates higher risk.

- Risk de-escalation **MUST be justified and recorded** in the Decision Receipt.

**Invariant Mapping:** I1, I6

---

## 7.3 Action Type (Required)

The **Action Type** identifies the semantic category of the proposed operation.

Action Types enable:

- policy reuse across systems
- regulator interpretation
- cross-domain analytics

Certain Action Types **MAY** inherently require governed state transition validation prior to execution.

Implementations **SHOULD** treat Action Types involving operational, financial, physical, legal, security, or irreversible consequence as governance-sensitive execution categories.

---

### 7.3.1 Baseline Action Types

The following Action Types are **RECOMMENDED baseline categories**.

- NOTIFY
- REORDER
- DISCARD
- UNLOCK
- LOCK
- SHUT\_OFF
- START
- SCHEDULE
- UPDATE\_RECORD
- DELETE\_RECORD
- AUTHORIZE\_ACCESS
- CALL\_EMERGENCY

Implementations **MAY introduce domain-specific Action Types**, provided they:

- preserve semantic clarity
- do not overload baseline meanings
- are recorded verbatim in the Decision Receipt

---

## 7.4 Decision Mode (Required)

The **Decision Mode** specifies the **level of autonomy permitted for execution**.

Decision Mode is an explicit governance control and **MUST be honored by the Executor**.

---

### 7.4.1 Decision Modes

A governed decision **MUST be classified as exactly one** of the following modes.

Mode	Identifier	Description
Autonomous	AUTONOMOUS	System executes without human confirmation
Assisted	ASSISTED	Human confirmation required prior to execution
Manual	MANUAL	System may recommend but <b>MUST NOT</b> execute

---

### 7.4.2 Decision Mode Constraints

- Decision Mode **MUST be compatible with Risk Class**.
- Higher Risk Classes **SHOULD default to reduced autonomy**.
- Decision Mode changes **MUST be recorded as governance outcomes** (for example via **DEGRADE**).

**Invariant Mapping:** I1, I8

### 7.4.3 Formation and Transition Sensitivity

Certain governed decisions **MAY** require enhanced governance preservation requirements based on:

- Risk Class
- Action Type
- Payload Kind
- execution sensitivity
- consequence irreversibility
- regulatory requirements
- authority scope

Implementations **MAY** require enhanced:

- formation integrity preservation
- admissibility verification
- governed state transition validation
- evidence sufficiency guarantees
- receipt persistence guarantees

prior to execution authorization.

Governed decisions involving operational, legal, financial, physical, security, or material consequence **SHOULD** be treated as governed transition-sensitive decisions.

Where transition-sensitive decisions are identified, the resulting Decision Receipt **SHOULD** preserve sufficient governance metadata to support:

- authority reconstruction
- admissibility reconstruction
- formation inspection
- governed state transition verification
- audit replay

## 7.5 Classification Integrity and Disclosure

The following rules apply to all compliant implementations.

- Classification fields **MUST be included in every Decision Receipt.**
- Classification **MUST be machine-verifiable.**
- Classification changes across retries or escalations **MUST be traceable.**
- Classification logic **SHOULD be policy-driven or declarative.**

Classification **MUST NOT be:**

- inferred post-hoc
- overridden silently
- omitted for performance reasons

## 7.6 Relationship to Policy Evaluation

Decision Classification:

- **precedes policy evaluation**
- **informs constraint severity**
- **determines escalation thresholds**

Policies **MAY** reference:

- specific Risk Classes
- Action Types
- Decision Modes
- combinations thereof

This separation ensures that **classification remains stable while policy evolves independently**.

---

## 8. The Decision Receipt (Normative Core)

The **Decision Receipt** is the primary artifact defined by the **R.E.T.N.A Protocol**.

All systems claiming **R.E.T.N.A compliance MUST emit a Decision Receipt for every governed decision**.

The **Governor** is the authoritative issuer of Decision Receipts for governed decision classes.

The Decision Receipt provides:

- **verifiable provenance**
- **enforceable accountability**
- **durable auditability**

for decisions constructed by AI or agentic systems and applied to real-world state.

A Decision Receipt is the canonical record of:

- what decision was proposed
- which actors participated
- which evidence was used
- which policies and constraints were evaluated

- which governance outcome was produced
- whether execution occurred
- how integrity and retention were preserved

The Decision Receipt therefore serves as the Protocol's **portable unit of AI accountability**.

---

## 8.1 Decision Receipt Requirements

A Decision Receipt **MUST**:

1. be emitted for every governed decision, regardless of outcome
2. capture sufficient information to explain what was proposed, why it was evaluated, how the outcome was determined, and whether consequential execution was authorized or refused
3. be structured for both machine verification and human inspection
4. be persistable and retrievable according to the declared compliance profile

Decision Receipts **MUST be generated prior to execution for ALLOW outcomes and in lieu of execution for DENY, DEFER, ESCALATE, or DEGRADE outcomes.**

Decision Receipts **MUST** contain sufficient information to support:

- audit and investigation
- policy verification
- evidence traceability
- authority verification
- admissibility reconstruction
- governed state transition verification
- execution authorization verification
- interoperability across compliant systems

Decision Receipts **MUST** preserve sufficient governance context to explain:

- how the candidate decision was formed
- which authority conditions governed evaluation and execution eligibility
- why the decision was admissible, inadmissible, deferred, degraded, escalated, or denied
- what governed state transition was authorized, constrained, blocked, invalidated, or refused
- whether consequential execution was permitted, constrained, blocked, degraded, or invalidated
- which governance conditions materially influenced execution authorization

Decision Receipts **SHALL** function as governance authorization artifacts rather than passive audit records alone.

Where governed consequence is permitted under the Protocol, the associated Decision Receipt **MUST** preserve sufficient metadata to reconstruct:

- the authority basis for execution eligibility
- the admissibility basis for execution authorization
- materially relevant formation context where applicable
- the governed state transition associated with consequential execution
- the governance outcome and associated Reason Codes

Decision Receipts **MUST NOT** depend on access to proprietary model internals, unrestricted chain-of-thought, or vendor-specific reasoning artifacts in order to be interpreted, verified, or audited.

**Invariant Mapping:** I2, I3, I5, I7, I9, I10, I11, I12

---

## 8.2 Receipt Lifecycle

Every Decision Receipt **MUST progress through the following logical lifecycle states:**

1. **Proposed** — the decision has been submitted for governance evaluation
2. **Evaluated** — policies and constraints have been applied
3. **Resolved** — a Governance Outcome has been determined
4. **Executed** (*optional*) — the action has been applied
5. **Finalized** — the receipt has been sealed for audit and retained according to policy

Receipts **MAY**:

- be updated with execution status, or
- be linked to a follow-on execution receipt

provided that **trace continuity is preserved**.

If receipt state changes occur after initial issuance, those changes **MUST** preserve a verifiable relationship to the original receipt through explicit linkage or chaining.

A lifecycle transition **MUST NOT** erase prior state.

## 8.3 Canonical Receipt Structure (v0.1)

The following fields constitute the **minimum canonical structure** of a Decision Receipt.

Serialization format **is not mandated**.

Implementations **MUST define deterministic canonicalization rules** sufficient to ensure:

- stable hashing
- reproducible verification
- consistent attestation or signature behavior
- interoperable audit export across systems

Implementations **MAY** add additional fields, provided such additions do not alter the semantics of required canonical fields.

Canonical receipt ordering **SHOULD** preserve governance evaluation order where practical.

The canonical receipt structure reflects the governance chain defined by this Protocol, preserving the relationship between identity, authority, decision formation, policy evaluation, admissibility determination, governed state transition eligibility, governance outcome, and execution consequence.

Implementations **MAY** serialize fields differently provided canonical semantics and required relationships remain preserved.

---

### 8.3.1 Receipt Header

The **Receipt Header** defines the identity and protocol versioning metadata of the receipt.

Field	Requirement
receipt_id	REQUIRED — globally unique
version	REQUIRED — protocol version
created_at	REQUIRED — ISO 8601 timestamp
expires_at	OPTIONAL — for deferred decisions

Additional header metadata **MAY** be included, provided canonical verification remains stable.

### 8.3.2 Identity Block

The **Identity Block** identifies the primary protocol actors associated with the receipt.

<b>Actor</b>	<b>Fields</b>
Producer	id, type, version
Governor	id, type, version, profile
Executor	id, type, version ( <i>OPTIONAL if not executed</i> )

All actor identifiers **MUST** be stable enough to support:

- cross-receipt correlation
- audit and investigation
- provenance tracking

Identity records **MAY** include deployment or environment identifiers where necessary for audit.

---

### 8.3.3 Authority Context Block

The Authority Context Block records the authority basis under which the governed decision was evaluated for admissibility and execution eligibility.

Where authority materially influences governance evaluation or execution eligibility, the Decision Receipt **MUST** preserve sufficient authority metadata to support auditability, admissibility reconstruction, and execution verification.

The Authority Context Block **MAY** include:

<b>Field</b>	<b>Requirement</b>
authority_id	REQUIRED where applicable
authority_type	REQUIRED where applicable
authority_scope	REQUIRED where applicable
delegation_chain	OPTIONAL
edge_policy_result	OPTIONAL
verification_status	REQUIRED where applicable

Field	Requirement
authority_expiration	OPTIONAL

Authority metadata **MUST** remain consistent with the Governance Outcome represented by the receipt.

Ambiguous, unverifiable, contradictory, expired, delegated-beyond-scope, or otherwise unresolved authority conditions **MUST NOT** independently authorize governed execution.

**Invariant Mapping:** I9, I12

---

### 8.3.4 Decision Envelope

The **Decision Envelope** describes the governed decision itself.

Field	Requirement
decision_id	REQUIRED — stable across retries
risk_class	REQUIRED
action_type	REQUIRED
decision_mode	REQUIRED
payload_kind	REQUIRED where supported by the implementation
proposed_action	REQUIRED — structured summary
target	REQUIRED — affected entity

The Decision Envelope **MUST** describe the candidate action independently of any internal model reasoning strategy.

If a **Payload Kind** is assigned, it **MUST** reflect the semantic category of the governed request and **MUST** remain consistent with downstream policy evaluation.

---

#### 8.3.4.1 Payload Kind Semantics

Payload Kind identifies the semantic category of the governed request for routing, privilege determination, and policy applicability.

Payload Kind supplements the core classification triad of Risk Class, Action Type, and Decision Mode. It does not replace those fields and **MUST NOT be interpreted as a substitute for risk classification.**

Where supported by an implementation, Payload Kind **MUST be recorded in the Decision Envelope** and **MUST remain stable throughout policy evaluation for the evaluated request.**

Payload Kind **MAY** be used to:

- resolve policy applicability
- determine whether an edge or actor is authorized to carry a given request category
- distinguish governance artifacts, execution requests, evidence references, and observational payloads

Where a request belongs to a privileged category under local policy, ambiguous or missing Payload Kind resolution **MUST NOT authorize privileged execution.**

---

### 8.3.5 Formation Context Block

The Formation Context Block preserves governance-relevant information describing how the candidate decision was shaped prior to admissibility evaluation and Governance Boundary enforcement.

- Formation preservation requirements **MAY** vary according to:
  - Compliance Profile
  - Risk Class
  - Payload Kind
  - Action Type
  - domain-specific governance policy
  - transition sensitivity

The Formation Context Block **MAY** include:

Field	Requirement
formation_id	OPTIONAL
formation_trace_ref	OPTIONAL
orchestration_lineage	OPTIONAL
participating_models	OPTIONAL
participating_tools	OPTIONAL
governance_constraints_applied	OPTIONAL
formation_integrity_status	OPTIONAL

Formation preservation **MUST NOT** require unrestricted disclosure of proprietary reasoning internals or unrestricted chain-of-thought visibility.

Implementations **SHOULD** preserve only the governance-relevant formation context necessary for auditability, admissibility reconstruction, execution integrity verification, and governed state transition inspection.

---

### 8.3.6 State Context

The **State Context** records the before/after references and the contextual state relevant to governance evaluation.

Field	Requirement
state_snapshot_ref_before	OPTIONAL
state_snapshot_ref_after	OPTIONAL
state_delta_summary	REQUIRED
environment_context	REQUIRED ( <i>redactable</i> )

The State Context **MUST** provide sufficient information to explain what changed, or what was expected to change, if the governed action were executed.

Where raw state data is sensitive, implementations **SHOULD** prefer references and bounded summaries over embedded disclosure.

---

### 8.3.7 Participants

All materially participating entities involved in **Decision Construction** **MUST** be declared.

The **Participants Block** **MUST** support declaration of at least:

- agents\_involved
- models\_involved
- tools\_invoked
- external\_services

Each participant entry **MUST** include:

- a stable identifier
- a declared role or purpose
- a version, where applicable

The set of participants **MAY** also be represented through a **Dispatch Trace**, provided the canonical receipt remains interpretable independently.

**Invariant Mapping:** I3

---

### 8.3.8 Evidence Block

Each Decision Receipt **MUST reference the evidence used during Decision Construction and governance evaluation.**

Each evidence\_item **MUST** include:

Field	Requirement
evidence_id	REQUIRED
type	REQUIRED
source	REQUIRED
timestamp	REQUIRED
hash	REQUIRED for Profiles B/C
location_ref	REQUIRED ( <i>non-public</i> )
summary	REQUIRED ( <i>bounded</i> )

Evidence items **MUST** be sufficiently specific to support retrieval, integrity validation, and audit review under the applicable compliance profile.

Where possible, evidence references **SHOULD** be preferred over raw embedded data.

**Invariant Mapping:** I4, I7

### 8.3.9 Policy Evaluation Block

The **Policy Evaluation Block** records the governance rules applied to the decision.

The receipt **MUST disclose all Policy Bundles applied.**

Each policy\_bundle **MUST** include:

Field	Requirement
bundle_id	REQUIRED
bundle_version	REQUIRED
authority	REQUIRED

Each applied constraint **MUST** include:

- constraint\_id
- result (pass / fail / unknown)
- severity
- reason\_codes

Risk assessment fields **MAY** include:

- risk\_score
- risk\_factors

Where applicable, the Policy Evaluation Block **SHOULD** record **Edge Policy** results governing whether the source actor was authorized to produce or route the decision.

Policy evaluation records **MUST** remain structured and machine-readable.

**Invariant Mapping:** I5, I6

---

### 8.3.10 Admissibility Block

The Admissibility Block records the admissibility determination produced during governance evaluation prior to execution authorization.

Admissibility evaluation determines whether the candidate decision satisfied the governance conditions required for governed consequence and governed state transition eligibility under the applicable policy, authority, evidence, formation, classification, and risk conditions.

The Admissibility Block **MAY** include:

<b>Field</b>	<b>Requirement</b>
admissibility_status	REQUIRED
admissibility_reason_codes	REQUIRED
blocking_constraints	OPTIONAL
authority_validation_result	OPTIONAL
formation_integrity_result	OPTIONAL
evidence_sufficiency_result	OPTIONAL
execution_eligibility	OPTIONAL
admissibility_timestamp	REQUIRED

Admissibility outcomes **MUST** remain traceable and reconstructable from the Decision Receipt.

Where admissibility cannot be established, governed execution **MUST NOT** proceed.

Execution eligibility **MAY** differ from admissibility status where policy, authority, persistence, escalation, degradation, or governed state transition conditions constrain consequential execution after admissibility evaluation.

**Invariant Mapping:** I10, I12

---

### 8.3.11 Outcome Block

The **Outcome Block** records the final result of governance evaluation.

<b>Field</b>	<b>Requirement</b>
governance_outcome	REQUIRED
outcome_reason_codes	REQUIRED
required_human_confirmation	REQUIRED
fallback_path	OPTIONAL

Field	Requirement
execution_status	REQUIRED ( <i>executed / not_executed / pending</i> )
human_override	REQUIRED if applicable

The Governance Outcome **MUST** be one of the canonical protocol outcomes:

- ALLOW
- DENY
- ESCALATE
- DEFER
- DEGRADE

At least one **Reason Code MUST** accompany every outcome.

Where human review, override, or fallback behavior occurs, the receipt **MUST** record that fact explicitly.

---

### 8.3.12 Integrity and Audit Block

The **Integrity and Audit Block** records the integrity protections and retention controls associated with the receipt.

Field	Requirement
receipt_hash	REQUIRED
previous_receipt_hash	OPTIONAL
signature	REQUIRED for Profile C
retention_class	REQUIRED
access_class	REQUIRED

Implementations **MUST** ensure that integrity metadata is sufficient to detect unauthorized modification of a receipt.

Where receipt chaining is used, previous\_receipt\_hash **MAY** be used to preserve continuity across related governance events.

**Invariant Mapping: I7**

---

**8.3.13 Canonical Receipt Serialization**

Implementations **MUST** define deterministic canonicalization rules sufficient to ensure stable hashing, reproducible verification, and interoperable audit export.

For JSON serialization, a conforming canonicalization profile **MAY** use the following rules:

- UTF-8 encoding
- canonical field names as defined by this section
- lexicographic key ordering
- no insignificant whitespace

Where a receipt hash is recorded, the receipt hash **MUST be computed over the canonical representation of the receipt with the “integrity\_receipt\_hash” field blanked or excluded prior to hashing.**

Receipt serialization **MAY** vary by implementation, but validators and auditors **MUST** be able to determine which canonicalization profile was used.

If an implementation supports multiple canonicalization profiles, the applied profile **MUST be recorded in the receipt integrity metadata.**

---

**8.3.14 Canonical Receipt Example (Informative)**

The following example illustrates a canonical Decision Receipt structure aligned with this specification.

```
{
  "receipt_id": "retna_8c1d0f8f08f2408b85f0f3650eb4e4e1",
  "version": "0.1",
  "created_at": "2026-03-15T01:00:00Z",
  "expires_at": null,

  "identities": {
    "producer": {
      "id": "device_integration_agent",
      "type": "agent",
```

```
"version": "2.x"
},
"governor": {
  "id": "retna_governor",
  "type": "governor",
  "version": "0.1",
  "profile": "RETNA-A"
},
"executor": {
  "id": "central_ai_agent",
  "type": "agent",
  "version": "2.x"
}
},

"authority_context": {
  "authority_id": "policy_authority_001",
  "authority_type": "system",
  "authority_scope": "inventory_reorder",
  "delegation_chain": [],
  "edge_policy_result": "authorized",
  "verification_status": "verified",
  "authority_expiration": null
},

"decision_envelope": {
  "decision_id": "dec_123",
  "risk_class": "R2_FINANCIAL",
  "action_type": "REORDER",
  "decision_mode": "ASSISTED",
  "payload_kind": "execution_request",
  "proposed_action": {
    "summary": "reorder milk"
  },
  "target": {
    "id": "inventory_management_agent",
    "type": "agent"
  }
}
},

"state_context": {
  "state_snapshot_ref_before": null,
  "state_snapshot_ref_after": null,
  "state_delta_summary": {
    "item": "milk",
    "level": "low"
  }
}
```

```
},
"environment_context": {
  "estimated_cost_usd": 4.99
}
},

"participants": {
  "agents_involved": [],
  "models_involved": [],
  "tools_invoked": [],
  "external_services": []
},

"evidence": [],

"policy_evaluation": {
  "policy_bundles": [
    {
      "bundle_id": "RETNA_BASE",
      "bundle_version": "0.1",
      "authority": "VIS"
    }
  ],
  "constraints": []
},

"admissibility": {
  "admissibility_status": "ADMISSIBLE",
  "admissibility_reason_codes": [
    "RC_POLICY_PASS"
  ],
  "blocking_constraints": [],
  "authority_validation_result": "valid",
  "formation_integrity_result": "verified",
  "evidence_sufficiency_result": "sufficient",
  "execution_eligibility": "eligible",
  "admissibility_timestamp": "2026-03-15T01:00:00Z"
},

"outcome": {
  "governance_outcome": "ALLOW",
  "outcome_reason_codes": [
    "RC_POLICY_PASS"
  ],
  "required_human_confirmation": false,
  "fallback_path": null,
```

```
"execution_status": "authorized",
"human_override": null
},

"integrity": {
  "profile": "RETNA-A",
  "canonicalization": "json:sort_keys,separators=(',',':),utf8",
  "previous_receipt_hash": null,
  "receipt_hash": "<sha256>"
},

"dispatch_trace": {
  "trace_id": "trace_123",
  "path": [
    "device_integration_agent",
    "retna_governor",
    "central_ai_agent"
  ]
}
}
```

This example is illustrative. It does not mandate a transport envelope, storage layout, or implementation language.

---

## 8.4 Receipt Immutability and Tamper Awareness

Decision Receipts **MUST** be treated as append-only artifacts.

- Receipts **MUST** be tamper-evident.
- Any mutation **MUST** result in a new receipt or a chained reference.
- Silent modification **constitutes a protocol violation**.

Receipts **MUST NOT** be altered in place in a manner that invalidates prior auditability.

Where receipt amendment is necessary, the amended state **MUST** be represented through an additional receipt or verifiable chain reference.

---

## 8.5 Minimal Receipt (RETNA-A)

Implementations claiming **RETNA-A** compliance **MUST emit, at minimum**, the following fields:

- receipt\_id
- version
- created\_at
- producer.id
- governor.id
- decision\_id
- risk\_class
- action\_type
- decision\_mode
- proposed\_action summary
- state\_delta\_summary
- agents\_involved
- policy\_bundle identifiers
- governance\_outcome
- reason\_codes
- admissibility\_status

The following fields **MUST** be included **where materially applicable to governance evaluation or consequential execution**:

- **authority\_context**, where authority materially influences governance evaluation or the decision involves privileged execution
- **formation\_context**, where formation preservation is route-required, profile-required, risk-sensitive, or materially relevant to how the candidate decision was shaped
- **governed\_state**\_transition metadata, where the governed action is capable of operational, legal, financial, physical, security, or material consequence
- **persistence metadata**, sufficient to indicate receipt persistence status and persistence target where consequential execution depends upon receipt durability

Where supported by an implementation, **Payload Kind SHOULD** be included in the minimal receipt profile.

For consequential governed actions, the minimal **RETNA-A** receipt **MUST** preserve sufficient metadata to determine:

- who asserted or carried the authority context
- whether the candidate decision was admissible
- whether required formation context was preserved
- what governed state transition was authorized, denied, deferred, degraded, blocked, invalidated, or refused
- whether receipt persistence completed prior to consequential execution

RETNA-A is intended to preserve baseline governance accountability while maintaining low-friction adoption of the Protocol.

RETNA-A does **not** require unrestricted disclosure of proprietary model reasoning, unrestricted chain-of-thought, or raw sensitive evidence.

Implementations **MAY** satisfy minimum disclosure requirements through bounded summaries, stable references, governance-relevant metadata, or equivalent governance artifacts.

**Invariant Mapping:** I2, I3, I5, I7, I9, I10, I11, I12

---

### **Transition Note (Non-Normative)**

With the Decision Receipt defined, R.E.T.N.A establishes a portable, attestable unit of AI accountability.

Subsequent sections specify:

- reason code standardization
  - policy interfaces
  - evidence handling
  - interoperability
  - conformance testing
- 

## **9. Canonical Reason Codes (Normative)**

This section defines the baseline set of **canonical reason codes** used by the **R.E.T.N.A Protocol** to represent governance outcomes and execution-relevant governance semantics in a structured and interoperable manner.

Reason codes provide a standardized governance semantics layer **between policy evaluation, admissibility determination, governed state transition evaluation, and the interpretation of governance outcomes** by humans, systems, validators, and auditors.

All systems claiming **R.E.T.N.A compliance MUST emit canonical reason codes when producing Decision Receipts.**

Canonical reason codes **MAY** influence execution eligibility, remediation posture, escalation behavior, degradation posture, persistence handling, or governed state transition authorization within compliant implementations.

Reason codes enable:

- deterministic explanation of governance outcomes
- machine-verifiable audit trails
- interoperable interpretation across independent systems

Implementations **MAY extend the canonical reason code set**, but **MUST NOT alter the meaning or semantics of baseline codes defined by this Protocol.**

Reason codes are **normative artifacts** and form part of the **Decision Receipt Outcome Block.**

---

## 9.1 Purpose of Reason Codes

Canonical reason codes serve four primary purposes:

1. **Consistent interpretation of governance outcomes across systems**
2. **Support for automated audit, analytics, and alerting pipelines**
3. **Regulatory review without disclosure of proprietary model reasoning**
4. **Stable signals for downstream workflows**, including escalation, retry, fallback, or remediation

Reason codes are **not narrative explanations.**

They are machine-readable assertions that identify the **authority, formation, admissibility, policy, evidence, risk, persistence, system, or governance condition** that materially influenced the governance outcome or execution posture.

Reason codes therefore act as the **canonical interpretation interface** between **governance evaluation, execution posture, and external observability.**

## 9.2 Outcome Reason Codes (Baseline Set)

The following baseline reason codes are defined by the Protocol.

All **R.E.T.N.A-compliant implementations MUST support these codes**.

Additional reason codes **MAY be emitted in combination with these canonical codes**, provided canonical codes remain present when applicable.

### 9.2.1 Policy Evaluation Codes

These codes describe the outcome of policy constraint evaluation.

Code	Meaning
<b>RC_POLICY_PASS</b>	All applicable constraints evaluated successfully
<b>RC_POLICY_FAIL</b>	One or more blocking constraints failed
<b>RC_POLICY_PARTIAL</b>	Non-blocking constraints failed or warnings were issued

Policy evaluation codes **SHOULD accompany every Decision Receipt** unless a higher-level failure prevented evaluation.

### 9.2.2 Evidence and Confidence Codes

These codes describe conditions related to the **availability, freshness, or integrity of evidence inputs**.

Code	Meaning
<b>RC_INSUFFICIENT_EVIDENCE</b>	Required evidence was missing or incomplete
<b>RC_LOW_CONFIDENCE</b>	Confidence threshold required for action was not met

Code	Meaning
<b>RC_EVIDENCE_STALE</b>	Evidence used during evaluation was outdated or expired
<b>RC_EVIDENCE_INTEGRITY_FAIL</b>	Evidence integrity verification failed

Evidence codes **SHOULD be emitted when evidence deficiencies materially influenced the governance outcome.**

### 9.2.3 Risk and Safety Codes

These codes describe governance decisions influenced by **risk evaluation or safety constraints.**

Code	Meaning
<b>RC_RISK_THRESHOLD_EXCEEDED</b>	Computed risk exceeded allowed policy threshold
<b>RC_SAFETY_GUARDRAIL</b>	Safety constraint or guardrail was triggered
<b>RC_REGULATORY_RESTRICTION</b>	Regulatory or compliance rule prevented execution

Risk and safety codes **MUST be emitted when a governance decision was influenced by safety, compliance, or risk thresholds.**

### 9.2.4 System and Dependency Codes

These codes describe system conditions or environmental dependencies that influenced the outcome.

Code	Meaning
<b>RC_EXTERNAL_DEPENDENCY_FAILURE</b>	Required external system was unavailable or failed
<b>RC_RATE_LIMITED</b>	Execution deferred due to rate limiting or throughput constraints
<b>RC_SYSTEM_DEGRADED</b>	System health condition triggered reduced autonomy

System reason codes **SHOULD be emitted when governance decisions are influenced by operational system state.**

### 9.2.5 Governance Control Codes

These codes describe governance decisions related to **human oversight, escalation, or fallback behavior.**

Code	Meaning
<b>RC_REQUIRES_CONFIRMATION</b>	Human confirmation required before execution
<b>RC_ESCALATION_REQUIRED</b>	Escalation to a higher authority required
<b>RC_FALLBACK_APPLIED</b>	Safe fallback path selected
<b>RC_HUMAN_OVERRIDE</b>	Human operator overrode the governance outcome

Governance control codes **MUST be emitted when governance outcomes involve escalation, confirmation, or override behavior.**

### 9.2.6 Authority Integrity Codes

These codes describe governance conditions related to authority disclosure, authority validity, delegation posture, or execution authorization scope.

Code	Meaning
<b>RC_AUTHORITY_REQUIRED</b>	Required authority disclosure was missing
<b>RC_AUTHORITY_INVALID</b>	Authority failed structural or verification requirements
<b>RC_AUTHORITY_UNVERIFIED</b>	Authority could not be verified
<b>RC_AUTHORITY_SCOPE_VIOLATION</b>	Authority exceeded declared execution scope
<b>RC_AUTHORITY_EXPIRED</b>	Authority context expired before execution authorization
<b>RC_AUTHORITY_ESCALATED</b>	Higher authority review required

Authority reason codes **MUST** be emitted where authority materially influenced governance eligibility or execution authorization.

---

### 9.2.7 Admissibility and Formation Codes

These codes describe governance conditions related to admissibility determination, formation integrity, or candidate decision formation.

Code	Meaning
<b>RC_ADMISSIBILITY_REQUIRED</b>	Required admissibility disclosure missing
<b>RC_ADMISSIBILITY_FAIL</b>	Candidate decision failed admissibility evaluation
<b>RC_FORMATION_REQUIRED</b>	Required formation metadata missing
<b>RC_FORMATION_INVALID</b>	Formation metadata malformed or insufficient
<b>RC_FORMATION_CONTEXT_LOST</b>	Required proposal-shaping context unavailable

Admissibility and formation reason codes **SHOULD** be emitted where proposal permissibility, provenance, or shaping integrity materially influenced governance outcomes.

---

### 9.2.8 Governed State Transition Codes

These codes describe governance conditions related to the authorization, refusal, degradation, or invalidation of governed state transitions.

Governed state transition codes distinguish successful message forwarding from governed consequence authorization.

Successful transport, routing, output generation, or forwarding **SHALL NOT** independently imply governed execution success.

Governed consequence requires explicit governed state transition resolution.

Governed state transition authorization **MUST** remain traceable to a valid Decision Receipt and associated governance outcome.

Code	Meaning
<b>RC_STATE_TRANSITION_AUTHORIZED</b>	Candidate decision authorized for governed state transition
<b>RC_STATE_TRANSITION_DENIED</b>	Candidate decision denied state transition
<b>RC_STATE_TRANSITION_DEFERRED</b>	Candidate decision deferred pending additional governance resolution
<b>RC_STATE_TRANSITION_DEGRADED</b>	Candidate decision allowed only under degraded posture
<b>RC_STATE_TRANSITION_BLOCKED_NO_AUTHORITY</b>	Missing required authority
<b>RC_STATE_TRANSITION_BLOCKED_NO_FORMATION</b>	Missing or invalid formation metadata
<b>RC_STATE_TRANSITION_BLOCKED_NO_ADMISSIBILITY</b>	Missing or invalid admissibility
<b>RC_STATE_TRANSITION_BLOCKED_NO_RECEIPT</b>	Required Decision Receipt unavailable
<b>RC_STATE_TRANSITION_BLOCKED_PERSISTENCE_FAILURE</b>	Receipt persistence unavailable under fail-closed posture
<b>RC_STATE_TRANSITION_INVALID_RESULT</b>	Invalid resulting state for current governance posture

Governed state transition codes **MUST** be emitted where a governed action is capable of producing operational, legal, financial, physical, security, or material system consequence.

Successful transport, routing, forwarding, or output generation **SHALL NOT** independently imply governed execution authorization. Governed consequence requires explicit governed state transition resolution supported by a valid Decision Receipt and persistence posture consistent with applicable governance policy.

---

### 9.2.9 Continuity and State Drift Codes — Reserved Canonical Set

Continuity and state drift codes describe governance conditions where assumptions, authority posture, admissibility context, telemetry, environmental state, dependency state, or execution continuity materially changed between governance authorization and consequential execution opportunity.

The following continuity and state drift reason codes are reserved canonical codes for implementations that support continuity-sensitive governance profiles. Implementations that do

not yet enforce continuity validation **MUST NOT** emit these codes as active enforcement outcomes unless the corresponding runtime validation exists.

Continuity and state drift codes **MUST** be emitted where a governed decision relied upon assumptions, authority posture, telemetry, environmental state, admissibility context, or dependency conditions that materially changed before consequential execution.

Implementations operating under continuity-sensitive governance profiles **SHOULD** require revalidation where drift materially impacts authorization posture, admissibility integrity, execution safety, telemetry trust, dependency reliability, or consequence legitimacy.

Code	Meaning
<b>RC_ADMISSIBILITY_CONTEXT_EXPIRED</b>	Admissibility context expired before consequential execution authorization could be completed.
<b>RC_CONTINUITY_REVALIDATION_REQUIRED</b>	Governance revalidation is required due to elapsed continuity window, changed conditions, or degraded execution confidence.
<b>RC_DEPENDENCY_STATE_DRIFT</b>	Required dependency state materially changed after governance authorization.
<b>RC_CONTINUITY_WINDOW_EXCEEDED</b>	Maximum allowed continuity interval between authorization and execution opportunity was exceeded.
<b>RC_CONTINUITY_DRIFT_DETECTED</b>	Material execution-context drift was detected before consequential execution.
<b>RC_AUTHORITY_CONDITION_DRIFT</b>	Authority conditions changed after authorization and must be re-evaluated before execution.
<b>RC_ENVIRONMENTAL_ASSUMPTION_INVALIDATED</b>	Execution assumptions were invalidated by environmental, operational, or contextual change.
<b>RC_TELEMETRY_CONTINUITY_FAILURE</b>	Required telemetry continuity was unavailable, corrupted, incomplete, or no longer sufficient to support execution.

A conforming implementation **MAY** introduce continuity-sensitive enforcement gradually, provided that reserved continuity codes are not emitted as active governance outcomes until the runtime can detect, validate, and explain the corresponding continuity condition.

### 9.3 Constraint Severity Codes

Constraint severity defines how a policy constraint influences the final governance outcome.

Severity precedence **MUST** be applied when multiple constraints influence the same decision.

The following severity levels are normative.

Severity	Meaning
<b>INFO</b>	Informational only; no effect on governance outcome
<b>WARN</b>	Outcome allowed but flagged
<b>BLOCK</b>	Outcome <b>MUST</b> be denied or escalated
<b>DEFER</b>	Outcome <b>MUST</b> be deferred pending evidence or time
<b>DEGRADE</b>	Outcome allowed only with reduced autonomy

Constraint severity **MUST be recorded for each evaluated constraint within the Decision Receipt Policy Evaluation Block.**

Severity classification enables systems and auditors to distinguish between:

- blocking policy failures
- warning-level policy signals
- advisory constraints

### 9.4 Reason Code Usage Rules

The following rules apply to all compliant implementations.

- **At least one outcome reason code MUST appear in every Decision Receipt.**
- Reason codes **MUST remain reproducible and validator-consistent under identical governance conditions.**
- Governed consequential actions **SHOULD** include reason codes sufficient to reconstruct **authority posture, admissibility basis, and governed state transition outcome.**
- Reason codes **MUST be deterministic given identical inputs, policy bundles, and configuration state.**
- Reason codes **MUST be sufficient to justify the governance outcome.**

- Reason codes **MUST NOT expose proprietary model internals or private reasoning artifacts.**

Reason codes **SHOULD be:**

- concise
- enumerable
- stable across protocol versions
- interpretable without system-specific context

Where multiple factors influenced the outcome, **multiple reason codes MAY be emitted.**

## 9.5 Vendor Extensions

Vendors **MAY define additional reason codes** provided that:

- canonical codes are always preserved
- extended codes **do not redefine baseline meanings**
- extended codes **are namespaced**

Recommended namespacing format:

VENDORNAME\_REASON\_CODE

Example:

ACME\_MODEL\_TIMEOUT  
ACME\_SENSOR\_DRIFT

Extended reason codes **MUST be clearly distinguishable from canonical protocol codes in the Decision Receipt.**

---

## 9.6 Relationship to Audit and Compliance

Canonical reason codes enable:

- **cross-system audit correlation**
- **automated compliance validation**
- **regulator-friendly inspection**
- **post-incident root cause analysis**

Canonical reason codes also enable **validator parity, runtime reproducibility, governance testability, and deterministic replay** of governed authorization posture.

Because reason codes are standardized, **Decision Receipts can be interpreted across vendors, deployments, and domains without prior knowledge of system internals.**

This standardization ensures that governance behavior remains **transparent, explainable, and verifiable** even in heterogeneous AI ecosystems.

---

## Transition Note (Non-Normative)

With reason codes standardized, the Protocol now defines:

- what was proposed
- how it was classified
- which policies were applied
- why it was allowed, denied, escalated, deferred, degraded, blocked, invalidated, or refused
- whether governed consequence was authorized
- which governance conditions prevented or constrained execution

The next section specifies how **governed decisions flow through the system**, making governance behavior observable, testable, and enforceable across implementations.

---

## 10. Protocol Flows (Normative)

This section defines the mandatory **execution flows** for governed decisions under the **R.E.T.N.A Protocol**.

Protocol flows specify:

- how decisions traverse the **Governance Boundary**
- how governance outcomes are enforced
- how **Decision Receipts** are emitted and finalized
- how continuity is preserved across retries, escalations, or deferrals

All systems claiming **R.E.T.N.A compliance MUST implement the applicable protocol flows corresponding to each governance outcome.**

Execution flows **MUST ensure that no governed decision can reach execution without first producing a valid Decision Receipt and governance outcome.**

Execution authorization **MUST NOT** be inferred from transport completion, orchestration continuity, routing success, message delivery, or output generation.

Governed execution **MUST** instead derive from explicit governance resolution resulting in an admissible governed state transition represented by a valid Decision Receipt.

Where governance integrity conditions remain unresolved, implementations **MUST refuse consequential execution.**

## 10.1 Standard Flow — Governed Execution (ALLOW)

This flow represents the canonical path for a governed decision that is **authorized and executed.**

### Flow Steps

#### 1. Decision Construction

A **Producer** constructs a **Decision Proposal** containing:

- the proposed action
- contextual state information
- evidence references
- intended target of the action

The proposal **MUST contain sufficient information for governance evaluation.**

---

#### 2. Submission to Governance Boundary

The Producer **submits the Decision Proposal to the Governor.**

The Producer **MUST NOT execute the proposed action while governance evaluation is pending.**

Execution **MUST remain blocked until an explicit governance outcome is returned.**

---

### 3. Classification

The Governor **assigns the required Decision Classification fields**:

- Risk Class
- Action Type
- Decision Mode

These fields **MUST be recorded in the Decision Receipt**.

---

### 4. Policy Evaluation

The Governor evaluates all applicable **Policy Bundles and Constraints**, including:

- policy rules
- risk thresholds
- safety guardrails
- regulatory constraints
- dependency and system health signals

Constraint evaluation results **MUST be recorded within the Decision Receipt Policy Evaluation Block**.

---

### 5. Governance Outcome

If all blocking constraints pass, the Governor determines the outcome:

`governance_outcome = ALLOW`

Reason codes describing the evaluation results **MUST be included**.

---

### 6. Receipt Emission

**A Decision Receipt MUST be emitted prior to execution.**

The receipt **MUST contain**:

- classification fields
  - policy evaluation results
  - reason codes
  - authorization context
- 

## 7. Governed State Transition Resolution

Following Decision Receipt emission, the system **MUST** resolve whether **consequential execution is admissible as a governed state transition**.

Governed consequence **MUST NOT** proceed unless required **authority, formation, admissibility, and persistence conditions remain satisfied**.

Where governed state transition resolution succeeds, execution **MAY** proceed under the authorization represented by the Decision Receipt.

Successful routing, forwarding, orchestration completion, transport success, message delivery, or output generation **MUST NOT independently authorize consequential execution**.

**Invariant Mapping:** I9, I10, I11, I12

---

## 8. Execution

The **Executor** receives authorization and **MUST verify the Decision Receipt** before performing the action.

The Executor **MUST NOT execute a governed decision without valid authorization**.

---

## 9. Finalization

Execution status **MUST be recorded** in the Decision Receipt.

The receipt **MUST then be finalized and persisted** in the **Receipt Store**.

---

**Invariant Mapping:** I1, I2, I6

---

## 10.2 Deny Flow (DENY)

This flow applies when execution **is not permitted**.

### Flow Characteristics

- Execution **MUST NOT occur**.
- Decision Receipt emission **MUST still occur**.

### Flow Steps

1. Decision Proposal submitted to Governor
2. Classification and policy evaluation performed
3. Governance outcome determined as:

governance\_outcome = DENY

4. Decision Receipt emitted containing:
  - failing constraints
  - applicable reason codes
  - execution\_status = not\_executed
5. Receipt persisted in the Receipt Store

Denied decisions **MUST remain observable through their Decision Receipts**.

**Invariant Mapping:** I1, I2, I5

---

## 10.3 Escalation Flow (ESCALATE)

This flow applies when governance requires **review by a higher authority**.

### Escalation Targets MAY Include

- human reviewer

- supervisory agent
- safety or compliance system
- external authorization service

### **Flow Requirements**

Execution **MUST NOT occur until escalation is resolved.**

A Decision Receipt **MUST record the escalation event.**

### **Required Receipt Fields**

The receipt **MUST contain:**

governance\_outcome = ESCALATE

And include:

- escalation target identifier
- reason for escalation
- additional evidence requested (if applicable)

Escalation resolution **MAY produce:**

- a new Decision Receipt referencing the original, or
- a resolved outcome appended through receipt chaining.

**Invariant Mapping:** I2, I8

## **10.4 Defer Flow (DEFER)**

This flow applies when execution is delayed due to **incomplete or unstable conditions.**

### **Common Deferral Causes**

- insufficient evidence
- stale system state
- external dependency outage
- timing or cooldown rules

### **Flow Requirements**

The Decision Receipt **MUST specify a path to resolution.**

Execution **MUST NOT occur until governance evaluation is re-run.**

### **Required Receipt Fields**

The receipt **MUST contain:**

governance\_outcome = DEFER

And include:

- expires\_at or recheck\_at
- required evidence list
- deferral reason codes

A deferred decision **MUST be re-evaluated through the Governance Boundary prior to execution.**

---

## **10.5 Degrade Flow (DEGRADE)**

This flow applies when **execution may proceed but autonomy is reduced.**

### **Degradation Examples**

- requiring human confirmation
- selecting a safer fallback action
- reducing operational scope
- limiting resource usage

### **Flow Requirements**

Degradation **MUST be explicitly recorded in the Decision Receipt.**

Reduced autonomy **MUST be enforced by the Executor.**

### **Required Receipt Fields**

The receipt **MUST contain:**

governance\_outcome = DEGRADE

And include:

- fallback path
- autonomy restrictions
- confirmation requirements

Execution **MUST follow the reduced-autonomy constraints defined by the receipt.**

**Invariant Mapping:** I1, I8

---

## 10.6 Retry and Continuity Rules

When a decision is retried:

- decision\_id **MUST remain stable**
- new receipts **MUST reference prior receipts**
- classification changes **MUST be disclosed**
- governance outcomes **MUST be independently justified**

Retries **MUST NOT overwrite or invalidate prior Decision Receipts.**

Silent retries without receipt emission **constitute a protocol violation.**

## 10.7 Failure Handling

If the **Governor becomes unavailable:**

- execution **MUST fail closed for governed decisions**
- no implicit allow behavior **is permitted**

If **receipt persistence fails:**

- execution **MUST NOT occur**
- the failure **MUST be observable to operators or monitoring systems**

Systems **MUST NOT execute governed decisions without a persisted Decision Receipt.**

---

## 10.8 Flow Determinism and Observability

Protocol flows **MUST remain observable through Decision Receipts**.

For identical inputs, evidence references, and policy bundles:

- governance outcomes **SHOULD be reproducible within declared bounds of nondeterminism**

Flow paths **MUST be auditable post hoc** without requiring access to internal model state.

Decision Receipts therefore provide the **complete observable history of governance behavior**.

---

### Transition Note (Non-Normative)

With protocol flows defined, the **R.E.T.N.A Protocol now specifies**:

- what must happen
- when it must happen
- what evidence must exist afterward

The following section defines **how policies interface with governance evaluation**, ensuring modular policy design and interoperable enforcement across independent implementations.

---

## 11. Policy Interface (Normative)

This section defines the required interface between **governance logic** and **policy logic** within the **R.E.T.N.A Protocol**.

The R.E.T.N.A Protocol **does not mandate a specific policy language, rule engine, or constraint syntax**. Instead, the Protocol defines a **stable, interoperable policy interface** that allows policies to be authored, evaluated, versioned, and audited independently of implementation details.

This separation is **intentional and foundational**.

It ensures that:

- governance behavior remains **auditable and deterministic**

- policy authorship remains **independent of execution infrastructure**
- policy systems may evolve **without requiring protocol changes**

The Policy Interface therefore acts as the **contract between policy evaluation and governance enforcement**.

---

## 11.1 Policy Interface Objectives

The Policy Interface **MUST enable**:

- policy portability across systems and vendors
- deterministic governance outcomes
- transparent auditability without exposing proprietary logic
- independent evolution of policy and execution systems

The Policy Interface **MUST NOT**:

- assume a specific programming language
- require a particular rule engine or policy DSL
- embed policy logic directly within execution components

Policy evaluation **MUST occur within the Governance Boundary** and **MUST remain separable from Decision Construction and Decision Execution**.

## 11.2 Policy Bundle Concept

Policies are grouped into **Policy Bundles**.

A **Policy Bundle** is a **versioned, immutable collection of constraints** evaluated together during governance evaluation.

Policy Bundles allow:

- modular policy deployment
- independent policy governance
- regulatory overlays
- version-controlled policy evolution

A governance evaluation **MUST reference the Policy Bundles applied**.

### 11.2.1 Policy Bundle Requirements

Each Policy Bundle **MUST include**:

Field	Requirement
bundle_id	Stable unique identifier
bundle_version	Immutable version identifier
authority	Publisher or owning entity
effective_scope	Domains, risk classes, or action types covered

Policy Bundles **MUST be immutable once published**.

Updates **MUST result in a new bundle version**.

Implementations **MAY maintain multiple active bundle versions**, provided that the applied version **is disclosed in the Decision Receipt**.

---

## 11.3 Constraints

A **Constraint** is the smallest enforceable policy unit.

Constraints represent **individual evaluable rules** within a Policy Bundle.

Constraints **MUST be**:

- deterministic
- side-effect free
- independently evaluable

Constraints **MUST NOT perform execution actions or modify system state**.

Constraints exist solely to **inform governance decisions**.

---

### 11.3.1 Constraint Interface

Each constraint evaluation **MUST return the following fields**:

<b>Field</b>	<b>Description</b>
result	pass   fail   unknown
severity	INFO   WARN   BLOCK   DEFER   DEGRADE
reason_code	canonical or extended reason code
explanation_ref	optional structured reference

explanation\_ref **MUST NOT contain free-form narrative text**.  
It **SHOULD reference structured artifacts such as documentation IDs, policy references, or audit entries**.

Constraints **MUST NOT**:

- directly trigger execution
- mutate system state
- override governance outcomes outside defined severity semantics

### 11.4 Policy Evaluation Contract

The **Governor MUST evaluate all applicable Policy Bundles** for a given decision.

Policy applicability **MUST be resolved before constraint evaluation**.

For each constraint evaluated, the Governor **MUST record**:

- constraint identifier
- evaluation result
- severity classification
- associated reason code

The Governor **MUST determine the governance outcome by applying the most restrictive applicable severity**.

#### **Severity Precedence (Normative)**

The following severity precedence **MUST be applied**:

Priority	Severity
1	BLOCK
2	DEFER
3	DEGRADE
4	WARN
5	INFO

This precedence ordering **is normative and MUST NOT be altered**.

---

## 11.5 Determinism and Reproducibility

For identical:

- Decision Classification
- evidence references
- Policy Bundles
- configuration state

policy evaluation results **SHOULD be reproducible**.

If nondeterminism exists, it **MUST be explicitly declared and traceable**.

This requirement supports:

- audit replay
  - incident investigation
  - regulatory verification
  - governance debugging
- 

## 11.6 Policy Disclosure in Decision Receipts

Decision Receipts **MUST disclose**:

- which Policy Bundles were applied
- which constraints were evaluated
- which constraints influenced the final governance outcome

Decision Receipts **MUST NOT embed**:

- raw policy logic
- proprietary scoring formulas
- internal decision trees
- implementation-specific rule syntax

This requirement ensures **governance transparency without intellectual property leakage**.

---

## 11.7 Policy Authority and Trust

Policy Bundles **MUST identify their issuing authority**.

Authorities **MAY include**:

- system operators
- enterprises
- regulators
- industry consortiums

Governors **MAY enforce trust rules**, including:

- accepting bundles only from approved authorities
- prioritizing regulatory policies over local policies
- applying domain-specific trust hierarchies

Trust evaluation itself **MAY be policy-driven**.

---

## 11.8 Policy Scope and Applicability

Policies **MAY apply conditionally** based on:

- Risk Class
- Action Type

- Decision Mode
- operational domain or environment
- actor identity

Policy scope resolution **MUST occur prior to constraint evaluation.**

Scope resolution decisions **MUST be auditable.**

---

## 11.9 Extensibility

Implementations **MAY extend the Policy Interface** to support advanced policy models, including:

- richer explanation references
- probabilistic constraints
- temporal policies
- cross-receipt correlation

Extensions **MUST NOT:**

- alter baseline constraint semantics
- bypass constraint severity precedence
- suppress required policy disclosures

All extensions **MUST remain compatible with the canonical Decision Receipt structure.**

### Transition Note (Non-Normative)

With the Policy Interface defined, the **R.E.T.N.A Protocol establishes governance without lock-in:**

- policy authors remain independent
- execution systems remain replaceable
- governance enforcement remains deterministic

The next section specifies **how evidence is captured, referenced, and verified**, completing the core governance triangle:

classification → policy → evidence

---

## 12. Evidence Handling (Normative)

This section defines the requirements for **capturing, referencing, protecting, and verifying evidence** used during **Decision Construction** and evaluated at the **Governance Boundary**.

Within the R.E.T.N.A Protocol, **evidence is treated as a first-class governance input**, not merely as diagnostic logging.

Evidence referenced during governance evaluation **MUST therefore satisfy integrity, traceability, and auditability requirements defined by the applicable Compliance Profile**.

Evidence referenced by a Decision Receipt **constitutes the factual basis upon which the governance outcome was determined**.

Evidence handling **MAY** additionally include **governance-relevant formation references** sufficient to support admissibility verification, orchestration traceability, authority reconstruction, and governed state transition inspection.

The Protocol does not require disclosure of proprietary reasoning internals or unrestricted chain-of-thought visibility.

Formation evidence **SHOULD** instead preserve **governance-relevant shaping context necessary for auditability and execution integrity verification**.

### 12.1 Evidence Principles

All compliant implementations **MUST adhere to the following principles**.

#### 1. Sufficiency

Evidence referenced by a Decision Receipt **MUST be sufficient to justify the resulting governance outcome**.

If evidence is incomplete or unverifiable, governance evaluation **MUST account for this condition**, potentially resulting in outcomes such as **DEFER, ESCALATE, or DENY**.

#### 2. Minimization

Decision Receipts **SHOULD reference evidence rather than embed raw data whenever feasible**.

Embedding raw evidence **SHOULD be avoided when external references are sufficient for audit and verification purposes**.

### 3. Integrity

Evidence references **MUST be tamper-evident** at the integrity level defined by the applicable Compliance Profile.

Integrity guarantees **MAY include hashing, immutable storage, or cryptographic attestation mechanisms.**

### 4. Traceability

Evidence referenced during governance evaluation **MUST be linkable to the Decision Receipt that relied upon it.**

Evidence items **MUST include identifiers that allow auditors to trace the relationship between evidence and governance outcomes.**

---

## 12.2 Evidence Representation

Evidence **MAY be represented in one of two forms.**

### 1. Embedded Summary

A **bounded, non-sensitive representation of evidence** included directly in the Decision Receipt.

Embedded summaries **SHOULD contain only the minimum information necessary for audit interpretation.**

### 2. External Reference

A **pointer to evidence stored outside the Decision Receipt.**

External references **MAY include:**

- internal datastore identifiers
- URI-like references
- immutable object storage references
- ledger or archival storage references

External references **SHOULD be used for Compliance Profiles B and C**, where evidence integrity verification is required.

Decision Receipts **MUST explicitly indicate which representation form is used.**

---

## 12.3 Evidence Item Requirements

Each evidence item referenced by a Decision Receipt **MUST include the following fields.**

Field	Requirement
evidence_id	Stable unique identifier
type	Evidence category (see Section 12.6 Evidence Types (Baseline))
source	Originating system, sensor, or actor
timestamp	Time of capture or generation
location_ref	Internal reference or URI-like pointer
summary	Bounded, human-readable description

For **Compliance Profiles B and C**, each evidence item **MUST additionally include:**

Field	Requirement
integrity_ref	Hash or equivalent integrity verification reference

The integrity reference **MUST enable verification that the referenced evidence has not been altered since capture.**

---

## 12.4 Evidence Integrity and Verification

Evidence integrity guarantees vary by **Compliance Profile.**

### 12.4.1 Integrity Guarantees by Profile

#### Profile A — Basic Governance

Evidence storage **SHOULD support append-only logging.**

Integrity guarantees **MAY rely on system-level logging controls.**

---

#### Profile B — Enterprise Governance

Evidence storage **MUST be tamper-evident.**

Acceptable mechanisms **MAY include:**

- cryptographic hashes
- checksums
- immutable object storage
- write-once storage systems

Evidence integrity **MUST be verifiable after receipt emission.**

#### Profile C — Regulatory Governance

Evidence **MUST support cryptographic attestation or equivalent integrity anchoring.**

Acceptable mechanisms **MAY include:**

- digital signatures
- cryptographic proof chains
- distributed ledger anchoring
- trusted timestamping services

Integrity verification **MUST be independently reproducible by auditors.**

---

### 12.4.2 Verification Failures

If evidence integrity verification fails:

- the governance evaluation **MUST record the integrity failure**
- appropriate reason codes **MUST be emitted** (e.g., RC\_EVIDENCE\_INTEGRITY\_FAIL)
- execution **MUST NOT proceed for governed decision classes**

Integrity verification failure **SHOULD be treated as a high-risk anomaly condition.**

---

## 12.5 Evidence Lifecycle and Retention

Evidence storage and access **MUST respect declared governance policies**, including:

- retention\_class
- access\_class
- applicable regulatory or contractual requirements

Evidence **MAY be handled independently of Decision Receipts**, provided that receipt integrity remains verifiable.

Permitted lifecycle operations **MAY include**:

- retention in external storage systems
- redaction after receipt finalization
- archival storage migration
- controlled destruction according to retention policies

Receipt references **MUST remain valid for the declared retention period.**

---

## 12.6 Evidence Types (Baseline)

The following evidence types define **baseline categories** recognized by the R.E.T.N.A Protocol.

Implementations **MAY extend this list**, provided baseline semantics remain unchanged.

Type	Identifier	Description
Sensor	SENSOR	Physical or digital sensor readings
User Input	USER_INPUT	Explicit user instructions or confirmations
State Snapshot	STATE_SNAPSHOT	Captured system or environment state
Retrieved Document	RETRIEVED_DOC	Retrieved reference material (e.g., RAG output)
Derived Signal	DERIVED_SIGNAL	Computed features, scores, or anomaly signals
External Event	EXTERNAL_EVENT	Webhooks, alerts, or third-party signals

Additional evidence types **MAY be introduced through extension mechanisms.**

---

## 12.7 Evidence Minimization and Privacy

To protect privacy and reduce exposure of sensitive data:

Decision Receipts **MUST NOT embed sensitive raw data when references are sufficient.**

Evidence summaries **MUST be bounded and redactable.**

Access to raw evidence **MUST be governed independently from receipt visibility.**

Implementations **SHOULD support:**

- selective disclosure mechanisms
- redaction policies
- differential access control based on role or authority

Privacy protections **MUST NOT compromise the integrity or traceability of governance decisions.**

---

## 12.8 Evidence and Reproducibility

Where feasible, evidence references **SHOULD enable**:

- re-evaluation of governance decisions
- audit replay
- post-incident forensic analysis

Evidence references **SHOULD remain accessible for the declared retention period**.

If full reproduction is not possible due to data volatility or external system constraints, the Decision Receipt **MUST disclose this limitation**.

This disclosure **MUST allow auditors to understand the evidentiary limitations of the decision**.

---

## Transition Note (Non-Normative)

With evidence handling defined, the protocol now specifies:

- what governance decisions rely upon
- how that reliance is recorded
- how evidence integrity is preserved over time

The next section formalizes **trust, security, and privacy guarantees**, ensuring that the protocol remains **defensible in adversarial or regulated environments**.

---

## 13. Trust, Security, and Privacy (Normative)

This section defines the **minimum trust, security, and privacy guarantees** required of any system claiming **R.E.T.N.A Protocol compliance**.

The R.E.T.N.A Protocol treats **governance as a security-critical control plane**.

Failures of governance enforcement **MUST therefore be treated as safety-critical failures**, not merely operational defects.

Failures of governed state transition enforcement **MUST** be treated as consequential governance failures capable of producing unauthorized operational, legal, financial, physical, security, or material system impact.

Implementations **MUST** therefore ensure that consequential execution cannot proceed under unresolved authority, admissibility, formation, evidence, integrity, or receipt persistence conditions.

Systems implementing the protocol **MUST ensure that governance decisions cannot be bypassed, altered, or executed without proper authorization and receipt traceability.**

## 13.1 Authentication and Authorization

### 13.1.1 Actor Authentication

Actors participating in the governance workflow **MUST be authenticated prior to performing protocol actions.**

Specifically:

- The Governor **MUST authenticate Producers** submitting Decision Proposals.
- The Executor **MUST authenticate the Governor's authorization** prior to executing a governed action.

Acceptable authentication mechanisms **MAY include:**

- mutual TLS (mTLS)
- signed identity tokens
- hardware-backed identities
- cryptographic API credentials
- equivalent strong authentication mechanisms

Decision Proposals submitted by **unauthenticated or unverifiable actors MUST be rejected.**

---

### 13.1.2 Authorization Scope

Authorization **MUST enforce strict role separation** within the governance workflow.

The following authorization model **MUST be enforced:**

<b>Role</b>	<b>Authorization</b>
Producer	<b>MAY</b> construct and submit Decision Proposals
Governor	<b>MAY</b> evaluate policies and issue governance outcomes
Executor	<b>MAY</b> execute actions only when explicitly authorized

Producers **MUST NOT** execute governed actions.

Executors **MUST NOT** perform policy evaluation.

Governors **MUST NOT** execute actions directly unless explicitly operating in a combined role with enforced logical separation.

Privilege escalation **MUST** be prevented by design.

**Invariant Mapping:** I1, I8

---

## 13.2 Least Privilege and Separation of Duties

Implementations **MUST** adhere to the principle of least privilege.

Specifically:

- Producers **MUST NOT** possess execution credentials
- Executors **MUST NOT** possess policy evaluation logic
- Receipt Stores **MUST NOT** permit unauthorized modification

Where system components are **co-located within the same service or runtime**, implementations **MUST** preserve logical separation of duties.

This separation **MUST** remain enforceable through authentication, authorization, or runtime isolation mechanisms.

## 13.3 Tamper Resistance and Integrity

### 13.3.1 Receipt Integrity

All **Decision Receipts MUST be tamper-evident.**

Integrity guarantees **vary by Compliance Profile.**

#### Profile A

Append-only receipt storage **SHOULD be used.**

#### Profile B

Tamper-evident storage **MUST be implemented**, including mechanisms such as:

- cryptographic hashes
- immutable storage
- write-once logging

#### Profile C

Cryptographic attestation or equivalent ledger anchoring **MUST be implemented.**

If receipt tampering is detected, the system **MUST:**

- record the integrity violation
- surface the violation to governance monitoring systems
- treat the event as a governance failure

---

### 13.3.2 Evidence Integrity

Evidence integrity requirements are defined in **Section 12 (Evidence Handling)** and **MUST be enforced consistently with the declared Compliance Profile.**

Integrity verification failures **MUST influence governance outcomes.**

For governed decision classes, integrity verification failures **MUST prevent execution.**

---

## 13.4 Privacy Controls

### 13.4.1 Data Minimization

Implementations **MUST minimize exposure of personal or sensitive data.**

Decision Receipts **SHOULD reference evidence rather than embed raw data when possible.**

Evidence summaries **MUST be bounded and redactable.**

Access to raw data **MUST be governed independently of receipt visibility.**

### 13.4.2 Access Classification

Decision Receipts **MUST include an access\_class field** describing the confidentiality level of the receipt.

Example classifications **MAY include:**

- public
- internal
- restricted
- regulatory

Access enforcement mechanisms **MUST be auditable.**

---

### 13.4.3 Retention Classification

Decision Receipts **MUST include a retention\_class field** specifying retention requirements.

The retention classification **MUST define:**

- minimum retention duration
- maximum retention duration
- archival or deletion rules

Retention policies **MUST comply with:**

- contractual obligations
- regulatory requirements

- applicable Policy Bundles

## 13.5 Human Override and Accountability

Human intervention **MAY override governance outcomes** where operational or regulatory policies permit.

When a human override occurs:

- the override **MUST be recorded**
- the identity or role of the overrider **MUST be disclosed**, subject to privacy rules
- the justification **MUST be captured via reason codes**

Human overrides **MUST NOT**:

- erase or delete prior Decision Receipts
- silently alter governance outcomes
- bypass audit traceability requirements

**Invariant Mapping:** I2, I7

---

## 13.6 Threat Model (Baseline)

The R.E.T.N.A Protocol is designed to mitigate the following baseline threats:

- unauthorized execution of actions
- policy bypass attempts
- evidence forgery
- untraceable decision execution
- post-hoc alteration of decision history
- silent escalation of autonomous authority

This threat model **defines the minimum security assumptions of the protocol**.

Threats outside this baseline **MAY be addressed through domain-specific extensions**.

---

## 13.7 Fail-Closed Behavior

For governed decision classes, the system **MUST fail closed** when any of the following conditions occur:

- Governor unavailability
- Decision Receipt persistence failure
- integrity verification failure
- authentication failure
- authorization failure

Under fail-closed conditions:

- governed actions **MUST NOT execute**
- the failure condition **MUST be observable**

Implicit “allow” behavior **is a protocol violation.**

---

## Transition Note (Non-Normative)

With trust, security, and privacy defined, the **R.E.T.N.A Protocol establishes governance enforcement as a secure control plane.**

The protocol therefore guarantees that governed decisions remain:

- enforceable
- auditable
- privacy-aware
- resilient under adversarial conditions

The remaining sections address **interoperability, conformance testing, and protocol evolution**, completing the protocol’s path toward **standards-grade specification.**

---

## 14. Interoperability and Extensibility (Normative)

This section defines how **R.E.T.N.A-compliant systems interoperate across vendors, deployments, and domains**, and how the Protocol may be extended **without fragmentation or vendor lock-in**.

The R.E.T.N.A Protocol is designed as **a governance protocol rather than a product architecture**. Interoperability is therefore **a primary requirement**, not an implementation convenience.

Any system claiming protocol compliance **MUST preserve interoperability at the level of canonical governance artifacts**, regardless of internal system design.

---

### 14.1 Interoperability Principles

All compliant implementations **MUST adhere to the following principles**.

#### Semantic Stability

The meaning of **canonical protocol fields MUST remain consistent across implementations**.

Implementations **MUST NOT reinterpret or redefine the semantics of baseline protocol fields**.

#### Artifact Portability

Decision Receipts **MUST be interpretable outside the originating system**.

External systems **MUST be able to evaluate the structure, integrity, and governance outcome of a receipt without requiring proprietary internal logic**.

#### Vendor Neutrality

Protocol compliance **MUST NOT depend on proprietary software components**.

Any conforming system **MUST be able to validate Decision Receipts produced by another conforming implementation**.

#### Forward Compatibility

Protocol extensions **MUST NOT break baseline conformance**.

Implementations **MUST tolerate unknown non-critical fields** without failing validation.

---

## 14.2 Canonical Artifact Interoperability

### 14.2.1 Decision Receipt as the Interoperable Unit

The **Decision Receipt** is the **atomic unit of interoperability** within the R.E.T.N.A Protocol.

A system **MUST be considered interoperable at the protocol level** if it:

1. emits Decision Receipts conforming to **Section 8**, and
2. preserves all **required semantics and integrity guarantees**.

Protocol interoperability **does not depend on**:

- internal architecture
  - programming language
  - runtime environment
  - deployment model
- 

### 14.2.2 Serialization and Canonicalization

The R.E.T.N.A Protocol **does not mandate a specific serialization format**.

However, implementations **MUST define and publish canonicalization rules** sufficient to ensure cross-system verification.

Canonicalization rules **MUST specify**:

- canonical field ordering
- normalization rules for values
- canonical hashing inputs

These rules **MUST enable**:

- consistent receipt hashing

- cross-system verification
  - independent third-party audit validation
- 

### 14.2.3 Receipt Validation and Canonical Verification (Normative)

A conforming implementation **MUST** support validation of Decision Receipts at the artifact layer.

Receipt validation **MUST** be capable of verifying at least the following:

- presence of all required canonical receipt fields
- presence of required substructures for identities, decision envelope, policy evaluation, outcome, integrity metadata, and dispatch trace
- canonical hash correctness where a receipt hash is present
- profile-specific integrity expectations applicable to the declared compliance profile

A validator **MUST** reject receipts that:

- omit required governance-critical fields
- omit required actor identity disclosure
- omit a governance outcome
- omit integrity metadata required by the declared profile
- contain an invalid canonical receipt hash

Validation behavior **MUST be based on observable receipt content and MUST NOT require access to proprietary model internals.**

---

## 14.3 Vendor Extensions

### 14.3.1 Permitted Extensions

Implementations **MAY extend the protocol** by introducing additional fields or categories, including:

- additional evidence types
- additional reason codes
- domain-specific constraints
- supplemental risk metrics
- auxiliary metadata fields

Extensions **MUST remain compatible with baseline protocol semantics.**

### 14.3.2 Extension Constraints

Extensions **MUST NOT**:

- alter the meaning of baseline protocol fields
- omit required canonical fields
- redefine baseline reason codes
- bypass governance outcomes
- interfere with Decision Receipt integrity verification

All extensions **MUST be**:

- namespaced
- explicitly declared
- safely ignorable by baseline validators

Unknown extension fields **MUST NOT invalidate otherwise compliant Decision Receipts**.

---

### 14.3.3 Payload Kind Registry

Payload Kinds constitute a controlled vocabulary governing semantic request categories used for routing, privilege determination, and policy applicability.

Implementations supporting Payload Kind **MUST** define a registry or equivalent controlled list of recognized values.

Payload Kind registries **MUST** preserve the following rules:

- baseline meanings **MUST** remain stable
- privileged categories **SHOULD** be explicitly declared
- unknown or ambiguous privileged categories **MUST** fail closed where required by policy
- implementation-specific extensions **MUST** be documented and **SHOULD** be namespaced when shared across system boundaries

Changes to baseline Payload Kind semantics constitute protocol-level changes and **MUST NOT be introduced silently**.

---

### 14.3.4 Registry-Governed Extension Points

The following protocol surfaces are registry-governed extension points:

- canonical reason codes
- payload kinds
- evidence type extensions
- action type extensions
- domain profile identifiers

For all such extension points:

- baseline protocol semantics **MUST NOT** be redefined
- required canonical fields **MUST NOT** be bypassed
- implementation-specific values **SHOULD** be documented in machine-readable form where practical
- unknown non-critical values **MAY** be ignored only where such behavior is safe
- unknown privileged categories **MUST fail closed where required by policy**

This section does not require a specific governance body for registry management in v0.1, but implementers **SHOULD** maintain transparent versioned records of registry changes.

---

## 14.4 Capability Advertisement

Governors **SHOULD advertise their protocol capabilities** in a machine-readable form.

Capability descriptions **MAY include**:

- supported Compliance Profile (RETNA-A, RETNA-B, RETNA-C)
- supported policy bundle formats
- supported attestation mechanisms
- supported receipt export formats

Capability discovery **enables**:

- automated system integration
  - toolchain compatibility
  - cross-vendor orchestration
-

## 14.5 Compliance Profiles (Normative)

The R.E.T.N.A Protocol defines Compliance Profiles to support interoperable adoption across heterogeneous systems with differing governance, operational, risk, and consequence requirements.

Compliance Profiles define the minimum governance, admissibility, persistence, authority, formation, and execution requirements applicable to a conformant implementation.

Higher-order profiles **SHALL** preserve the requirements of lower-order profiles unless explicitly superseded.

The Protocol defines the following baseline profiles:

- **RETNA-A — Foundational Governance Profile**
- **RETNA-B — Governed Execution Profile**
- **RETNA-C — Critical Consequence Profile**

### 14.5.1 Baseline Compliance Profile Requirements

Capability	RETNA-A	RETNA-B	RETNA-C
Decision Receipt emission	REQUIRED	REQUIRED	REQUIRED
Receipt persistence	REQUIRED	REQUIRED	REQUIRED
Canonical Reason Codes	REQUIRED	REQUIRED	REQUIRED
Admissibility evaluation	MINIMAL	REQUIRED	REQUIRED
Real-time admissibility enforcement	OPTIONAL	REQUIRED	REQUIRED
Authority validation	LIMITED	REQUIRED	REQUIRED

Capability	RETNA-A	RETNA-B	RETNA-C
Formation integrity preservation	OPTIONAL	REQUIRED WHERE MATERIAL	REQUIRED
Governed state transition resolution	OPTIONAL	REQUIRED	REQUIRED
Execution authorization enforcement	LIMITED	REQUIRED	REQUIRED
Consequential execution blocking	OPTIONAL	REQUIRED	REQUIRED
Evidence retention	MINIMAL	MODERATE	HIGH
Deterministic authorization	OPTIONAL	REQUIRED	REQUIRED
Persistence-before-consequence	OPTIONAL	REQUIRED	REQUIRED
Runtime governance validation	OPTIONAL	STRONG	FULL
Human escalation pathways	OPTIONAL	RECOMMENDED	REQUIRED

### 14.5.2 RETNA-A — Foundational Governance Profile

RETNA-A establishes the minimum viable governance posture required for baseline protocol accountability.

RETNA-A implementations **MUST** emit Decision Receipts and canonical Reason Codes and **MUST** preserve sufficient governance metadata to support auditability and interoperability.

RETNA-A **MAY** operate with limited authority validation, minimal admissibility enforcement, bounded persistence guarantees, and reduced runtime governance enforcement.

**Cross-Reference:** §§8.5, 11, 16

### 14.5.3 RETNA-B — Governed Execution Profile

RETNA-B establishes the minimum governance posture required for governed consequential execution.

RETNA-B implementations **MUST** support admissibility evaluation, governed state transition resolution, authority validation, execution authorization enforcement, and persistence-backed governance verification.

Consequential execution **MUST** remain governance-bound under RETNA-B.

**Cross-Reference:** §§7, 8, 11, 16

---

### 14.5.4 RETNA-C — Critical Consequence Profile

RETNA-C establishes the highest baseline governance posture for systems operating in materially consequential domains including physical, financial, healthcare, legal, industrial, infrastructure, or safety-sensitive environments.

RETNA-C implementations **MUST** enforce full admissibility, authority validation, governed state transition integrity, persistence guarantees, runtime governance validation, and execution refusal where governance conditions cannot be resolved.

**Cross-Reference:** §§7, 11, 13, 16

## 14.6 Cross-System Validation

Any third-party system **SHOULD be able to:**

1. ingest a Decision Receipt
2. verify structural protocol compliance
3. validate integrity claims
4. interpret the governance outcome
5. verify declared Compliance Profile conformance

This validation **MUST NOT require access to proprietary internal decision logic.**

Cross-system validation is essential for:

- regulators

- auditors
  - insurance providers
  - incident response teams
- 

## 14.7 Backward and Forward Compatibility

### 14.7.1 Backward Compatibility

Minor protocol versions **MUST**:

- preserve required canonical fields
- preserve semantic meaning of existing fields
- allow previously issued Decision Receipts to remain valid

Breaking semantic changes **MUST only occur in major protocol revisions**.

---

### 14.7.2 Forward Compatibility

Implementations **MUST ignore unknown fields** unless those fields are explicitly marked as **critical extensions**.

This requirement enables:

- gradual ecosystem adoption
  - safe experimentation with extensions
  - protocol evolution without disruption
- 

## 14.8 Domain Profiles (Optional)

Industries **MAY** define **domain-specific profiles** layered on top of the R.E.T.N.A Protocol to address sector-specific governance, operational, regulatory, safety, or execution requirements.

Example domain profiles include:

- healthcare governance profile
- financial services governance profile

- industrial automation governance profile
- smart infrastructure governance profile
- smart appliance governance profile

Domain profiles **MUST**:

- inherit baseline protocol invariants
- declare any additional governance, admissibility, authority, evidence, or execution constraints
- remain interoperable at the Decision Receipt level
- declare the minimum required Compliance Profile applicable to the domain

Domain profiles **MUST NOT**:

- alter baseline protocol semantics
- redefine canonical governance outcomes
- weaken required protocol invariants
- bypass Governance Boundary enforcement

Illustrative minimum profile expectations **MAY** include:

<b>Domain</b>	<b>Minimum Compliance Profile</b>
Healthcare	RETNA-C
Financial Services	RETNA-C
Industrial Automation	RETNA-B or RETNA-C
Smart Infrastructure	RETNA-B or RETNA-C
Smart Appliance Systems	RETNA-A or RETNA-B

Domain profile requirements **MAY exceed baseline Compliance Profile requirements** where operational consequence, regulatory obligation, safety posture, or execution risk necessitates stronger governance guarantees.

---

## Transition Note (Non-Normative)

With interoperability and extensibility defined, the R.E.T.N.A Protocol establishes a **stable governance core with controlled evolution**, a prerequisite for long-term industry adoption.

The next section defines **conformance and compliance verification**, completing the protocol's enforceability loop.

---

## 15. Conformance and Testability (Normative)

This section defines the requirements for **demonstrating conformance to the R.E.T.N.A Protocol** and establishes **testable criteria** for validation by:

- internal engineering teams
- third-party auditors
- regulatory authorities

A system **MUST NOT claim R.E.T.N.A Protocol compliance** unless it can demonstrably satisfy the requirements defined in this section.

Conformance verification **MUST be based on observable protocol behavior**, not on internal implementation details.

---

### 15.1 Conformance Claims

A system **MAY claim conformance to the R.E.T.N.A Protocol only if it satisfies all of the following conditions:**

1. It **implements all Normative (MUST/SHALL) requirements** applicable to its declared Compliance Profile.
2. It **emits Decision Receipts conforming to Section 8.**
3. It **upholds all System Invariants (I1–I12)** defined by the protocol.
4. It **passes the conformance tests defined in this section.**

Conformance declarations **MUST include:**

Field	Requirement
protocol_version	Implemented protocol version (e.g., RETNA v0.1)
compliance_profile	One of: RETNA-A, RETNA-B, or RETNA-C
governed_scope	Declared scope of governed decision classes

Systems **MUST NOT advertise protocol compliance without declaring these attributes.**

## 15.2 Mandatory Conformance Tests

The following conformance tests **MUST be supported by all compliant implementations.**

These tests validate that the **governance control plane cannot be bypassed and remains observable.**

### 15.2.1 Governance Boundary Integrity Test

#### Objective

Verify that governed decisions cannot execute without governance evaluation.

#### Test Procedure

Attempt to execute a governed action **without submitting a Decision Proposal to the Governor.**

#### Expected Result

- Execution **MUST fail.**
- No system side effects **MUST occur.**
- The failure **MUST be observable.**

**Invariant Mapping:** I1

---

### 15.2.2 Receipt Emission Test

#### Objective

Verify that all governed decisions produce a Decision Receipt.

#### Test Procedure

Submit governed decisions that result in each possible governance outcome:

- ALLOW
- DENY
- ESCALATE

- DEFER
- DEGRADE

### Expected Result

- Exactly **one Decision Receipt MUST be emitted per governance evaluation.**
- Receipts **MUST persist regardless of outcome.**

### Invariant Mapping: I2

---

## 15.2.3 Receipt Completeness Test

### Objective

Verify that Decision Receipts include all required canonical fields.

### Test Procedure

Inspect emitted receipts for required fields defined in **Section 8**.

### Expected Result

- All required canonical fields **MUST be present.**
  - Silent omission of required fields **MUST NOT occur.**
- 

## 15.2.4 Policy Trigger Test

### Objective

Verify correct constraint enforcement during governance evaluation.

### Test Procedure

Introduce a policy constraint designed to fail.

### Expected Result

- The governance outcome **MUST reflect the constraint severity.**

- Appropriate **reason codes MUST be emitted**.

**Invariant Mapping:** I5

---

### 15.2.5 Deterministic Outcome Test

#### Objective

Verify governance reproducibility.

#### Test Procedure

Submit identical Decision Proposals containing:

- identical evidence references
- identical policy bundles
- identical configuration state

#### Expected Result

- Governance outcomes **SHOULD be identical**.
- Any nondeterminism **MUST be explicitly declared**.

**Invariant Mapping:** I6

---

### 15.2.6 Receipt Validator Conformance Test

#### Objective

Verify that a conforming validator can correctly accept valid receipts and reject materially invalid receipts.

#### Test Procedure

Provide the validator with:

1. a structurally complete receipt with a valid canonical hash
2. a receipt missing one or more required canonical fields
3. a receipt with an invalid canonical hash

4. a receipt missing required profile-specific integrity elements for the declared profile

### **Expected Result**

The validator **MUST**:

- accept the valid receipt
- reject the structurally incomplete receipt
- reject the receipt with an invalid canonical hash
- reject the receipt that fails profile-specific integrity requirements

Validation failure **MUST be observable and machine-readable.**

---

## **15.2.7 Governance Integrity Validation**

Governance integrity validation verifies that a conforming implementation enforces the Protocol's execution integrity requirements beyond passive receipt emission or post-execution observability.

Conforming implementations **MUST** demonstrate that governed execution cannot proceed where required governance conditions remain unresolved, invalid, inadmissible, unverifiable, or non-persisted.

Governance integrity validation **MUST** evaluate the implementation's ability to:

- enforce authority integrity
- enforce admissibility requirements
- preserve formation integrity where applicable
- refuse invalid governed state transitions
- prevent execution under receipt persistence failure conditions

Validation behavior **MUST** be observable, deterministic where applicable, and traceable through Decision Receipts, Reason Codes, validator outputs, or equivalent governance artifacts.

---

## **15.2.8 Authority Integrity Validation Test**

Objective

Verify that governed execution cannot proceed under unresolved, invalid, expired, contradictory, or unverifiable authority conditions.

Test Procedure

Submit governed decision proposals containing:

- invalid authority assertions
- missing authority metadata
- expired authority context
- unauthorized actor relationships
- Edge Policy violations

Expected Result

- governed execution **MUST** fail closed
- execution authorization **MUST NOT** be inferred from transport position, routing behavior, orchestration continuity, or system adjacency
- the failure **MUST** be observable through governance artifacts and associated Reason Codes

**Invariant Mapping:** I9, I12

---

### 15.2.9 Admissibility Enforcement Validation Test

Objective

Verify that governed execution cannot proceed unless admissibility evaluation is explicitly resolved.

Test Procedure

Submit governed decisions with:

- incomplete admissibility metadata
- unresolved policy evaluation conditions
- insufficient evidence
- unresolved risk conditions
- invalid governance classification state

Expected Result

- execution **MUST NOT** proceed
- admissibility failure **MUST** be represented through governance metadata and Reason Codes
- implementations **MUST** fail closed where admissibility cannot be established

**Invariant Mapping:** I10, I12

### 15.2.10 Formation Integrity Validation Test

#### Objective

Verify that governance-sensitive decisions preserve sufficient formation context where required by policy, Risk Class, Payload Kind, or governance profile.

#### Test Procedure

Submit governed decision proposals lacking materially required formation metadata, orchestration lineage references, participating actor disclosures, or governance shaping context.

#### Expected Result

- implementations enforcing formation integrity requirements **MUST** reject or invalidate the governed decision
- missing formation integrity conditions **MUST** be observable through governance outputs or validator behavior
- execution **MUST NOT** proceed where formation integrity requirements remain unresolved

#### Invariant Mapping: I11

---

### 15.2.11 Governed State Transition Blocking Test

#### Objective

Verify that transport success, forwarding completion, orchestration continuity, output generation, routing success, or persistence alone cannot independently authorize governed consequence.

#### Test Procedure

Attempt consequential execution under conditions where:

- no admissible governed state transition exists
- receipt persistence is incomplete
- authority conditions are invalid
- execution authorization is unresolved
- governance outcome integrity is missing or contradictory

#### Expected Result

- consequential execution **MUST** be refused
- successful transport semantics **MUST NOT** independently constitute governed execution success

- governed consequence **MUST** require explicit governed state transition authorization

**Invariant Mapping:** I1, I10, I12

---

### 15.2.12 Receipt Persistence Refusal Test

Objective

Verify that consequential execution cannot proceed where required Decision Receipt persistence fails.

Test Procedure

Simulate:

- receipt storage failure
- receipt persistence interruption
- integrity persistence mismatch
- receipt durability verification failure

Expected Result

- execution **MUST** fail closed
- implementations **MUST** prevent consequential execution until receipt persistence succeeds
- persistence failure **MUST** be observable and auditable

**Invariant Mapping:** I2, I7, I12

---

### 15.2.13 Runtime Governance Validation Doctrine

Implementations **MAY** enforce governance integrity through runtime validation mechanisms operating prior to consequential execution.

Such mechanisms **MAY** include:

- governance validators
- admissibility validators
- authority verification pipelines
- receipt integrity validators
- governed state transition validators
- policy enforcement gateways

- persistence verification controls

Where implemented, runtime governance validation **MUST** remain aligned with the protocol invariants, Decision Receipt semantics, canonical Reason Codes, and Governance Boundary enforcement requirements defined by this specification.

---

## 15.3 Profile-Specific Tests

Additional conformance tests apply depending on the declared **Compliance Profile**.

### 15.3.1 Profile B (RETNA-B) Tests

Implementations claiming **RETNA-B** compliance **MUST additionally pass the following tests**.

#### Evidence Integrity Test

##### Test

Introduce an evidence hash mismatch.

##### Expected Result

- The mismatch **MUST be detected**.
- The failure **MUST be recorded in the Decision Receipt**.

#### Tamper-Evident Storage Test

##### Test

Attempt unauthorized modification of stored Decision Receipts.

##### Expected Result

- The modification **MUST be detectable**.
- 

### 15.3.2 Profile C (RETNA-C) Tests

Implementations claiming **RETNA-C** compliance **MUST additionally pass the following tests**.

### Attestation Verification Test

#### Test

Verify receipt signatures or ledger anchors.

#### Expected Result

- Attestation verification **MUST succeed for valid receipts.**
- 

### Human Override Recording Test

#### Test

Perform a governance override via authorized human intervention.

#### Expected Result

- The override **MUST produce explicit receipt entries.**
- 

### Third-Party Verifiability Test

#### Test

Provide receipts to an external validation system.

#### Expected Result

- The external validator **MUST be able to verify receipt integrity and governance outcomes.**

## 15.4 Negative Testing Requirements

Conformance testing **MUST include negative scenarios**, including:

- missing evidence
- malformed receipts

- unauthorized Producers
- unavailable Governor
- Decision Receipt persistence failure

For all such conditions, the system **MUST fail closed**.

Governed actions **MUST NOT execute under these failure conditions**.

---

## 15.5 Audit Replay Capability

Implementations **SHOULD support audit replay capabilities**, including:

- replay of Decision Receipts
- reconstruction of decision context (within privacy limits)
- export of receipts for regulatory or incident review

Replay capability **MAY be partial** where evidence volatility prevents full reconstruction, provided the limitation **is disclosed in the Decision Receipt**.

---

## 15.6 Independent Verification

A compliant implementation **SHOULD be independently verifiable** by:

- internal audit teams
- third-party assessors
- regulators

Verification **MUST be possible without access to**:

- proprietary model internals
- source code
- private policy logic

Decision Receipts **MUST provide sufficient information to enable verification of governance outcomes**.

---

## Transition Note (Non-Normative)

With conformance and testability defined, the R.E.T.N.A Protocol now specifies **how governance trust is proven rather than merely asserted**.

The remaining sections address:

- implementation guidance
- protocol versioning and evolution
- ecosystem adoption strategies

These complete the protocol's transition from **conceptual governance framework to deployable infrastructure specification**.

---

## 16. Reference Implementation Guidance (Non-Normative)

This section provides **implementation guidance** for adopting the **R.E.T.N.A Protocol** in real-world systems.

This section is **non-normative** and **does not impose additional requirements** beyond those defined in earlier sections of this specification.

The purpose of this guidance is to:

- accelerate adoption
- reduce the risk of incorrect implementations
- demonstrate practical deployment patterns
- illustrate feasibility without introducing vendor lock-in

### 16.1 Architectural Placement

R.E.T.N.A is designed to operate at the **Governance Boundary**, positioned between **Decision Construction** and **Decision Execution**.

Several architectural deployment patterns are commonly used.

#### Gateway Pattern

The **Governor operates as an authorization gateway** positioned in front of execution systems.

Execution requests pass through the Governor before any action is performed.

---

### Library Pattern

The Governor is implemented as a **governance library embedded inside an application**, invoked synchronously before execution.

This approach may be suitable for tightly integrated systems.

---

### Service Pattern

The Governor operates as a **standalone governance service**, accessed via an API by Producers and Executors.

This pattern supports:

- distributed systems
  - multi-agent architectures
  - cross-service governance enforcement
- 

All deployment patterns are acceptable provided that **System Invariants (I1-I12) remain preserved**.

---

### 16.1.1 Runtime Governance Enforcement (Illustrative)

Reference implementations **MAY** enforce governance integrity through runtime validation mechanisms capable of:

- refusing inadmissible execution
- invalidating unresolved governed state transitions
- enforcing receipt-backed authorization
- verifying authority integrity
- validating formation metadata requirements

- enforcing admissibility resolution requirements
- validating governance classification consistency
- preventing execution under receipt persistence failure conditions

Such enforcement mechanisms **MAY** operate:

- synchronously at the Governance Boundary
- asynchronously through governance validation pipelines
- through execution gateways
- through policy enforcement services
- through validator-based governance layers
- through receipt validation and persistence verification workflows

provided that consequential execution remains governance-bound under the protocol invariants defined by this specification.

Reference implementations **SHOULD** distinguish:

- transport success from governed consequence
- orchestration continuity from execution authorization
- output generation from governed state transition authorization
- routing completion from admissibility satisfaction

Governed execution **SHOULD** therefore derive from explicit governance resolution rather than communication-layer success semantics alone.

Reference implementations **MAY** additionally support:

- governed state transition validation
- authority verification pipelines
- formation integrity enforcement
- runtime admissibility validation
- execution refusal semantics
- receipt persistence enforcement
- canonical governance artifact validation
- validator-driven governance observability

These mechanisms strengthen interoperability between specification doctrine and operational enforcement behavior while preserving implementation flexibility across heterogeneous architectures.

## 16.2 Minimal Reference Components

A minimal reference implementation typically contains the following components.

### 1. Governance Interface

Accepts **Decision Proposals** and returns governance outcomes.

Typical responsibilities include:

- receiving Decision Proposals
  - evaluating proposals against policies
  - emitting Decision Receipts
- 

### 2. Policy Evaluation Module

Responsible for:

- loading Policy Bundles
- evaluating Constraints
- producing reason codes
- determining constraint severity results

### 3. Receipt Store

Responsible for:

- persisting Decision Receipts
  - enforcing immutability or tamper-evidence
  - enabling audit retrieval
- 

### 4. Evidence Reference Layer

Responsible for:

- generating evidence identifiers
- storing or referencing raw artifacts
- producing integrity hashes where required

## 5. Optional Attestation Module

Used primarily for **RETNA-C deployments**.

Responsibilities may include:

- signing Decision Receipts
  - anchoring receipts to immutable ledgers
  - producing cryptographic verification artifacts
- 

## 16.3 Decision Proposal Interface (Illustrative)

A typical **Decision Proposal** may include:

- proposed action summary
- target entity or resource
- evidence references
- relevant environment context
- requesting Producer identity

Implementations typically ensure that:

- proposals are structured
  - proposals are authenticated
  - proposals are immutable once submitted
- 

## 16.4 Receipt Emission Strategy

Decision Receipts are generally emitted **prior to execution** for governed decisions.

Typical emission patterns include:

<b>Outcome</b>	<b>Emission Timing</b>
ALLOW	emitted prior to execution
DENY	emitted immediately
DEGRADE	emitted prior to degraded execution
ESCALATE	emitted upon escalation
DEFER	emitted upon deferral

Execution systems typically **delay action until receipt emission succeeds**, ensuring governance traceability.

---

## 16.5 Receipt Storage Patterns

Common storage approaches for Decision Receipts include:

- append-only databases
- write-once object stores
- ledger-backed systems
- SIEM-integrated audit stores

Operational implementations often support:

- receipt retrieval by identifier
  - filtering by risk class or outcome
  - export for audit or regulatory review
- 

## 16.6 Validation and Tooling

Implementers are encouraged to provide tooling such as:

- receipt schema validators
- integrity verification utilities

- receipt viewers or dashboards
- automated conformance test harnesses

These tools ideally operate **independently of proprietary execution logic**, allowing external validation.

---

### 16.6.1 Reference Validator Tooling

Reference implementations may provide tooling such as:

- receipt schema validators
- canonical hash verification utilities
- profile-aware receipt validators
- machine-readable registry exports for payload kinds and reason codes

Such tooling is particularly useful for CI, audit replay workflows, cross-system verification, and independent conformance assessment.

Where provided, reference tools **SHOULD remain protocol-first** and **SHOULD avoid coupling validation semantics to proprietary execution logic**.

---

## 16.7 HomeSphere AI as a Reference Implementation (Informative)

HomeSphere AI provides an example implementation demonstrating the R.E.T.N.A Protocol across:

- multi-agent orchestration
- physical device actuation
- real-time sensor evidence ingestion
- policy-driven governance evaluation
- auditable Decision Receipt generation

HomeSphere illustrates **one possible deployment architecture**.

R.E.T.N.A remains **independent of HomeSphere** and can be implemented by any system that satisfies the protocol requirements.

## 16.8 Common Implementation Pitfalls

Implementers commonly encounter several avoidable mistakes.

Examples include:

- treating Decision Receipts as operational logs rather than governance artifacts
- emitting receipts after execution instead of before execution
- embedding sensitive data unnecessarily within receipts
- silently retrying decisions without emitting new receipts
- coupling policy logic directly to execution logic

Avoiding these pitfalls improves **auditability, interoperability, and protocol compliance**.

---

## 17. Licensing and Implementation Rights (Normative)

The R.E.T.N.A Protocol Specification defines a proprietary decision governance framework developed and owned by Value Intelligence Solutions Inc.

This document is made available for informational, evaluative, and research purposes. Access to this specification does not grant any rights to implement, reproduce, modify, distribute, or commercialize the protocol or any derivative system without explicit authorization.

Implementations of the R.E.T.N.A Protocol in production, commercial, or operational environments **MAY** require a valid license issued by Value Intelligence Solutions Inc.

All conformant implementations **MUST preserve the core invariants, governance boundary enforcement, decision receipt requirements, and policy evaluation mechanisms defined in this specification.**

**The following restrictions apply:**

- Unauthorized implementation of the R.E.T.N.A Protocol for commercial or operational use is prohibited
- Creation of derivative governance frameworks that replicate the core architectural model, decision receipt structure, or policy enforcement mechanisms may be subject to enforcement
- Redistribution of this specification without proper attribution is not permitted

**Permitted uses include:**

- Internal evaluation and research
- Academic study and analysis

- Non-commercial prototyping for validation purposes

Organizations seeking to implement the R.E.T.N.A Protocol in production systems **SHOULD obtain a formal license or partnership agreement with Value Intelligence Solutions Inc.**

Value Intelligence Solutions Inc. reserves the right to define licensing tiers, compliance requirements, certification programs, and conformance validation processes for implementations of the protocol.

Nothing in this specification shall be interpreted as granting implicit rights to deploy, commercialize, or operationalize the R.E.T.N.A Protocol without authorization.

All trademarks, system designs, architectural patterns, and governance mechanisms described herein remain the exclusive intellectual property of Value Intelligence Solutions Inc.

---

## 18. Versioning and Evolution (Normative)

This section defines the **versioning model, compatibility rules, and evolution guarantees** of the **R.E.T.N.A Protocol**.

The protocol is intended to function as a **long-lived governance standard**. Protocol evolution **MUST therefore remain predictable, transparent, and safe** for adopters operating in **regulated, safety-critical, or mission-critical environments**.

Protocol changes **MUST preserve the auditability and interpretability of previously issued Decision Receipts**.

---

### 18.1 Version Identifier Format

R.E.T.N.A protocol versions **MUST follow the format**:

major.minor

Example versions include:

0.1

1.0

1.1

Each **Decision Receipt MUST include the protocol version** under which the receipt was emitted.

This version identifier **enables validators to interpret receipts according to the appropriate protocol semantics.**

---

## 18.2 Minor Version Rules

Minor version increments (for example 0.1 → 0.2 or 1.0 → 1.1) represent **backward-compatible protocol evolution.**

Minor versions **MAY introduce:**

- additional optional fields
- clarifications of semantic definitions
- new canonical reason codes
- new evidence or action type categories
- expanded non-normative guidance

Minor versions **MUST NOT:**

- remove required canonical fields
- alter the meaning of existing fields
- break Decision Receipt integrity verification
- invalidate previously issued Decision Receipts

Receipts emitted under earlier minor versions **MUST remain valid and interpretable.**

---

## 18.3 Major Version Rules

Major version increments (for example 0.x → 1.0 or 1.x → 2.0) represent **significant protocol evolution.**

Major versions **MAY introduce:**

- new required fields
- removal or deprecation of legacy constructs
- new compliance profiles
- refinements to protocol invariants with explicit migration paths

Major versions **MUST**:

- publish explicit migration guidance
- preserve interpretability of previously issued Decision Receipts
- avoid silent semantic changes

Even when implementations evolve, **receipt verification and audit interpretation MUST remain possible across protocol versions.**

---

## 18.4 Deprecation Policy

When protocol elements are deprecated, including:

- fields
- reason codes
- protocol behaviors

the following rules apply:

- deprecation **MUST be explicitly declared** in the specification
- deprecated elements **MUST remain accepted for a defined compatibility period**
- Decision Receipts **SHOULD indicate deprecated usage where applicable**

Silent deprecation **MUST NOT occur.**

---

## 18.5 Receipt Compatibility Across Versions

Validators **MUST support cross-version receipt validation.**

Specifically, validators **MUST**:

- accept Decision Receipts emitted under earlier protocol versions
- ignore unknown non-critical fields

- preserve receipt hash verification using the canonicalization rules declared by the emitting version

Decision Receipts **MUST NOT require re-issuance solely due to protocol upgrades.**

Previously issued receipts **remain valid governance artifacts.**

## 18.6 Governance of the Protocol Itself

Evolution of the R.E.T.N.A Protocol **SHOULD occur through a transparent governance process.**

Such a process typically includes:

- public change proposals
- published rationale for protocol modifications
- versioned specification releases

Protocol stewardship **MAY be administered by:**

- an independent foundation
- an industry consortium
- a neutral steward entity with publicly declared governance rules

Regardless of the governing body, protocol evolution **SHOULD prioritize interoperability, stability, and audit continuity.**

---

## 18.7 Stability Guarantees

The R.E.T.N.A Protocol provides the following long-term stability guarantees.

1. Decision Receipts **remain auditable indefinitely.**
2. Governance outcomes **remain interpretable across protocol versions.**
3. Policy logic **remains decoupled from protocol evolution.**
4. Conformance claims **remain verifiable over time.**

Protocol changes **MUST preserve these guarantees.**

---

## Transition Note (Non-Normative)

With versioning and evolution defined, the R.E.T.N.A Protocol ensures that:

- early adopters are protected from breaking changes
- regulated deployments remain valid across protocol revisions
- the protocol can evolve without fragmentation

The final section addresses **adoption strategy and ecosystem positioning**, completing the **R.E.T.N.A v0.1 specification**.

---

## 19. Adoption Path (Non-Normative, Strategic)

This section outlines a **practical staged adoption strategy** for the **R.E.T.N.A Protocol**, intended to accelerate industry uptake while preserving interoperability, regulatory trust, and long-term standardization potential.

The adoption model is designed to:

- minimize operational disruption
- reduce integration friction
- allow organizations to realize immediate value
- support gradual progression toward full governance enforcement

R.E.T.N.A can be adopted incrementally without requiring full system refactoring.

### 19.1 Adoption Philosophy

R.E.T.N.A adoption is guided by three core principles.

#### Receipts First

Adoption typically begins by **emitting Decision Receipts**, even before governance enforcement is enabled.

This allows organizations to establish **decision traceability and audit visibility** immediately.

---

#### Governance Gradually Enforced

Governance controls may initially operate in **observational mode**, with enforcement enabled progressively as confidence and operational maturity increase.

---

### **Protocol Before Platform**

The protocol is designed to remain **valuable independently of any particular product or vendor implementation**.

Organizations should be able to implement the protocol using their existing systems and architectures.

---

## **19.2 Stage 0 — Observability-Only Adoption**

### **Objective**

Establish visibility into autonomous or semi-autonomous decisions without altering execution behavior.

### **Typical Characteristics**

- Decision Receipts emitted alongside normal execution
- Governance outcomes recorded but not enforced
- Receipts analyzed for operational insights

### **Benefits**

- no operational risk
- immediate decision audit trail
- organizational familiarity with governance concepts

## **19.3 Stage 1 — Governance Boundary Introduction**

### **Objective**

Introduce the **Governance Boundary** for selected decision classes.

### **Typical Characteristics**

- decision risk classification enabled
- policy evaluation activated
- governance outcomes recorded but not enforced

### **Benefits**

- visibility into policy violations
  - improved operational awareness
  - early regulatory readiness
- 

## **19.4 Stage 2 — Selective Enforcement**

### **Objective**

Enable governance enforcement for higher-risk decisions.

### **Typical Characteristics**

- blocking enabled for higher-risk decision classes (e.g., R3 / R4)
- escalation and deferral workflows active
- human confirmation required where appropriate

### **Benefits**

- meaningful risk reduction
  - improved compliance posture
  - clear accountability boundaries
- 

## **19.5 Stage 3 — Full Compliance Profiles**

### **Objective**

Achieve **RETNA-B** or **RETNA-C** compliance.

### **Typical Characteristics**

- evidence integrity verification
- tamper-evident receipt storage

- cryptographic attestation or ledger anchoring
- third-party verification capability

### Benefits

- regulator-grade governance
  - audit defensibility
  - improved insurance and liability posture
- 

## 19.6 Tooling and Ecosystem Enablement

Adoption is typically accelerated when an ecosystem provides supporting tooling such as:

- Decision Receipt validators
- receipt visualization tools
- audit export utilities
- conformance test harnesses
- policy authoring and validation tools

Such tooling ideally remains **protocol-first and vendor-neutral**.

## 19.7 Industry Standardization Path

A potential long-term trajectory for protocol maturation may include:

1. public publication of the protocol specification
2. independent reference implementations
3. formation of a neutral stewardship body or working group
4. engagement with formal standards organizations
5. alignment with regulatory frameworks and guidance

The protocol has been designed to support such standardization efforts **without requiring architectural redesign**.

---

## 19.8 Strategic Positioning

The R.E.T.N.A Protocol positions **Decision Governance** as a foundational layer of the emerging AI infrastructure stack.

Within this stack, R.E.T.N.A can be understood as analogous to:

- **TLS** for secure transport
- **OAuth** for identity delegation
- **OpenTelemetry** for observability

In this context, **Decision Receipts function as a portable governance artifact**, enabling trust, accountability, and verification across organizational boundaries.

---

## 19.9 HomeSphere AI as a Catalyst (Informative)

HomeSphere AI provides a working example demonstrating how the R.E.T.N.A Protocol can be applied in a real-world system involving:

- multi-agent orchestration
- physical device actuation
- sensor-derived evidence inputs
- policy-driven governance decisions
- auditable Decision Receipt generation

HomeSphere serves as an **illustrative reference implementation**, not as a required dependency.

R.E.T.N.A remains **fully implementable by any system satisfying the protocol requirements**.

## Closing Statement (Non-Normative)

R.E.T.N.A Protocol v0.1 introduces **Decision Governance as a first-class, enforceable, and auditable capability** for AI systems operating in real-world environments.

By centering governance on **portable Decision Receipts**, the protocol enables:

- trust without vendor lock-in
- accountability without pervasive surveillance
- innovation without compromising safety

The protocol provides a foundation upon which organizations, regulators, and industry partners can build **reliable and trustworthy autonomous systems**.

## Appendix A — Glossary

### Access Classification

A policy-defined categorization determining which actors or systems may access specific governance artifacts, evidence references, or Decision Receipts.

---

### Action Type

A classification describing the type of action proposed by a decision, such as purchase, notification, device actuation, or modification of system access controls.

---

### Actor

An entity participating in a governed decision process, including systems, services, agents, devices, or human participants.

---

### Actor Authentication

The process of verifying the identity of an actor participating in a governed decision process.

---

### Admissibility

The governance determination that a candidate decision satisfies the conditions required for governed consequence under applicable authority, policy, evidence, formation, classification, and risk constraints.

---

### Admissibility Outcome

The recorded governance determination establishing whether a candidate decision satisfies the conditions required for governed consequence. Admissibility outcomes **SHALL** remain reconstructable through the Decision Receipt.

**Cross-Reference:** §§5, 8.3.10, 9, 11

---

## Anomaly Signal

A derived signal indicating unusual, abnormal, or potentially unsafe conditions that may influence policy evaluation or risk classification.

---

## Attestation

Optional cryptographic proof verifying the authenticity and integrity of a Decision Receipt or associated governance artifact.

---

## Authorization Boundary

The enforcement interface at which a proposed decision must obtain authorization before execution.

In this Protocol, the Authorization Boundary is implemented as the Governance Boundary.

---

## Authorization Scope

The defined set of actions, resources, or decision categories an authenticated actor is permitted to initiate or execute.

---

## Authority Integrity

The condition under which the authority basis for a governed decision is explicit, valid, verifiable, and within authorized scope. Ambiguous, expired, contradictory, or unresolved authority **SHALL NOT** independently authorize execution.

---

---

## Authority Provenance

The identifiable basis from which execution authority originates, including asserted authority source, scope, delegation posture, and verification status where applicable.

**Cross-Reference:** §§7 (19), 8.3.3, 9.2.6

---

## Backward Compatibility

The ability of newer protocol implementations to interpret and process artifacts created by earlier protocol versions.

---

## Canonical Artifact

A governance artifact whose structure and semantics are defined by the Protocol and are intended to be interoperable across implementations.

---

## Canonical Reason Code

A standardized machine-readable code representing the outcome or justification of a governance decision.

---

## Canonical Serialization

The standardized representation of protocol artifacts ensuring consistent interpretation across systems.

## Canonicalization Rules

The set of rules defining how protocol artifacts must be structured and serialized to ensure interoperability.

## Capability Advertisement

A mechanism by which a system declares supported protocol features, compliance profiles, or extension capabilities.

---

## Compliance Profile

A defined implementation posture specifying the minimum governance, persistence, evidence, admissibility, and execution requirements applicable to a conformant deployment of the Protocol (for example RETNA-A, RETNA-B, or RETNA-C).

**Cross-Reference:** §§8.5, 15, 16

---

## Conformance Claim

A declaration that a system implementation satisfies the requirements of a specified protocol version and compliance profile.

---

## Conformance Test

A test verifying that an implementation behaves according to protocol requirements.

---

## Consequential Execution

Execution capable of producing Governed Consequence within a target environment or system. Consequential execution **MUST** remain governance-bound under the Protocol.

**Cross-Reference:** §§5, 7 (I1, I12), 11

---

## Constraint

An enforceable rule evaluated during policy evaluation that restricts or conditions decision execution.

---

## **Constraint Evaluation**

The process of evaluating individual constraints within a policy bundle.

---

## **Constraint Severity**

A classification indicating the enforcement priority of a constraint.

Canonical severity levels:

BLOCK  
DEFER  
DEGRADE  
WARN  
INFO

---

## **Context Capture**

The process of collecting relevant state, signals, and environmental information required for decision evaluation.

---

## **Decision Class**

A classification grouping decisions according to operational domain or governance requirements.

---

## **Decision Construction**

The process of forming a candidate action or plan using inputs, signals, models, or rules.

---

## **Decision Execution**

The act of applying an authorized decision to the environment.

---

## **Decision Governance**

The protocol-governed process by which candidate decisions are evaluated, authorized, and executed according to defined policies.

---

## **Decision Mode**

A configuration defining the operational posture of a system, such as fully autonomous operation, human-in-the-loop governance, or advisory mode.

---

## **Decision Receipt**

A structured artifact containing provenance information, evidence references, policy evaluations, authorization outcomes, and associated reason codes for a governed decision.

---

## **Deterministic Outcome Test**

A conformance test verifying that identical inputs and policies always produce the same governance outcome.

---

## **Dispatch Trace**

A graph-like representation of the actors and processes involved in constructing and authorizing a governed decision.

---

## Domain Profile

A domain-specific extension of the protocol defining additional governance rules or compliance requirements.

---

## Edge Policy

A governance rule specifying which actors may issue or execute particular categories of decisions.

---

## Embedded Evidence Summary

A compact representation of evidence included directly within a Decision Receipt.

---

## Execution Consequence

The resulting state, effect, or external outcome produced following authorized consequential execution. Execution consequence SHALL remain attributable to a corresponding Decision Receipt.  
**Cross-Reference:** §§5, 11

---

## Evidence

Inputs used during decision construction or policy evaluation.

Evidence may include sensor signals, user commands, retrieved documents, state snapshots, or derived signals.

---

## Evidence Identifier

A unique identifier assigned to an evidence artifact.

---

### **Evidence Integrity Reference**

A reference or verification mechanism used to confirm the integrity of stored evidence.

---

### **Evidence Integrity Verification**

The process of confirming that referenced evidence has not been altered or corrupted.

---

### **Evidence Item**

A single piece of evidence used in decision construction or policy evaluation.

---

### **Evidence Lifecycle**

The full lifecycle of evidence artifacts including creation, reference, verification, retention, and eventual expiration.

---

### **Evidence Reference**

A pointer or identifier referencing stored evidence used during decision construction or governance evaluation.

---

### **Evidence Representation**

The format used to encode evidence artifacts for governance processing.

---

### **Evidence Retention**

The policy-defined duration for which evidence artifacts must be preserved.

---

## **Evidence Reproducibility**

The ability to reproduce or reconstruct evidence used in decision evaluation.

---

## **Evidence Source**

The origin of a piece of evidence, such as a sensor, external system, or user input.

---

## **Evidence Type**

A standardized category describing the origin or nature of evidence.

Baseline evidence types include:

SENSOR  
USER\_INPUT  
STATE\_SNAPSHOT  
RETRIEVED\_DOC  
DERIVED\_SIGNAL  
EXTERNAL\_EVENT

---

## **Execution Refusal**

The prevention of consequential execution when governance authorization cannot be established or required governance conditions remain unresolved.

---

## **Executor**

The actor responsible for carrying out an authorized decision.

---

## **Fail-Closed Behavior**

A security principle requiring systems to deny execution when governance validation cannot be completed successfully.

---

## **Formation Context**

Governance-relevant information describing how a candidate decision was shaped prior to admissibility evaluation and Governance Boundary enforcement. Formation Context MAY include upstream actors, constraints, environmental assumptions, or shaping inputs.

**Cross-Reference:** §§5, 7 (I11), 8.3.5

---

## **Formation Integrity**

The preservation of sufficient governance-relevant context to explain how a candidate decision was shaped prior to admissibility evaluation and Governance Boundary enforcement.

---

## **Forward Compatibility**

The ability of older implementations to safely process artifacts generated by newer protocol versions.

---

## **Governance Boundary**

The mandatory interface between Decision Construction and Decision Execution where governance evaluation occurs.

---

## **Governance Boundary Integrity**

The condition under which governed execution remains blocked until required authority, admissibility, policy, and state transition conditions have been resolved. Governance Boundary

integrity SHALL prevent unauthorized consequential execution.

**Cross-Reference:** §§2, 5, 7 (I1, I2, I10, I12)

---

## **Governance Boundary Integrity Test**

A conformance test verifying that execution cannot occur without passing through the Governance Boundary.

---

## **Governance Outcome**

The final authorization result produced by governance evaluation.

Canonical outcomes:

ALLOW

DENY

ESCALATE

DEFER

DEGRADE

---

## **Governed Consequence**

An operational, legal, financial, physical, security, environmental, or materially consequential outcome requiring explicit governance authorization prior to execution. Governed consequence SHALL NOT be inferred from transport, routing, forwarding, or output generation alone.

**Cross-Reference:** §§5, 7 (I12), 11

---

## **Governed State Transition**

The formally evaluated transition between prior state and candidate consequential state that is explicitly authorized, denied, deferred, degraded, escalated, blocked, invalidated, or refused under the Protocol.

**Cross-Reference:** §§5, 7 (I12), 8.3.11, 11

---

## **Governor**

The actor responsible for evaluating policies and producing governance outcomes.

---

## **Governed Execution**

Execution of a decision after successful authorization at the Governance Boundary.

---

## **Human Review**

A governance escalation mechanism requiring human approval prior to execution.

---

## **Incremental Adoption**

A staged implementation approach in which governance capabilities are introduced progressively.

---

## **Integrity and Audit Block**

The portion of a Decision Receipt responsible for integrity verification and audit metadata.

---

## **Interoperability**

The ability of independent systems to exchange and validate Decision Receipts and governance artifacts.

---

## **Least Privilege**

A security principle requiring actors to possess only the permissions necessary to perform their roles.

---

## **Major Version**

A protocol version increment indicating incompatible changes.

---

## **Minor Version**

A protocol version increment indicating backward-compatible enhancements.

---

## **Minimal Receipt (RETNA-A)**

A minimal Decision Receipt structure required for basic governance compliance.

---

## **Observability-Only Adoption**

An adoption stage in which Decision Receipts are generated for visibility and auditing without enforcing governance outcomes.

---

## **Outcome Block**

The portion of a Decision Receipt recording the governance outcome and associated reason codes.

---

## **Payload Kind**

A classification identifying the semantic category of a decision request used to route governance evaluation.

---

## **Policy**

A set of enforceable rules evaluated at the Governance Boundary.

---

## **Policy Applicability**

The scope within which a policy bundle is applied.

---

## **Policy Authority**

The system or entity responsible for defining and distributing governance policies.

---

## **Policy Bundle**

A collection of policy rules, constraints, and configuration parameters used during governance evaluation.

---

## **Policy Bundle Version**

The version identifier associated with a policy bundle.

---

## **Policy Constraint**

An individual rule within a policy restricting execution behavior.

---

## **Policy Evaluation**

The process of applying policy rules and constraints to a candidate decision.

---

## **Policy Evaluation Block**

The section of a Decision Receipt documenting policy evaluation results.

---

## **Privileged Execution**

Execution of a decision that alters external state and therefore requires explicit authorization.

---

## **Producer**

The actor responsible for constructing a candidate decision.

---

## **Protocol Version**

The version identifier associated with a specific revision of the Protocol.

---

## **Reason Code**

A standardized machine-readable code representing the result of policy evaluation.

---

## **Receipt Compatibility**

A standardized machine-readable code representing governance semantics, including policy evaluation, admissibility, authority, evidence, persistence, governed state transition, or execution

posture conditions.

---

---

## **Receipt Emission**

The generation of a Decision Receipt as part of governance evaluation.

---

## **Receipt Integrity**

The property ensuring that a Decision Receipt cannot be altered without detection.

---

## **Receipt Immutability**

The property that once issued, a Decision Receipt cannot be modified.

---

## **Receipt Persistence**

The durable preservation of Decision Receipts as immutable governance artifacts sufficient to support auditability, admissibility reconstruction, execution verification, and reproducibility.

**Cross-Reference:** §§7 (I2, I6), 8.5, 11

---

## **Receipt Store**

A system responsible for persisting Decision Receipts for audit and verification.

## **Retention Classification**

A policy-defined category specifying how long governance artifacts must be retained.

---

---

## **Risk Class**

A classification representing the potential impact or severity of a decision.

---

### **Separation of Duties**

A governance principle requiring different actors to perform decision construction, evaluation, and execution.

---

### **Stability Guarantees**

Protocol assurances ensuring that implementations remain interoperable across version updates.

---

### **Tamper-Evident Storage**

A storage mechanism that detects unauthorized modification of governance artifacts.

---

### **Vendor Extension**

An optional protocol extension introduced by a specific implementation.

---

### **Vendor-Neutral**

A design principle ensuring the protocol can be implemented across independent systems without vendor lock-in.

---

## Appendix B — Payload Kind Registry (Informative Reference)

This appendix points to the machine-readable payload kind registry used by the reference implementation.

Suggested publication note:

The reference implementation publishes a machine-readable payload kind registry describing canonical payload categories, privilege-sensitive categories, and baseline semantics. This registry is informative for v0.1 and does not override the normative rules of Sections 8 and 14.

Download the recommended companion artifact:

[retna\\_payload\\_kind\\_registry.json](#)

---

## Appendix C — Canonical Reason Code Registry (Informative Reference)

This appendix points to the machine-readable canonical reason code registry used by the reference implementation.

Suggested publication note:

The reference implementation publishes a machine-readable reason code registry describing canonical baseline reason codes, category groupings, and short semantic descriptions. This registry is informative for v0.1 and does not override the normative rules of Section 9.

Download the recommended companion artifact:

[retna\\_reason\\_code\\_registry.json](#)