

# AI OSI Core Protocol v0.1

Decision Evidence & Temporal Validity Protocol

Standards Track Candidate

Last Updated: 23 January 2026

## Abstract

AI OSI Core defines a universal protocol for recording, preserving, and expiring decisions made by or with the assistance of artificial intelligence systems.

The protocol enables future reconstruction of *who decided what, using which inputs, under what constraints, and whether that decision remains valid*. It is tool-agnostic, organization-agnostic, and sector-agnostic.

AI OSI Core does not define safety, ethics, correctness, or compliance. It defines interoperable decision evidence and temporal validity semantics that downstream audits, regulators, insurers, and governance systems may rely upon.

## Design Principles (Normative)

AI OSI Core is governed by the following principles:

- **Universality** — usable by anyone, anywhere, for any AI system
- **Human-Executable** — implementable with paper, spreadsheets, or files
- **Binary Verifiability** — conformance is pass/fail only
- **Temporal Discipline** — decisions must expire or be replaced
- **Accountability Preservation** — a single accountable owner must be named
- **Non-Enforcement** — the protocol records decisions; it does not police behavior

## Non-Goals (Explicit)

AI OSI Core does not:

- define acceptable risk, safety, or compliance thresholds
- require disclosure of proprietary or confidential evidence
- guarantee honesty, completeness, or integrity of decision timelines
- enforce decision quality or ethical correctness
- replace audits, regulators, courts, or governance bodies

## Terminology (Normative)

The key words **MUST**, **MUST NOT**, **SHOULD**, and **MAY** are to be interpreted as described in RFC 2119.

## The Core Object: Decision Record

A **Decision Record (DR)** is a single, immutable record representing exactly one finalized decision.

All Decision Records **MUST** include the following fields:

1. **Record ID**

A globally unique identifier.

2. **Timestamp**

The exact UTC date and time when the decision was finalized.

3. **Executing System**

The AI system, software, or model used to produce or execute the decision, including version and hash where available. If unknown, this **MUST** be explicitly stated.

4. **Accountable Party**

The human, role, or organization accepting responsibility for the decision. This field **MUST** identify a single accountable owner.

5. **Action**

A plain-text description of the authorized decision or command.

6. **Inputs**

A list of data sources, documents, prompts, models, or artifacts relied upon. Each input **MUST** be referenceable by identifier, version, date, or hash.

7. **Constraints**

The explicit limits or conditions in force at the time of decision.

8. **Validity Limit**

A timestamp or objectively testable condition after which the decision is void. If a condition is used, it **MUST** be objectively verifiable at review time.

9. **Supersedes ID** (Optional)

The Record ID of a prior Decision Record that this record replaces.

## Protocol Rules (Strict)

- **Atomicity** — Each Decision Record **MUST** represent exactly one decision.
- **Immutability** — Once finalized, a Decision Record **MUST NOT** be edited or deleted.
- **Referenceability** — All listed Inputs **MUST** be retrievable; failure invalidates the decision.
- **Validity** — A decision is valid **ONLY** if the current time precedes its Validity Limit.
- **Supersession** — Superseded records are immediately invalid regardless of remaining time.

## Conformance (Binary)

An implementation is AI OSI Core v0.1 conformant for a decision if and only if:

- a Decision Record exists
- all required fields are present
- the Validity Limit has not been exceeded
- any superseded records are not relied upon

There are no maturity levels, partial credit, or scoring systems.

## Minimal Human Test (Operational)

An independent party **MUST** be able to answer the following using only the records:

1. What decision was made?
2. Who is accountable for it?
3. What AI system was involved?
4. What inputs were relied upon?
5. What constraints applied?
6. When does the decision expire?
7. Has it been replaced?

Failure to answer any question constitutes operational failure, even if records technically exist.

## Security, Privacy, and Misuse Considerations

AI OSI Core assumes adversarial conditions.

Implementations **SHOULD**:

- mark unknown or uncertain fields explicitly
- prevent AI systems from fabricating Decision Records or Inputs
- separate sensitive evidence from public records when required
- avoid implying that recording a decision implies correctness or safety

AI OSI Core does not guarantee honesty, completeness, or timeline integrity. It guarantees clarity and inspectability of recorded decisions.

## Publication and Canonical Sources (Non-Normative)

AI OSI Core v0.1 is published as a stable, standalone specification.

The canonical distribution sources are:

- Project website: <https://aiosi.org>
- Canonical repository: <https://github.com/danielpmadden/ai-osi-core>

The website provides orientation and background material. The repository is the authoritative source for the specification text, version history, reference implementations, and errata.

If a discrepancy exists between copies of this document, the version published in the canonical repository SHALL be considered authoritative.

## Appendix A — Operational Compliance Checklist (Non-Normative)

This checklist evaluates functional accessibility of Decision Records without adding new protocol requirements.

### A.1 Identity & Accountability

Record ID is unique within scope

Accountable Party names a reachable owner

Executing System is traceable in system inventory

### A.2 Evidence & Constraints

All listed Inputs are retrievable now

Constraints clearly bound AI behavior

Action description aligns with system outputs or logs

### A.3 Temporal Validity

Current UTC time precedes the Validity Limit

Condition-verification data is accessible if used

Superseded records are explicitly marked invalid

### A.4 Failure Modes

Failure	Operational Symptom	Protocol Violation
Black Hole	Inputs referenced but unreachable or deleted	Referenceability
Zombie Decision	Expired decision still relied upon	Validity
Ghost Owner	No real accountable owner exists	Accountability
Multi-Decision	One record contains multiple decisions	Atomicity

## Overall Verdict

AI OSI Core v0.1 establishes a universal accountability primitive. It does not judge decisions, but it ensures that decisions leave an inspectable, time-bounded trail that cannot be silently ignored.

*End of Specification.*