

Open request to NIST: a decision-record standard for the agent economy

Version 1.0 · July 29, 2026

Published as an open proposal on July 29, 2026. It has not yet been submitted to NIST.

Request

The NIST AI Agent Standards Initiative is developing an interoperable and secure foundation for agents that act on a user's behalf. This open request asks NIST to evaluate a complementary technical object: a portable, versioned decision record that preserves what a consequential agent saw, predicted, selected, rejected, and later learned.

The proposed Decision Record & Control Standard Consumer E-Commerce profile (DRCS-EC) is a candidate for a bounded public pilot. It is a working specification whose field performance, privacy properties, interoperability, and production cost require testing.

Why a decision record belongs in the standards discussion

Authentication can show which agent acted. Authorization can show that the user delegated a task. A transaction receipt can show the selected outcome. These objects usually do not show the serious option set, rejected alternatives, scoring trace, feed provenance, or the agent's prediction before it acted.

A common decision-record shape would allow independent evaluators to apply different authorized tests to the same evidence. The standard would define the evidence interface and conformance rules. Marketplaces, users, auditors, regulators, courts, procurers, and researchers would retain authority over their own evaluation rules and consequences.

This separation aligns with the NIST AI Risk Management Framework's Govern, Map, Measure, and Manage functions: governance sets authority and accountability; mapping defines the decision context and affected parties; measurement tests the record and outcome; management supplies response, review, and correction paths.

Proposed technical work program

NIST should consider a public workshop and reference implementation program covering:

1. Core schema

- stable record, decision, agent, principal, and task identifiers;
- schema, policy, model, scoring-rule, and evidence-feed versions;
- delegation, authority, consent, constraints, and purpose;

- evidence and provenance;
- uncertainty and a scoreable pre-decision prediction;
- selected and rejected serious options with reasons;
- affected stakeholders and declared system-welfare direction;
- review, correction, reconsideration, and reversal conditions; and
- append-only amendments and supersession links.

2. Consumer e-commerce profile

- seller and product identity;
- price and total landed cost;
- shipping and delivery terms;
- return and warranty terms;
- payment path and checkout friction;
- reliability and safety inputs;
- affiliation, placement, sponsorship, or marketplace signals used by the agent;
- chosen option;
- rejected serious alternatives;
- recomputable scoring trace; and
- outcome follow-up against the user's instruction.

3. Conformance and interoperability

- normative JSON schema and human-readable rendering;
- required-object and type validation;
- deterministic test fixtures;
- pre-decision creation and timestamp tests;
- reproducible scoring-trace tests;
- cross-provider export and import;
- version compatibility and migration behavior;
- signature, hash-link, and tamper-evidence profiles;
- error, correction, revocation, and supersession semantics; and
- open negative fixtures showing incomplete, contradictory, stale, or misleading records.

4. Privacy and security

- data minimization by evaluation purpose;
- local-first or selective disclosure where appropriate;
- separation of identity, commercial, and evaluation data;

- access logging and reader authorization;
- retention and deletion rules;
- data-subject correction and appeal;
- resistance to fabricated alternatives, hidden feed substitution, omitted material options, and post-decision rewriting; and
- threat models for collusion among agents, platforms, merchants, readers, and record services.

5. Testing, evaluation, verification, and validation

The pilot should measure whether the record helps a declared evaluator answer a declared question. Candidate questions include instruction adherence, steering, omitted serious alternatives, score reproducibility, feed provenance, outcome follow-up, and repeated cross-provider patterns. Each test should publish its false-positive and false-negative boundary and distinguish technical conformance from policy or legal judgment.

Evidence and limits

The July 2026 working paper *Standardize the Record, Not the Reader* supplies a controlled computational result.

- The evidence path alone has no pricing effect under common seeds.
- Forecast scoring has no pricing effect in this implementation.
- Welfare-priced reward changes the private objective and conduct.
- Outcome-contingent restrictions change conduct at higher tested audit intensities; targeting matters there relative to dose-matched random interruption.
- The rejected-alternatives rule and several-reader score add no effect because their implemented triggers do not activate.

The findings show why conformance and consequence must be tested separately. They do not establish field performance, prevalence, legal sufficiency, durable filing, reader diversity, privacy adequacy, or production cost. A NIST pilot could supply that missing evidence.

Requested outputs

A bounded NIST evaluation could produce:

1. a public problem statement and use-case boundary;
2. a draft DRCS Core and DRCS-EC crosswalk to the AI RMF and current agent standards work;
3. a normative schema and conformance corpus;
4. a reference producer, validator, and human-readable renderer;
5. a privacy and security threat model;
6. an interoperability exercise across independent implementations;
7. a field-pilot protocol with registered success and stop criteria;
8. a report of negative results and inactive tests; and

9. an open issue and revision ledger.

Sources and downloads

- [NIST AI Agent Standards Initiative](#)
- [NIST AI Risk Management Framework](#)
- [NIST AI RMF Playbook](#)
- [NIST Generative AI Profile, NIST AI 600-1](#)
- [Calvano/DRCS working paper, figures, methods, and downloads](#)
- [Calvano/DRCS v1.2 replication package](#)
- [DRCS-EC 1.0 technical profile](#)
- [DRCS Core 17-field public record guide](#)

Corrections and supersession

This is version 1.0. Corrections will be listed on the public request page and in the site corrections log. Any successor will identify the exact version it replaces, preserve this file, explain material changes, and link to the superseding version.

Feedback: <https://decisionaccounting.org/feedback/?topic=nist-open-request>