REPRODUCIBLE COVERAGE DECISIONS

RCD-1 — A Standard for Reproducibility, Traceability, and Accountability in Automated and Computer-Assisted Coverage Decision Systems

Working Draft, Version 0.9 · June 2026 · Status: Published for Public Comment

Convened and published for public comment by the **Healthcare Decision Integrity Working Group**, an open industry initiative convened by Avectic Corporation (Clarkston, Michigan). Participation in the Working Group is open to provider organizations, payers, technology vendors, professional associations, academic researchers, and regulators. This document is published under a Creative Commons Attribution 4.0 license. Comments: standard@avectic.com.

This standard is vendor-neutral. Conformance does not require any particular product, vendor, or commercial relationship. It specifies properties that any decision system operating in coverage workflows can be tested against, regardless of implementation.

1. Purpose and Scope

1.1 This standard defines requirements for systems that make, recommend, or materially assist coverage decisions in United States healthcare workflows, including prior authorization, utilization review, claims adjudication, downcoding, and medical-necessity determination (collectively, “coverage decisions”).

1.2 The standard exists because coverage decisions are consequential, governed by written policy, and subject to audit, appeal, regulation, and litigation. A coverage decision that cannot be reproduced and explained after the fact cannot be meaningfully appealed, audited, or defended. Recent enforcement actions, court orders compelling algorithm disclosure, and state certification statutes establish that decision systems in this domain will be examined; this standard defines what a system must be able to show when examined.

1.3 This standard applies to the decision layer of such systems. It does not prohibit the use of probabilistic technologies, including large language models, elsewhere in the workflow; it constrains where their outputs may carry decision authority (Clause 5.5).

1.4 This standard is designed to be adoptable by vendors as an engineering specification, by provider organizations and payers as procurement language, by auditors as a test protocol, and by regulators as a reference framework.

2. Normative Language

The key words MUST, MUST NOT, SHOULD, SHOULD NOT, and MAY in this document are to be interpreted as described in IETF RFC 2119. Requirements are individually numbered (R-, T-, P-, D-, U-, H-, M-) and are citable independently.

3. Terms and Definitions

Coverage decision. A determination, recommendation, or material input that affects whether a healthcare service, procedure, or product is authorized, paid, reduced, or denied under a governing policy.

Decision layer. The component(s) of a system in which the outcome of a coverage decision is determined from structured inputs and governing policy.

Understanding layer. The component(s) of a system that read, extract, structure, classify, or summarize source materials (clinical documents, records, communications) into the structured inputs consumed by the decision layer.

Decision-material field. Any output field whose value affects the decision outcome, including the determination itself, each criterion result, the policy revision applied, and the codes evaluated.

Policy corpus. The set of payer policies, medical-necessity criteria, coding rules, and regulatory requirements against which decisions are made.

Policy revision. An immutable, uniquely identified version of an element of the policy corpus, with effective dates.

Decision record. The durable artifact produced for each decision, sufficient to reconstruct the decision per Clause 5.4.

Re-execution. Running a previously decided case through the decision layer again, using the recorded inputs and the recorded policy revision(s).

Reproducibility. The property that re-execution yields outputs identical to the original on all decision-material fields.

4. Architectural Principle: Separation of Understanding and Decision

4.1 Reading clinical reality is inherently interpretive; deciding under written policy is not. Systems within scope MUST be architected such that the understanding layer and the decision layer are distinct, independently testable components, and such that the decision layer satisfies the requirements of Clause 5 regardless of how the understanding layer is implemented.

> *NOTE — This principle is technology-neutral. It permits any technique, including machine learning and large language models, in the understanding layer, and permits any implementation of the decision layer (rules engines, decision tables, executable policy logic) that satisfies Clause 5.*

5. Requirements

5.1 Reproducibility

R-1 The decision layer MUST be deterministic: given identical structured inputs and identical policy revision(s), it MUST produce identical values on all decision-material fields, on every execution, indefinitely.

R-2 Decision outcomes MUST NOT depend on sampling temperature, random seeds, model version drift, wall-clock time, execution order, infrastructure, or any other source of run-to-run variance.

R-3 The system MUST support re-execution of any historical decision on demand throughout the retention period (D-4), without reliance on external services whose behavior may have changed.

R-4 Where the governing policy has changed since the original decision, re-execution MUST apply the originally recorded policy revision, and MAY additionally report the outcome under current policy, clearly distinguished.

5.2 Traceability

T-1 Every decision-material field MUST be traceable to (a) the specific policy revision(s) applied, and (b) the specific input evidence relied upon, with stable references to both.

T-2 For each criterion evaluated, the decision record MUST state the criterion, its source in the policy corpus, the evidence considered, and the result.

T-3 Where a decision is adverse in whole or part, the record MUST contain the specific reason(s) at the criterion level, in terms consistent with the governing policy.

5.3 Policy Versioning

P-1 The policy corpus MUST be maintained under immutable version control. Each revision MUST carry a unique identifier, content hash, effective date range, and provenance (source document and authority).

P-2 Every decision MUST be bound at execution time to the policy revision(s) in force for the case, and the binding recorded.

P-3 Superseded revisions MUST remain retrievable and executable for the full retention period.

5.4 Decision Records

D-1 A decision record MUST be produced for every coverage decision, written at decision time, and immutable thereafter (append-only corrections permitted, with the original preserved).

D-2 The record MUST be sufficient for a qualified third party to answer, without access to the vendor: what was decided, on what evidence, under which policy revision, by which system version, with which human accountable, and to re-execute the decision per R-3.

D-3 Records MUST be available in both human-readable and machine-readable form.

D-4 Records, bound inputs, and policy revisions MUST be retained for no less than ten (10) years, or longer where governing law requires.

5.5 Probabilistic Components

U-1 Probabilistic components (including large language models) MAY be used in the understanding layer. Their outputs MUST be expressed as structured, human-verifiable assertions before consumption by the decision layer.

U-2 Probabilistic outputs MUST NOT directly determine any decision-material field. The decision layer MUST treat them as evidence inputs subject to its own deterministic evaluation.

U-3 For each decision, the record MUST identify the model(s) and version(s) used in the understanding layer and preserve their material outputs as part of the bound inputs.

U-4 Operators MUST monitor understanding-layer components for drift and error, with documented procedures for correction; corrections MUST NOT silently alter historical records.

> *NOTE — U-2 is the load-bearing clause of this standard. A system in which a probabilistic component selects, scores, or phrases the determination itself cannot satisfy R-1 and is out of conformance regardless of measured accuracy.*

5.6 Human Accountability

H-1 Every coverage decision MUST have an identified, qualified human accountable for it, recorded in the decision record. Automated processing MUST NOT be the sole basis for an adverse determination.

H-2 The accountable human MUST have access, at the time of decision, to the full trace required by Clause 5.2, presented intelligibly.

H-3 Override of a system result by the accountable human MUST be supported, recorded, and attributed.

5.7 Monitoring and Attestation

M-1 Operators MUST attest conformance annually, in the form of Annex A, executed by an officer of the organization.

M-2 Attestation MUST be supported by Re-Run Test results (Clause 7) executed within the attestation period.

M-3 Material nonconformities discovered between attestations MUST be documented, remediated, and disclosed in the next attestation.

6. Conformance Levels

Level	Name	Requirements
RCD-1 / L1	Traceable	Clauses 5.2, 5.3, 5.4, 5.6 in full. Decisions are explained and documented, but reproducibility is not claimed.
RCD-1 / L2	Reproducible	L1 plus Clauses 5.1 and 5.5 in full. The decision layer is deterministic and probabilistic components carry no decision authority.
RCD-1 / L3	Independently Verified	L2 plus the Re-Run Test (Clause 7) executed by an independent qualified party, and annual attestation (Clause 5.7) published or furnished to customers on request.

Levels exist to create an adoption path. L1 is achievable by well-run incumbent platforms today; L2 and L3 are architectural claims that can be verified, not asserted.

7. The Re-Run Test

7.1 Sample. From the attestation period, select the greater of 1,000 decisions or 5% of decision volume, by random sampling stratified across decision types and outcomes. The operator MUST NOT curate the sample.

7.2 Procedure. Re-execute each sampled decision per R-3/R-4 using recorded inputs and recorded policy revisions.

7.3 Pass criterion. Conformance at L2/L3 requires identity on all decision-material fields for **100% of re-executed decisions**. There is no statistical tolerance: a single divergent decision-material field is a nonconformity, because divergence is an architectural property, not a defect rate.

7.4 Disclosure boundary. The test verifies reproducibility without requiring publication of proprietary policy logic. Operators MAY conduct L3 verification under confidentiality with the independent party. Reproducibility, not source disclosure, is the conformance object.

> *NOTE — 7.4 answers the trade-secret objection: recent litigation demonstrates that secrecy does not survive discovery. A system that can be re-executed under NDA protects its operator far better than one that can only be described.*

8. Regulatory Alignment

RCD-1 is designed so that conformance produces, as a by-product, the evidence demanded by current and emerging law:

Authority	Demand	RCD-1 clause(s)
CMS-0057-F (eff. Jan 2026)	Specific reason for every denial; decision timeframes; transparency	T-3, D-1–D-3 (reasons are produced at criterion level, at decision time)
CMS MA guidance on AI in PA	Determinations must reflect the individual's circumstances, not group datasets	T-1, T-2, U-2 (evidence-bound, per-case evaluation)
Alabama SB 63 (eff. Oct 2026)	Annual certification; accuracy monitoring; no	M-1–M-3, U-4, Annex A

	group-dataset reliance	
Indiana HB 1271	No AI as sole basis for downcoding without professional review	H-1–H-3
California SB 1120 and successors	Clinician accountability for utilization decisions	H-1–H-3, D-2
Judicial discovery (e.g., 2026 disclosure order, D. Minn. litigation)	Algorithms used in denials are discoverable	Clause 7 (re-execution under NDA as the defensible posture)

9. Governance and Invitation

This Working Draft is published for a 90-day public comment period. The Working Group invites co-conveners from provider organizations, payers, professional associations (including the AMA, HFMA, MGMA, AAPC, and NICA), standards bodies, vendors, and patient advocates. Version 1.0 will be ratified by the assembled Working Group, and no single organization, including the convener, will hold veto over its content. The standard will remain free to implement, with no licensing fee, certification monopoly, or membership requirement.

Annex A — Attestation of Conformance (Template)

“[Organization] attests that, for the period [dates], the system [name, version] operated in conformance with RCD-1 at Level [1/2/3]; that the Re-Run Test described in Clause 7 was executed on [date] on a sample of [n] decisions with [result]; that all decision records required by Clause 5.4 were produced and retained; and that the undersigned officer is accountable for the accuracy of this attestation. [Name, title, date.]”