

# Appendix G — Governance and References

Status: v1.0 — Reading Edition (rev. 8) | Drafted: 2026-05-30

Source markdown: `./decision-provenance-standard-v1.0-appendix-G-governance-and-references.md`

---

Companion to the Decision Provenance Standard v1.0; tracks core revision rev. 8.

*Appendix G aggregates governance and reference material from multiple origin sections of the Standard; each fragment carries a back-pointer to its origin anchor. References here to the Standard's core sections (§1–§7, §10.7, §11.1/§11.2/§11.5) resolve against the core Reading Edition (rev. 8); references to Companions A, B, C resolve against those documents.*

**Lettering note.** Lettered **G**; Appendices A–F are reserved.

**Normative annex (NOT informative).** Sections §G.7.5, §G.7.6, and §G.7.7 below are the Standard's version-stability rules, relocated byte-identical from core §7.5–§7.7. **They remain normative and binding** on any Charter or release that claims conformance under core §7; they are aggregated here for length, not demoted to reference material. Core §7 retains a normative back-pointer to this annex (see core §7.5). The remainder of this Appendix G (§G.1, §G.11.3, §G.11.4, and Section 12 References) is informative/reference material.

---

## Appendix G — Governance and References

**Disclaimer pointer.** See top of the core Standard for the load-bearing UPL firewall and Jurisdiction Assumed. Section 12 of this appendix is a bibliography; it cites the frameworks the Standard converses with and does not characterize what those frameworks substantively require, certify, or attest. Substantive engagement with each framework lives in Companion A (Regulatory Cross-References).

---

### G.7 Version-Stability Rules (Normative Annex)

*Back-pointer: §G.7.5–§G.7.7 are the relocated, byte-identical core §7.5–§7.7. They are **NORMATIVE**. Core §7 carries a binding back-pointer to this annex; the R-005 non-break commitment, the classifier-version-increment rule, and the classification-ambiguity arbiter declared here govern conformance under core §7 exactly as if they appeared inline in §7. Cross-references in this annex to Section 4, Section 6 §6.2, Section 7 §7.3.3, and Section 11 resolve against the core Standard; intra-annex sibling references use the §G.7.x form.*

## **G.7.5 Classifier-Version Increments and the R-005 Minor-Release Non-Break Commitment**

The Mode-Drift Composed Mitigation sub-spec's Layer 1 (statistical detection) trains an independent Mode-2-trained classifier. The training corpus is curated and disjoint from any AI worker output the classifier will later classify. The classifier is versioned. Each output emits `corpus_id` + `corpus_version` + `classifier_version` provenance fields per the sub-spec FINAL, so a downstream auditor can verify independence at any point. The classifier version increments as the corpus expands and the classifier retrains. Section 7 must answer one question: does a classifier-version increment count as a Conformance Level 2 break under the Standard's R-005 minor-release non-break commitment? I decide it normatively in this Subsection.

### ***G.7.5.1 The decision***

**A classifier-version increment is NOT a Conformance Level 2 break in the general case.** Section 7 grades against the structural fact `no_silent_mode_drift_in_sample` at the audit moment the Layer 3 sample audit runs, not against the Layer 1 classifier's identity at any given moment. A Charter whose Layer 3 audit ran clean against classifier version N and whose audit re-runs clean against classifier version N+1 grades the same at Level 2; a Charter whose audit re-runs and produces a peer-confirmed drift finding at version N+1 does not grade at Level 2 until the affected records are re-dispatched per the demotion mechanism. The grade is a fact about the audit outcome, not a fact about which classifier version produced the audit's input flags.

This is the (A) position from the CPO sub-spec sign-off — classifier-version increments do not constitute Level 2 breaks in the general case. I adopt it in Section 7 normative text on three grounds.

**First, the conformance reporter does not bind to classifier version.** The Level 2 reporter reads `no_silent_mode_drift_in_sample` from the Layer 3 audit output. That output is itself a peer-reviewer-confirmed disposition (per Mode-Drift sub-spec FINAL Layer 3 §7.3.3 above), and the peer reviewer is the named firing authority for the signal. A classifier-version increment that produces additional Layer 1 hard flags routes those flags to Layer 3 peer review. The peer review either confirms drift (Level 2 grade flips to fail until re-dispatch) or rejects the flag (Level 2 grade unchanged). The Level 2 grade follows the peer review's disposition, not the classifier's version number. Binding the grade to classifier version would bind it to a layer that is not the named firing authority.

**Second, Layer 1's phased deployment is itself a classifier-version progression.** The sub-spec FINAL deploys Layer 1 in three phases — detection-only weeks 1-3 (Layer A + B corpus), detection-only weeks 4-6 (Layer A + B + C corpus), enforcement-mode week 7+ (full corpus, full firing authority). Each phase transition is in effect a classifier-version increment. If such increments were R-005 minor-release breaks in the general case, the sub-spec's own phased deployment would generate three Level 2 breaks per Charter in the first 7 weeks. That is precisely the wrong kind of structural noise. The phased deployment is the architecturally correct rollout (per the CPO sub-spec sign-off Call (ii)), and Section 7 grading must accommodate it without registering a break at every phase boundary.

**Third, the false-positive-rate-shift threshold (the (B) position) cannot be specified at this altitude without arbitrariness.** The (B) position lets classifier-version increments achieve Level 2 breaks when false-positive rate moves more than X%, which requires fixing X. The CPO sub-spec sign-off rejected (B): specifying X without a Section 7 baseline for acceptable signal stability would be theater. Section 7 is now drafted, and the structural baseline is clear. The Level 2 grade follows the Layer 3 audit's peer-reviewer disposition. A false-positive-rate shift in Layer 1 produces additional flags that route through Layer 3. Those flags either confirm drift (and the Charter's Level 2 grade flips to fail until re-dispatch, per the standard mechanism) or reject as false positive (and the Level 2 grade is unchanged). The structural mechanism handles the rate shift without a numeric threshold pinned to classifier version.

#### ***G.7.5.2 The narrow exception***

The general-case rule above admits one narrow structural exception. **A classifier-version increment that materially changes the corpus disjointness property is a Conformance Level 2 break.** The disjointness property is load-bearing for Layer 1 — the classifier's training corpus must be disjoint from any AI worker output it will later classify, per the Mode-Drift sub-spec FINAL Layer 1, "Classifier training corpus" subsection. A classifier retrained on a corpus that includes AI worker output it has previously classified, or that includes Mode-1-declared records that should have dispatched as Mode 2 (and that the prior classifier missed), is a classifier whose disjointness has lapsed. Layer 1's safety property depends on disjointness; without it, the classifier is reading its own output and the post-close population sampling collapses to a self-consistency check rather than an independent detection layer.

A classifier-version increment that lapses disjointness is not a routine retrain — it is a structural amendment to Layer 1's architecture. Section 4 normative-text amendment process governs structural amendments to Mode-Drift sub-spec architecture, and a deployer whose classifier version has lapsed disjointness must route the change through the amendment process before reasserting Level 2 conformance.

The narrow exception protects against one failure mode: a Charter's Level 2 grade looks stable across classifier-version increments while the underlying detection layer has silently lost its independence property. The exception is structural. It does not require a numeric threshold, and it does not turn every retrain into a conformance break. The trigger is the disjointness lapse, recorded as a sub-spec amendment, with the amendment's `closed_at` timestamp as the moment the Level 2 grade re-asserts.

#### ***G.7.5.3 The R-005 watch-item***

The R-005 minor-release non-break commitment is the Standard's commitment that minor releases of the Standard, the Mode-Drift sub-spec, the Section 4 dispatch architecture, and the conformance-signal vocabulary do not break a deployer's pre-release Conformance Level grade except in the explicit cases Section 11 enumerates. The classifier-version-vs-R-005 question is the most-likely-contested case of the commitment per the Phase 0.5.F.4 panel, and the CPO sub-spec sign-off routed it here. The decision in §G.7.5.1 and the narrow exception in §G.7.5.2 together resolve the question structurally: classifier-version increments are not breaks in the general case, and only disjointness-lapsing increments are breaks.

The R-005 commitment is itself subject to the single-arbiter-vs-3-of-6-sub-panel question §G.7.7 below resolves. A clarifying-language edge case — a deployer who reads §G.7.5.1 and reasonably arrives at a different interpretation of "the general case" — is the precise surface §G.7.7's classification-ambiguity arbiter resolves. The arbiter's role is to declare, in cases where the rule's application to a specific classifier-version increment is contested, whether the increment is a routine retrain or a disjointness lapse.

---

## **G.7.6 The Minor-Release Non-Break Commitment as Conformance Property**

The R-005 minor-release non-break commitment is a Standard-level property: minor releases of the Standard's authoring artifacts (Sections 2 through 8 normative text, the Section 4 dispatch architecture, the Mode-Drift sub-spec, the conformance-signal vocabulary, the Article 50 conformance language) do not break a deployer's pre-release Conformance Level grade in the general case. The commitment is structural — it lives at the Standard altitude rather than the Charter altitude — and Section 7 grades against it through the conformance-reporter's behavior across releases.

A Charter whose pre-release Conformance Level grade was Level 3 against Standard release N grades at Level 3 against Standard release N+1 (minor release) unless the release explicitly enumerates the Charter's grade as broken in Section 11. Section 11 enumerates the explicit breaks per minor release; absence from the enumeration is the commitment that the grade carries forward. The conformance reporter reads the deployer's release version and the Charter's grade against that version, and the Standard's commitment is that the release version increment does not silently mutate the grade.

The commitment exists because conformance-level grading is consumed by counsel, auditors, and the deployer's accountable personnel as input to substantive work. A grade that mutated silently across releases would force every consumer to re-validate every Charter on every Standard release. That would defeat the structural-input value the grade is designed to provide. The commitment is the Standard's promise that the structural-input value is stable across minor releases, and its load-bearing surface is Section 11.

The §G.7.5 classifier-version-vs-R-005 decision applies the commitment from first principles to the Mode-Drift sub-spec's Layer 1 versioning. Subsequent applications — to the Section 4 dispatch architecture, the conformance-signal vocabulary, the Article 50 conformance language — follow the same structure. Routine versioning does not break grades. Structural amendments to load-bearing properties (disjointness for Layer 1; the dispatch state machine's actor read/write boundaries in Section 4; the five required Article 50 fields for the conformance language) do break grades, and Section 11 enumerates the breaks per release.

A clarifying-language edge case in any of these applications is the surface §G.7.7 below resolves through the classification-ambiguity arbiter.

---

## G.7.7 Classification-Ambiguity Arbiter

R-005 mitigation per Phase 0.5.F.4 names the surface at which the most-likely-contested case of the minor-release non-break commitment is resolved. The surface is the classification-ambiguity arbiter — the named authority who declares, in cases where the Standard's rule application to a specific case is contested, whether the case is routine or a structural amendment. The Phase 0.5 sub-spec sign-off recorded that the arbiter is a single agent and routed the normative declaration to Section 7.

### G.7.7.1 *The decision*

**The classification-ambiguity arbiter is a single agent (CPO), not a 3-of-6 sub-panel.** The arbiter operates at the same altitude the Charter's `accountable_owner` operates: one named human, accountable for the call, recorded in the audit trail. The arbiter's role is to declare whether a contested case is a routine versioning event (no Conformance Level break) or a structural amendment (Conformance Level break per Section 11). The declaration is itself a decision and produces a decision record per Section 6 §6.2, dispatched under a Charter the Standard's authoring envelope governs (the Standard-of-Standards Charter, governed by the Standard's own authoring team).

Three grounds support the single-agent structure.

**First, the arbiter's role is binary.** The contested case is either a routine versioning event or a structural amendment, and the arbiter declares one or the other. A 3-of-6 sub-panel introduces deliberation overhead. That overhead suits substantive multi-disciplinary decisions (Phase 0.5.F panels for the most-likely-contested rule formulations) but is excessive for binary classification decisions. The single-agent structure preserves decision-time velocity at the altitude where the structural mechanism's input-stability property depends on timely calls.

**Second, the arbiter's call is auditable in the same surface as Charter `accountable_owner` calls.** Section 6 §6.2 specifies the decision-record schema. The arbiter's calls produce records under the schema, with the arbiter's identity in `accountable_owner`, the contested case in `decision_statement`, the rule application in `options_considered`, and the call's reasoning in the substantive content fields. A 3-of-6 sub-panel would produce records spanning six accountable-owner identities (or a single-named representative aggregating six positions). That is a structural mismatch with the rest of the Standard's `accountable_owner` discipline.

**Third, the arbiter's call is reviewable.** A deployer who disputes an arbiter call routes the dispute through the Section 4 normative-text amendment process — the Standard's amendment surface for structural questions, itself a multi-stakeholder review. The single-agent arbiter is the first-pass decision-maker; the multi-stakeholder amendment process is the review surface. Folding the review into the arbiter role conflates the two altitudes and produces neither the velocity of a single-agent call nor the rigor of a multi-stakeholder amendment.

### G.7.7.2 *Operating mechanics*

The arbiter operates against a fixed scope: contested cases of Section 11's break enumeration as applied to specific releases, contested cases of §G.7.5's classifier-version disjointness exception, and contested cases of

analogous rules elsewhere in the Standard's versioning policy. The arbiter does not opine on Charter-altitude conformance — the Charter's accountable\_owner and the conformance-level reporter cover that altitude — and does not opine on substantive matters that fall to qualified personnel.

A deployer who believes a specific case is contested routes the question to the arbiter through a structured request: the contested case, the arbiter's role on it, the deployer's view of the rule application, and the structural reasoning. The arbiter produces the call within the declared SLA (TBD per Section 11 authoring; the SLA is a Section 11 question, not a Section 7 question) and records the call as a decision record per Section 6 §6.2. The deployer then either accepts the call (it binds the deployer's Conformance Level grade against the contested case) or routes the dispute through the Section 4 normative-text amendment process.

The arbiter's role is bounded. It is the resolution surface for classification ambiguity in versioning, not a general escalation surface. Substantive ambiguities about the Charter's decision class, regulatory framework applicability, or conformance-reporter implementation are out of scope; those route through the Charter's escalation rule, counsel review, and a reference implementation's own escalation surface respectively. The arbiter's authority is structural and narrow; the Section 7 normative text fixes the scope so the surface does not drift into general-purpose escalation work that other surfaces are designed for.

---

## G.1 Related Work (origin: core §1.7)

*Back-pointer: this section is the relocated core §1.7 "Related Work." Core §1.3's navigation pointer to the section-to-audience Reading Guide resolves to core §1.8 (which stays in core), not to this section.*

This Standard operates on a substrate of prior academic, standards-body, and authored work. The following are the named lineages this Standard converses with. Each entry is a citation with a one-to-two-sentence placement. Substantive engagement with the underlying doctrines is the territory of qualified personnel reading those works on their own terms.

**Singh, J., Cobbe, J., and Norval, C. — "Decision Provenance: Harnessing Data Flow for Accountable Systems."** *IEEE Access*, 2018–2019. This is the academic root of the vocabulary the Standard operationalizes. The paper introduces "decision provenance" as a concept for accountable systems; this Standard formalizes the concept into a record format with an explicit lifecycle, a dispatch grammar, and a conformance grading layer at the executive-decision altitude. The structural innovations in this Standard — the sequential lifecycle gated on named human affirmation (§5), the intentional non-coverage of real-time telemetry (§5.3), and the conformance-level grading (§7) — operationalize what Singh, Cobbe, and Norval named conceptually.

**W3C — PROV-AGENT working group output on agent provenance metadata** (2025). The W3C PROV-AGENT effort addresses provenance metadata for agent-mediated artifacts in distributed systems. This Standard's Article 50 disclosure metadata block (§4.6) carries a related metadata burden for AI-authored content reaching natural persons. The two efforts are at different altitudes — PROV-AGENT addresses machine-to-machine provenance traces; this Standard addresses human-judgment decision records — and the relationship is

complementary, not derivative. Where a deployer's implementation surfaces both altitudes (a Mode 2 record whose underlying agent activity was traced under PROV-AGENT-conformant tooling), the trace can be referenced from the Decision Provenance Standard record under §6 as supporting evidence.

**AGENTSAFE framework for agentic AI safety** (arxiv, December 2025). AGENTSAFE addresses safety properties of agentic AI systems at the system-design altitude. This Standard addresses the human-judgment decision provenance altitude that sits above agent-system safety concerns. A deployer running agent-mediated workflows under AGENTSAFE-aware tooling produces decisions about that tooling at the executive altitude; those decisions are the Standard's territory. The two efforts are complementary at different altitudes.

*\*Trammell, J. — Chief Executive Operating System\*\* (2023). Trammell's Chief Executive Operating System is CEO-seat prior work on executive operating systems. The Standard's altitude — open record format for human-judgment decisions under CC-BY 4.0 — is distinct from the operating-system framing in Trammell's work. Different seats, different altitudes.\** The "Operating System" framing has independent lineage in Trammell's prior work; the Standard's "Decision Provenance" framing does not claim derivation from it.

**Gartner — Bimodal IT (Mode 1 / Mode 2 origins).** The Standard's Mode 1 / Mode 2 dispatch grammar is a technical term-of-art for **dispatch authorship** (human-led with AI enforcement vs. AI-led with human review) in this Standard's normative text, defined in Section 4. The terminology has independent lineage in the Bimodal IT discourse where Mode 1 / Mode 2 names two distinct IT delivery cadences. The Standard's use is in a distinct technical sense and does not claim derivation from Bimodal IT. Readers familiar with Bimodal IT should hold the two usages as independent; the Standard's Section 2.2.5 and 2.2.6 are the binding definitions for the Standard's use of the terms.

This Related Work section is the Standard's structural distinction-from-prior-art statement. It is not adversarial and does not claim that any cited work is incomplete, deficient, or superseded by this Standard. The Standard occupies a specific altitude — open record format for human-judgment decisions at named executive accountability — that none of the cited works occupies; the citations are placement, not contestation.

---

### G.11.3 Voluntary Adoption (origin: core §11.3)

*Back-pointer: this section is the relocated core §11.3 "Voluntary Adoption." Core §11.1 (Trademark Convention), §11.2 (Steward Governance), and §11.5 (Section Closure) remain in the core Standard.*

The Decision Provenance Standard is **voluntary infrastructure**. An accountable leader **installs** the Standard at the organization — a leader's act, not a regulator's mandate. This Section governs the Standard's adoption discipline. There is no certification body, no auditor pool, no accreditation regime, no plan to create any of the above.

The Standard's records **inform** the deployer's counsel and auditors when they prepare evidence, certifications, or attestations **without satisfying** any regulatory obligation; the obligation belongs to the obligated party, and the full non-claim set is at core §1.4.2.



**Scope of the records.** The Standard's records describe organizational decisions, named accountable seats, and the structural process by which decisions reach the affirmed state. Records MAY identify the natural person occupying an accountable seat at the moment of affirmation; records SHOULD link primarily to the role the seat names rather than to the natural person. Where a deployer uses the records as an input to performance management, hiring, promotion, retention, termination, or any other employment decision affecting a natural person, the deployer's use is subject to the deployer's local employment, data-protection, and automated-decision-making law (including, where applicable, NYC Local Law 144, EU GDPR Article 22, works-council consultation regimes, and equivalent state and federal employment-discrimination law). The Standard does not authorize, validate, or recommend any such use; the Standard's records produce a structural substrate, and the deployer's qualified personnel — counsel, HR leadership, works councils where applicable — govern whether and how that substrate is consumed for employment purposes. **The role-discipline rule applies to AI workers as well as natural persons: where a Mode 2 dispatch chain operates, the `drafting\_authority` field names the deployer-authorized role primarily, and MAY name model + version only where regulatory traceability requires it.**

**Per-altitude consent posture.** Voluntary at the deployer altitude is necessary but not sufficient for lower-altitude records. Individual-altitude records require separately-obtained, revocable, use-case-scoped consent from the affirmer (not bundled with the offer letter, not buried in an employee handbook update). The affirmer can withdraw consent and have their individual-altitude record stream stopped (not retroactively erased, but no further records added) without that withdrawal being treated as performance evidence in itself. Default scope at altitudes below executive is team-level aggregation. **Team-level aggregation numerical floor — deployer-DPIA-determined.** The numerical floor for "team-level aggregation" is determined by the deployer's qualified Data Protection Officer (or equivalent counsel role under the deployer's local regime) through a Data Protection Impact Assessment per the deployer's local data-protection law (GDPR Article 35 in EU/UK installations, and analogous DPIA-equivalent assessments in other jurisdictions per Companion A §A.11). The Standard does not specify the floor (a French CSE consultation typically lands on  $n \geq 10$  per CNIL k-anonymity guidance, UK ICO monitoring guidance points toward  $n \geq 5$  with k-anonymity considerations, and German Betriebsrat practice varies by establishment); the Standard requires that the floor be deployer-DPIA-determined and recorded in the Charter's `works_council_consultation_record` field per §3.1 (where a works-council jurisdiction applies) or in the deployer's HR-of-record DPIA artifact (where no works-council jurisdiction applies). A deployer whose "team-level aggregation" floor is undocumented at the Charter's `fields-completed` lifecycle state has departed from §G.11.3 conformance and SHALL NOT self-declare Conformance Level 2 or above per §7. The access-policy layer per §6.2.3.1 SHALL enforce the deployer's DPIA-determined floor at record-write time for any record at altitude: team-leader or below; aggregation below the deployer's documented floor SHALL be rejected at write time. Individual-altitude records exist if and only if the four design choices in the Standard's installation guidance (default-team-aggregation, affirmer-ownership, use-case scope-limit by Charter declaration per §3, jurisdiction-calibrated retention) are observed by the deployer. **The four design choices are not aspirational: per §6.2.3.1, the access-policy layer SHALL enforce each binding at record-write time, record-read time, and affirmation time, and a deployer whose access-policy layer does not perform these bindings SHALL NOT self-declare Conformance Level 2 or above.** *"Where a deployer's use case for individual-altitude records would bring the records into scope of NYC Local Law 144*



*(Automated Employment Decision Tools — AEDT), the deployer's AEDT obligations (bias audit, public summary, candidate notice) belong to the deployer as AEDT operator under Companion A §A.11; the §G.11.3 consent posture and the §6.2.3.1 access-policy binding are structural inputs to the deployer's AEDT readiness work and do not satisfy AEDT obligations." "Where the deployer's use case for individual-altitude records would constitute a decision based solely on automated processing producing legal or similarly significant effects on a natural person under GDPR Article 22, the controller's Article 22 obligations (explicit consent, human-in-the-loop guarantees, Article 22(3) safeguards) belong to the controller per Companion A §A.11. The §5.2 affirmation requirement (an affirmation MUST be an affirmative human act, not a passive signal) is the structural primitive that distinguishes Decision Provenance Standard records from solely-automated decision-making, but does not by itself discharge Article 22 obligations; the controller's qualified personnel govern the Article 22 analysis." "Where individual-altitude records inform an employment-adverse decision, the Fair Credit Reporting Act (15 U.S.C. §1681 et seq.) and its state analogs (including California's ICRAA and CCRAA) may impose pre-adverse-action notice, copy-of-report, summary-of-rights, and 30-day dispute-window obligations on the deployer as the user of the output; whether the records constitute a 'consumer report' under §1681a is an unsettled question currently in litigation. Those obligations belong to the deployer and its counsel and are not addressed, discharged, or satisfied by the Standard's records." "Where an affirmer or a natural person described in an individual-altitude record exercises a right of erasure under EU GDPR Article 17 (or an analogous erasure right under another regime per Companion A §A.bis), the §5.5 redaction-event record pattern operationalizes the erasure of the named fields from the operational data store as the deployer's counsel determines the right applies, while the archival record is retained or purged under the Article 17(3) ground the deployer's counsel identifies; the erasure determination and the lawful-basis analysis are the deployer's counsel's and Data Protection Officer's, and are not made, discharged, or satisfied by the Standard's records (see Companion A §A.2.bis)."*

**Named employment counsel of record — Charter sub-field at HR-altitude.** Where a Charter under §3 creates records at altitude: function-leader, altitude: team-leader, or altitude: individual-professional describing natural persons, the Charter SHALL carry the named\_employment\_counsel\_of\_record field at the Charter's fields-completed lifecycle state per §3.3. The Charter cannot reach fields-completed — and therefore cannot dispatch decisions claimable as conformant under §7 — unless this field is populated. The field records: the natural-person name and admission jurisdiction of the employment counsel engaged by the deployer for the Charter's scope; the engagement reference (internal counsel role title or external engagement letter pointer); and the date of the counsel's most recent attestation per the re-attestation requirement below. Where the Charter's scope crosses jurisdictional employment-law regimes, the field SHALL name one counsel per jurisdiction the Charter's records describe natural persons in; a single counsel attestation does NOT cover jurisdictions the named counsel is not admitted in. The works-council consultation record sub-field established at §3.1 is the primary record of consultation completion and is a precondition to the named counsel's attestation: the named counsel SHALL NOT attest to the Charter's lawful scope unless the works-council sub-field is populated where the local regime listed in §G.11.3 paragraph 4 applies; where the works-council sub-field's outcome\_class is consultation-pending-with-good-faith-progress or consultation-completed-without-agreement, the named counsel's attestation SHALL explicitly note the outcome class and SHALL NOT advance the Charter past fields-

completed until the named counsel has determined under the local regime that advancement is lawful on the recorded outcome class. The named employment counsel of record SHALL re-attest annually that the Charter's scope remains lawful under the named counsel's admission jurisdiction at the time of re-attestation. The annual re-attestation is itself an affirmed record under a Charter governing HR-altitude counsel re-attestations. A Charter whose `named_employment_counsel_of_record.last_attestation_date` is more than 365 days stale at any record-write event SHALL be flagged by the access-policy layer per §6.2.3.1; records authored under a stale-attestation Charter SHALL NOT be claimed as conformant under §7 for the period of staleness. The deployer's HR-of-record may cure the staleness by obtaining a fresh attestation; records authored after the cure date are conformant, but records authored during the stale period remain non-conformant regardless of cure. A deployer that attempts to install an HR-altitude Charter without this field, without the works-council consultation record sub-field where the local regime requires it, or without an in-date annual re-attestation has authored a non-conformant Charter; records authored under such a Charter cannot claim §7 conformance at any Level, and the access-policy layer per §6.2.3.1 SHALL reject record-write operations under such Charters from the moment the Charter is detected as un-scaffolded.

**U.S. employment-litigation discoverability and the "informs without satisfying" firewall.** Decision records authored under this Standard at any altitude that describes a natural person — typically records at altitude: function-leader, altitude: team-leader, or altitude: individual-professional, but extending to altitude: executive records where the named accountable owner or affirmer is a natural person whose role status is the litigation subject — MAY become discoverable in U.S. employment-litigation proceedings, including wrongful-termination claims, employment-discrimination claims under federal anti-discrimination frameworks, retaliation claims, and equivalent state-law causes of action. The records' discoverability is a function of state-of-employment civil-procedure rules and the litigation forum's discovery doctrines, not a function of the Standard's seal-hash discipline or the §5.1(3) seal's tamper-evidence properties (the seal is tamper-evident on a stored, access-controlled record per §2.2.18, not cryptographic immutability against all attack surfaces). The seal preserves tamper-evidence of what the deployer recorded and when; it does not exempt the records from discovery, and a deployer that installs the Standard expecting the seal to function as a litigation shield has misread §5. Where a deployer is subject to a litigation hold or anticipates U.S. employment litigation, the deployer's qualified personnel — employment counsel of record under the strengthened §G.11.3 paragraph immediately above, and litigation counsel where engaged — govern the records' production, redaction, and use as litigation inputs. The "informs without satisfying" firewall (per §1.4.2 and Companion A §A.0 lead) applies symmetrically to wrongful-termination defense postures: conformance to the Standard does NOT establish, support, or substantiate a defense against wrongful-termination, employment-discrimination, or related claims. A deployer's installation of the Standard, the self-declaration of any §7 Conformance Level, the existence of a `named_employment_counsel_of_record` field, and the population of any §3 use-case scope-limit declaration are NOT, individually or collectively, a substantive defense to any employment-law claim. Deployer compliance with the Standard's record-keeping discipline is not a substitute for compliance with applicable U.S. federal or state employment law. Deployers facing or anticipating U.S. employment litigation SHALL engage U.S. employment counsel admitted in the relevant state(s) BEFORE relying on Standard-conformant records as a defensive posture. The Standard produces audit-ready inputs the named counsel

consumes; the defensive theory, the privilege and work-product analysis, and the admissibility determination remain the named counsel's substantive work, not the Standard's structural property.

*"Deployers operating in jurisdictions with works-council co-determination or consultation regimes — including German Betriebsverfassungsgesetz §87 (Mitbestimmung at the establishment level on technical devices monitoring employee performance or behavior), French Code du travail Article L. 2312-26 (CSE consultation on the introduction of new technologies), Dutch Wet op de ondernemingsraden equivalent provisions, and analogous regimes elsewhere in the EU/UK — SHALL complete works-council pre-consultation before authoring a Charter that creates records at function-leader altitude or below describing natural persons. The Standard's records inform the deployer's pre-consultation work (the Charter's use-case scope-limit declaration per §3.1, the per-altitude consent posture in this §G.11.3, and the access-policy binding per §6.2.3.1 are the structural inputs the works council reads); they do not satisfy the consultation obligation. The consultation obligation is the deployer's, the works council's response is the works council's, and the resulting agreement (or its absence) governs whether and how the deployer authors the Charter."*

**Consent withdrawal — audit-trail event.** Where the affirmer issues a consent withdrawal under this section, the deployer's HR-of-record records the withdrawal in a new affirmed record at function-leader altitude per §6.2.3 (the `consent_withdrawal_event` field). The withdrawal-event record is authored at altitude: function-leader by default, MAY be authored at altitude: executive where the deployer's HR-of-record IAM policy assigns that altitude, and SHALL NOT be authored at altitude: individual-professional — the original (now-stopped) individual-altitude stream is exactly what the withdrawal event documents the cessation of. The affirmer selects the `withdrawal_reason` enum value (affirmer-discretion, affirmer-disengagement-from-deployer, affirmer-objection-to-scope, or affirmer-other-with-free-text); the deployer's HR-of-record records the affirmer's selection but SHALL NOT characterize the withdrawal on the affirmer's behalf and SHALL NOT override an affirmer-selected value. Where the affirmer declines to select a value, the default value SHALL be affirmer-discretion. Where the affirmer selects the free-text variant, the affirmer's text is recorded verbatim up to a deployer-configured maximum length (recommended default: 280 characters); the deployer's HR-of-record SHALL NOT append analytical commentary, summary, paraphrase, or characterization to the affirmer's text, and where the affirmer's text exceeds the maximum, the deployer's tooling SHALL truncate at the maximum without paraphrase or summary, and the truncation event is itself recorded in the withdrawal-event record's `revision_history`. The withdrawal-event record's HR-of-record affirmation links primarily to the HR-of-record role under the deployer's IAM policy (e.g., "HR-of-record under Charter HR-2026-001"); it MAY additionally name the natural person occupying the role at the moment of affirmation where the deployer's regulatory or audit posture requires natural-person traceability. The role-primary linkage applies the role-discipline rule of this §G.11.3 to the withdrawal-event record's affirmer surface; it does not modify the §5.1(3) affirmation requirement that an affirmation is an affirmative human act by a named human. The prior individual-altitude record stream is sealed at its prior `seal_hash` per §5.1(3) and is not mutated, deleted, or retroactively erased by the withdrawal event; the withdrawal-event record is a new, separately-sealed structural fact recording the cessation of the prior stream. The withdrawal-event record carries no field that re-asserts substantive performance evidence, behavioral commentary, or evaluative content about the affirmer; the record's purpose is structural-fact recording — *that* the withdrawal occurred,

when it occurred, *that* HR-of-record affirmed the deployer's honoring of it — not characterization. The Charter under which the withdrawal-event record is authored SHALL explicitly enumerate "consent-withdrawal-event-recording" in its use-case scope-limit declaration per §3.1; where the Charter does not enumerate this use case, the withdrawal-event record cannot be authored under that Charter, and the deployer's HR-of-record authors a Charter or Charter-amendment that does enumerate it before the withdrawal event is recorded. The access-policy layer SHALL reject withdrawal-event records authored by any principal other than HR-of-record under the Charter's enumerated scope, and SHALL reject affirmation events on the withdrawal-event record by any principal other than the HR-of-record affirmer named in `acknowledgement_by_hr_of_record`. **Cessation NOT gated on Charter amendment.** The cessation of the affirmer's individual-altitude record stream SHALL NOT be gated on the Charter's enumeration of consent-withdrawal-event-recording. Where the Charter does not enumerate this use case at the moment the affirmer issues a withdrawal, the deployer's HR-of-record honors the withdrawal immediately under the deployer's local-law obligation (GDPR Article 7(3) in EU/UK installations, California CCPA right-to-delete equivalent, Israel PPL 5741-1981 equivalent, and analogous primary obligations elsewhere per Companion A §A.11): the affirmer's individual-altitude record stream stops at the moment of withdrawal, no further records are added under the affirmer's prior consent, and the prior records are sealed and retained per §5.1(3) without retroactive erasure. The structural-fact recording of the withdrawal event in a new affirmed record follows once the deployer's HR-of-record authors a Charter or Charter-amendment that enumerates consent-withdrawal-event-recording per §3.1; the recording obligation is the Standard's structural primitive, the cessation obligation is the deployer's primary legal obligation, and the two obligations operate on independent timelines. A deployer whose Charter has not enumerated the use case at the moment of first withdrawal SHALL document the cessation event in an interim HR-of-record artifact (e.g., the deployer's standard data-subject-rights tracking system) and SHALL author the Charter or Charter-amendment that enumerates the use case before the second withdrawal event arrives; recurring receipt of withdrawal events without Charter enumeration is a deployer compliance gap, not a Standard-permitted operating posture. The `consent_posture.withdrawal_state` enum value `withdrawn-stream-stopped` (§6.2.3) is operationally implemented through the redaction-event record pattern (§5.5 and §6.2.3): the withdrawal-event record described in this paragraph IS a `record_type: redaction_event` record whose `redaction_basis` is `counterparty_contractual_obligation` or the relevant jurisdictional erasure basis (e.g., `gdpr_article_17`, `ccpa_right_to_delete`, `israeli_privacy_law_erasure`) per the deployer's counsel's determination. The §6.2.3.2 access-policy binding on redaction events governs the withdrawal-event record's affirmation, and the §7 Level 2 signal `every_redaction_event_carries_operational_store_deletion_attestation` reads the withdrawal-event record's `operational_store_deletion_attestation` field at sample-level audit events.

The voluntary-adoption discipline carries three load-bearing properties:

**Active leader verb: install / installation.** A CPO **installs** the Standard at the seat. The verb is the Standard's chosen register for what a deployer does with it. **The Standard is never "complied with" or "audited against" or "Standard-mandated";** it is installed. A deployer who reads the Standard as a compliance regime has misread the Standard.

**Self-declared Conformance Levels.** An organization **self-declares** Conformance Level X against the Standard per §7. The Standard's Steward (per §11.2) does NOT certify, accredit, audit, stamp, or grade any organization. There is no certification body, no auditor pool, no plan to create either. **Etsion Brands does not certify deployer organizations.** Per IP Counsel R4 Finding 6 and the locked language at §7's lead paragraph, self-declared / non-certified usage does NOT trigger Lanham Act §1054 certification-mark obligations — that's load-bearing for the trademark posture established in §11.1.

**Vendors do NOT certify against the Standard.** A vendor MAY build "Standard-aware tooling" or "Standard-conformant tooling" as the vendor's own product, MAY market the tooling using the Standard's name per the permitted trademark uses in §11.1, and MAY build Standard-conformance into the tooling's runtime behavior. A vendor MAY NOT stamp a customer organization as "Standard-compliant," "Decision Provenance Standard™-certified," or any equivalent third-party-certification framing. **Stamp the tool, not the org.**

The Standard's records **inform** regulatory cross-references (NIST AI RMF, ISO/IEC 42001, EU AI Act, and equivalent frameworks per Companion A) **without satisfying** them. They are **audit-ready decision provenance** (per §1.4 and §2.2.7), never **legal evidence, compliance certification, or regulatory substitute**. Always pair "inform" with "without satisfying" on regulatory surfaces. Any "the Standard satisfies [framework]" framing is corrected on first sight, per the §2.3 vocabulary discipline.

**NIST-RMF-adjacent posture.** The Standard is positioned as voluntary infrastructure, parallel to NIST AI RMF's voluntary-framework posture. NIST AI RMF is a voluntary U.S. framework that informs AI risk management work without satisfying any specific regulatory obligation; this Standard occupies the same voluntary posture at the executive-decision-provenance altitude. Neither this Standard nor NIST AI RMF is a regime, a mandate, a certification scheme, or an audit obligation.

---

## G.11.4 Recognition of Self-Declaring Adopters (origin: core §11.4)

*Back-pointer: this section is the relocated core §11.4 "Recognition of Self-Declaring Adopters."*

Where an adopting organization self-declares a Conformance Level (§7) and publishes that self-declaration on a publicly accessible surface, the Steward MAY recognize the self-declaration on a community page hosted at the canonical Standard URL. **Recognition is an acknowledgment of the org's self-declaration; recognition is NOT certification, audit, or grading by the Steward.** The community page reads as a list of organizations that have publicly self-declared conformance, not as a list of organizations the Steward has audited.

**Recognition mechanics.** A deployer wishing to be recognized publishes:

1. The deployer's organization name and a public URL where the self-declaration can be read by a third party
2. The Charter or set of Charters covered by the self-declaration (named per the deployer's own naming convention; the Steward does not validate Charter contents)
3. The Conformance Level claimed (1, 2, or 3 per §7), with the date of self-declaration

4. A statement that the self-declaration is the deployer's, not the Steward's, and that the Steward has not audited the underlying records

The Steward reviews the public-URL contents for the structural elements above (name, URL, Charters covered, Level claimed, self-declaration statement) and recognizes the publication on the community page. The Steward's review is **structural, not substantive**: the Steward confirms that the publication contains the four elements, is hosted at a stable URL, and uses the Standard's vocabulary without misrepresentation. The Steward does NOT confirm that the underlying records actually support the Level claimed; that's the deployer's self-declaration to make.

**Structural-correctness review.** The Steward's review of a recognition publication is limited to the four structural elements enumerated above (deployer name, public URL, Charters covered, Conformance Level claimed and dated, self-declaration statement). The Steward MAY decline to list a publication that is missing one or more of the four structural elements, that is hosted at a non-resolving or non-stable URL, or that uses the Standard's name in a manner outside the §11.1 permitted trademark uses (for example, a publication that frames the listing as third-party certification by the Steward, or that brands a consulting service or training program in a manner suggesting Steward authorization per §11.1). The Steward does NOT review the underlying implementation, does NOT evaluate whether the deployer's records support the Conformance Level claimed, and does NOT opine on whether the self-declaration is substantively accurate; the substantive accuracy of the self-declaration is the deployer's responsibility per §G.11.3 and §7. Decisions to decline a listing are themselves recorded as decisions at the Steward altitude, in the Steward's own decision register; the Steward operates the Standard's authoring under a Charter consistent with §3 of the Standard.

**Founding Confirmer recognition (post-launch).** Where the Steward recognizes early adopters who self-declared organically in the first months after the Standard's publication, the framing is "*we recognized their self-declaration*" — never "*we certified them*" — per the §G.11.3 voluntary-adoption discipline.

---

## Section 12 — References (origin: core §12)

*Back-pointer: this section is the relocated core §12 "References."*

**Use of the Standard.** See top of the core Standard. Section 12 is a bibliography; it cites the frameworks the Standard converses with and does not characterize what those frameworks substantively require, certify, or attest. Substantive engagement with each framework lives in Companion A (Regulatory Cross-References).

**Jurisdiction Assumed:** as declared at the top of the core Standard.

---



## 12.1 Purpose

Section 12 is the bibliography of named regulatory frameworks, prior art, and source materials the Standard cites. It enumerates each citation with sufficient precision — issuing body, version or year, article or clause where applicable, and a stable pointer URL — for a reader, an audit chair, or a reader's counsel to locate the underlying source and confirm currency.

A reader's counsel turns to Section 12 to validate that the Standard engages real frameworks in the form those frameworks actually exist. Section 12 also makes the Standard's "input to regulatory work" framing legible. By listing the frameworks the Standard converses with, and by citing them with discipline without characterizing what they substantively require, Section 12 makes clear what the Standard is *not* replacing. The bibliography is a bibliography. It is not a regulatory cross-walk, a compliance crosswalk, or a substitution map. Companion A holds the substantive cross-reference territory and is authored under the verification chain named in §12.5 below.

Section 12 follows three discipline rules:

- Entries CITE; they do NOT CHARACTERIZE.
- Frameworks the Standard does not engage are not listed.
- Citations are version-specific.

---

## 12.2 Regulatory frameworks

The Standard's Companion A (Regulatory Cross-References) maps the Standard's structural requirements onto the named regulatory frameworks below. Section 12 cites those frameworks; Companion A engages them. The frameworks are grouped in three blocks: AI-specific frameworks (the AI/ISO trio sub-verified by Privacy Counsel); traditional internal-control and assurance frameworks (sub-verified by the General Counsel); and case law (sub-verified by the General Counsel). Each entry follows the citation form: issuing body — full title — version/year — article/clause where the Standard cites it — pointer URL — one-line neutral descriptor.

### 12.2.1 AI-specific frameworks (the AI/ISO trio)

**European Union — Regulation (EU) 2024/1689 (the EU AI Act).**

- Issuing body: European Parliament and Council of the European Union.
- Full title: Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending certain Union legislative acts (Artificial Intelligence Act).
- Version/year: 2024 (adopted 13 June 2024; published in the Official Journal 12 July 2024).
- Articles cited by the Standard: Article 14 (Human Oversight); Article 17 (Quality Management System); Article 50 (Transparency Obligations for Providers and Deployers of Certain AI Systems).



- Pointer URL: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>.
- Neutral descriptor: European Union regulation harmonising rules for the development, placing on the market, and use of artificial-intelligence systems.

**United States — National Institute of Standards and Technology, AI Risk Management Framework (AI RMF 1.0).**

- Issuing body: U.S. National Institute of Standards and Technology (NIST).
- Full title: Artificial Intelligence Risk Management Framework (AI RMF 1.0).
- Version/year: 1.0 (January 2023).
- Sub-functions cited by the Standard: the Manage function generally; specifically Manage 4.1.
- Pointer URL: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.
- Neutral descriptor: U.S. voluntary framework structuring how organizations identify, measure, and manage risks associated with artificial-intelligence systems.

**International — ISO/IEC 42001:2023.**

- Issuing body: International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC).
- Full title: ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system — Requirements (rendered per the ISO catalog page at the pointer URL below; readers verify the live rendering before citing in publication).
- Version/year: First edition, 2023.
- Pointer URL: <https://www.iso.org/standard/81230.html>.
- Neutral descriptor: International management-system standard specifying requirements for establishing, implementing, maintaining, and continually improving an artificial-intelligence management system within an organization.

## **12.2.2 Traditional internal-control and assurance frameworks**

**United States — Committee of Sponsoring Organizations of the Treadway Commission (COSO) — Internal Control – Integrated Framework (2013).**

- Issuing body: Committee of Sponsoring Organizations of the Treadway Commission.
- Full title: Internal Control – Integrated Framework.
- Version/year: 2013 edition (the framework was first issued in 1992 and updated in 2013; the 2013 edition is the version cited by the Standard).
- Pointer URL: <https://www.coso.org/guidance-on-ic>.
- Neutral descriptor: U.S. internal-control framework articulating five components and seventeen principles used by organizations to design, implement, and assess internal control over financial reporting and over operations and compliance.

**United States — Sarbanes-Oxley Act of 2002, Section 404.**

- Issuing body: 107th United States Congress.
- Full title: Sarbanes-Oxley Act of 2002, Public Law 107-204, Section 404 (Management Assessment of Internal Controls).
- Version/year: Enacted 30 July 2002; codified at 15 U.S.C. § 7262.
- Pointer URL: <https://www.govinfo.gov/content/pkg/PLAW-107publ204/pdf/PLAW-107publ204.pdf>.
- Neutral descriptor: U.S. federal statute, Section 404 of which addresses management assessment of internal control over financial reporting and the related external auditor attestation.

**United States — American Institute of Certified Public Accountants (AICPA) — SOC 2® Type II reports under the Trust Services Criteria.**

- Issuing body: American Institute of Certified Public Accountants (AICPA), Assurance Services Executive Committee.
- Full title: Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy (TSP Section 120, 2017 Trust Services Criteria, with revisions). SOC 2 Type II refers to a service auditor's report on the suitability of design and the operating effectiveness of controls over a defined period.
- Version/year cited: 2017 Trust Services Criteria (with 2022 revisions to Points of Focus). Readers verify currency at the AICPA pointer below before relying on a specific edition.
- Pointer URL: <https://www.aicpa-cima.com/topic/audit-assurance/audit-and-assurance-greater-than-soc-2>.
- Neutral descriptor: U.S. assurance reporting framework under which an independent CPA evaluates and reports on the design and operating effectiveness of a service organization's controls relevant to security, availability, processing integrity, confidentiality, or privacy.

### 12.2.3 Case law

*Case-law citations below follow Section 12's CITE-not-CHARACTERIZE rule (see §12.1). Each entry gives the formal citation, court, and year, plus a one-line neutral descriptor naming the subject domain. What these decisions hold, how they are interpreted, and how they apply to a specific organization or implementation is the territory of Companion A (Regulatory Cross-References) and, ultimately, of counsel admitted in the relevant jurisdiction.*

**In re Caremark International Inc. Derivative Litigation.**

- Court: Delaware Court of Chancery.
- Year: 1996.
- Citation pointer: 698 A.2d 959 (Del. Ch. 1996).
- Pointer URL: <https://courts.delaware.gov/opinions/>.
- Neutral descriptor: Delaware Court of Chancery decision in the subject domain of board-of-directors oversight duties under Delaware corporate law.

**Marchand v. Barnhill.**

- Court: Supreme Court of Delaware.

- Year: 2019.
  - Citation pointer: 212 A.3d 805 (Del. 2019).
  - Pointer URL: <https://courts.delaware.gov/opinions/>.
  - Neutral descriptor: Supreme Court of Delaware decision in the subject domain of board-of-directors oversight duties, decided in 2019.
- 

## 12.2.4 Normative-keyword foundation

**Bradner, S. — Key words for use in RFCs to Indicate Requirement Levels.**

- Issuing body: Internet Engineering Task Force (IETF), Network Working Group
  - Document type: Best Current Practice (BCP 14) / Request for Comments (RFC)
  - RFC number: 2119
  - Version/year: March 1997 (no successor; updated by RFC 8174 in May 2017 for ambiguity-handling on lowercase keywords)
  - Pointer URL: <https://www.rfc-editor.org/rfc/rfc2119>
  - Neutral descriptor: IETF Best Current Practice document defining the uppercase normative keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL as used in technical specifications. The Decision Provenance Standard™ adopts RFC 2119 as its normative-keyword foundation per the Normative Keywords block at the top of this Standard.
- 

## 12.3 Prior art and source materials

Prior art listed here is the published material the Standard builds on, names, or otherwise treats as a source. Inclusion here is a bibliography fact, not an endorsement of any author's substantive claims, and is not a derivative-work declaration. §12.3 only enumerates.

### 12.3.1 Origin / Acknowledgments

The Charter mechanism, Mode taxonomy, and decision-record discipline formalized in this Standard were first developed in the author's prior work on product-organization decision systems. That work is acknowledged here as origin; it is not required reading and the Standard's normative content is self-contained in Sections 1–11 and its companions.

### 12.3.2 Phase 0.5 preflight materials internal to this Standard's authoring

These are internal authoring substrates that Section 12 cites because Sections 2–9 reference them. They are not external prior art; they are shared constraints across the Standard's sections.

**Language Discipline Cheat Sheet (Phase 0.5.A).**

- Version: v1, 2026-04-28.
- Pointer: `working/preflight/language-discipline-cheat-sheet.md`.
- Neutral descriptor: Internal authoring constraint document specifying terms-to-use and terms-to-avoid for every section and every cross-reference; load-bearing for the Standard's "audit-ready provenance" framing.

#### **Standard ↔ Reference-Implementation Specification (Phase 0.5.B).**

- Version: v1, 2026-04-28.
- Pointer: `working/preflight/standard-reference-implementation-spec.md`.
- Neutral descriptor: Internal authoring constraint document defining the relationship between the Standard's Section 4 normative text and a reference implementation's runnable surfaces (Charter state model per Section 3, Article 50 disclosure metadata schema per Section 4 §4.6, Mode 1/2 dispatch state machine per Section 4, conformance-signal vocabulary per Section 7).

#### **Message Architecture Lock v1 (Phase 0.5.C).**

- Pointer: `working/preflight/message-architecture-lock-v1.md`.
- Neutral descriptor: Internal authoring constraint document locking the canonical Mode 1 / Mode 2 names and the accessible-alias pair, governing where each may appear.

### **12.3.3 Related Work Citations**

The following citations support the Related Work paragraph at §G.1 (origin §1.7). Each entry follows the same citation discipline as the regulatory frameworks in §12.2: issuing body, full title, version/year, pointer URL, neutral descriptor.

#### **Singh, Cobbe, and Norval — "Decision Provenance: Harnessing Data Flow for Accountable Systems."**

- Authors: J. Singh, J. Cobbe, C. Norval
- Publication: *IEEE Access*
- Years: 2018–2019
- Pointer URL: <https://ieeexplore.ieee.org/document/8395145>
- Neutral descriptor: Academic publication introducing "decision provenance" as a concept for accountable systems; the academic root of the vocabulary the Standard operationalizes.

#### **W3C — PROV-AGENT working group output.**

- Issuing body: World Wide Web Consortium (W3C)
- Full title: PROV-AGENT — Provenance for Agent-Mediated Artifacts (working group output)
- Year: 2025
- Pointer URL: <https://www.w3.org/community/prov-agent/>
- Neutral descriptor: W3C working group output on provenance metadata for agent-mediated artifacts in distributed systems.

### **AGENTSAFE — framework for agentic AI safety.**

- Publication: arxiv preprint
- Date: December 2025
- Pointer URL: arxiv preprint at <https://arxiv.org/> — the specific arxiv identifier for the AGENTSAFE framework as cited in §G.1 will be filled in at v1.0 publication; readers verifying citation currency between the v1.0 publication date and any subsequent rev. should consult the arxiv search interface for AGENTSAFE.
- Neutral descriptor: Framework for safety properties of agentic AI systems at the system-design altitude.

*\*Trammell, J. — Chief Executive Operating System.\**

- Author: Joel Trammell
- Year: 2023
- Pointer URL: *(book publisher URL to be confirmed at publication)*
- Neutral descriptor: CEO-seat prior work on executive operating systems; the Standard's altitude — open record format for human-judgment decisions — is distinct.

### **Gartner — Bimodal IT (Mode 1 / Mode 2 origins).**

- Issuing body: Gartner, Inc.
- Years of origination: ~2014 (with subsequent Gartner research publications)
- Pointer URL: <https://www.gartner.com/en/information-technology/glossary/bimodal>
- Neutral descriptor: Gartner research framing of two distinct IT delivery cadences as "Mode 1" and "Mode 2"; the term "Mode 1 / Mode 2" in this Standard is a distinct technical use at the dispatch-authorship altitude (per §2.2.5 and §2.2.6) and does not claim derivation from Bimodal IT.

---

## **12.4 Tooling and reference implementations**

§12.4 records that a reference implementation of the Standard's structural requirements exists. A reference implementation is not a conformance-certifying body, and using one does not by itself produce a Standard-conformant Charter. A reference implementation structures the inputs; conformance against any framework remains a determination by the deployer's qualified personnel.

### **Reference implementation.**

- License: MIT License (the reference implementation is code; the Standard's own text is separately licensed under Creative Commons Attribution 4.0 International (CC-BY 4.0)).
- Neutral descriptor: An open-source skill bundle implementing the structural requirements of the Standard's Sections 3 (Charter mechanism), 5 (decision-record schema), and 6 (conformance levels) as runnable skills. A reference implementation is not a substitute for the Standard's normative text and does not declare conformance with any external regulatory framework.

---

## 12.5 Versioning note and verification chain

### 12.5.1 Versioning note

Citations in §12.2 through §12.4 are version-specific. Frameworks evolve: regulations are amended, standards are revised, case law is interpreted by subsequent decisions, and reference implementations release new versions. A reader relying on any citation in Section 12 verifies currency at the issuing body's pointer URL before treating that citation as current. A bibliography entry that was accurate when this Standard was published can become stale. The entry is a starting point for the reader's verification, not a snapshot of authoritative current state.

Note on NIST AI RMF version scope: as of this draft, NIST has published a Generative AI Profile companion (NIST AI 600-1, July 2024) to the AI RMF 1.0. The Generative AI Profile is a companion document, not a successor version; AI RMF 1.0 (NIST AI 100-1, January 2023) remains the operative framework version Companion A engages with.

Where a citation in §12.2–§12.4 is marked [TBD], the authoring agent could not confirm a specific version, edition, or revision identifier with sufficient confidence to lock it. [TBD] markers make the gap visible to verifiers in the chain below; they are not placeholders to be filled silently. Each [TBD] is resolved at verification or escalated to the named verifier.

### 12.5.2 Verification chain

Section 12 is the bibliography backbone of the Standard. The substantive accuracy of each citation block is verified by the domain owner accountable for that block. Privacy Counsel verifies the AI-specific frameworks (EU AI Act Articles 14/17/50; NIST AI RMF 1.0 including Manage 4.1; ISO/IEC 42001:2023). The General Counsel verifies the traditional internal-control and assurance frameworks (COSO 2013; SOX § 404; SOC 2 Type II / Trust Services Criteria) and the case law (*In re Caremark*; *Marchand v. Barnhill*). The Compliance Officer verifies the prior-art and tooling entries, with the Chief Architect consulted on the reference-implementation entry. The CPO verifies architectural fidelity — whether Section 12 cites what the Standard's sections actually reference.

The verification chain is documented here so a reader of the Standard can see who validated each citation block. The Standard's reliability comes from the chain, not from any single agent's authorship.