Companion to the Decision Provenance Standard v1.2; tracks core revision rev. 10.

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. 10); 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 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 §G.11.3 restates some rules with requirement wording (SHALL, MUST or "requires"). Where a core section (for example §3.1, §5, §6.2, §6.2.3.1, §6.2.3.2, §7 or §11.1) states the same rule, the core section is the binding one. A requirement in §G.11.3 that no core section states is a recommendation and reads as SHOULD.

Appendix G — Governance and References

> **Disclaimer pointer.** See top of the core Standard for the not-legal-advice notice 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 core §7.5–§7.7. They are NORMATIVE. Core §7 carries a binding back-pointer to this annex; the minor-release 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 Minor-Release Non-Break Commitment

The Mode-Drift Composed Mitigation's Layer 1 (statistical detection, Section 4 §4.8.1) 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, 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 minor-release non-break commitment? This Subsection decides it normatively.

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.

Classifier-version increments do not constitute Level 2 breaks in the general case. Three grounds support this rule.

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 Layer 3, Section 4 §4.8.1, and §7.3.3), 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. Section 4 §4.8.1 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 minor-release breaks in the general case, the Standard'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, and Section 7 grading must accommodate it without registering a break at every phase boundary.

Third, the false-positive-rate-shift threshold cannot be specified at this altitude without arbitrariness. The false-positive-rate-shift threshold alternative lets classifier-version increments achieve Level 2 breaks when false-positive rate moves more than X%, which requires fixing X. Specifying X without a Section 7 baseline for acceptable signal stability would be theater. 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 Section 4 §4.8.1. 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. A deployer whose classifier version has lapsed disjointness records the lapse in a record under its own Charter and does not reassert Level 2 conformance until that record closes.

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 in that record, with the record's `closed_at` timestamp as the moment the Level 2 grade re-asserts.

G.7.5.3 The non-break watch-item

The minor-release non-break commitment is the Standard's commitment that minor releases of the Standard, the Mode-Drift Composed Mitigation (Section 4 §4.8), the Section 4 dispatch architecture, and the conformance-signal vocabulary do not break a deployer's pre-release Conformance Level grade. A change that would break one is a structural amendment and ships only in a major release (§G.7.7.1). The classifier-version question is the most-likely-contested case of the commitment. 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 non-break 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 clarify how the rule reads; whether a specific deployer's classifier-version increment is a routine retrain or a disjointness lapse is for the deployer to determine and record.

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

The 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 Composed Mitigation (Section 4 §4.8), the conformance-signal vocabulary, the disclosure-block requirement in Section 4 §4.6) 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 Level that a Charter could declare against Standard release N still holds against minor release N+1. A change that would break it is a structural amendment and ships only in a major release

(§G.7.7.1), whose release notes list what it breaks. 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 the release notes of each major release are where any break is recorded.

The §G.7.5 classifier-version decision applies the commitment from first principles to the Mode-Drift Composed Mitigation's Layer 1 versioning. Subsequent applications — to the Section 4 dispatch architecture, the conformance-signal vocabulary, the disclosure-block requirement — follow the same structure. Routine versioning does not break grades. Structural amendments to core properties (disjointness for Layer 1; the dispatch state machine's actor read/write boundaries in Section 4; the five required fields for the disclosure block) do break grades; they ship only in a major release, whose release notes list the breaks.

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

This Subsection 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 Steward decides, and records publicly, whether a change to the Standard is a patch, minor or major release. The call concerns the Standard's text. It does not grade, confirm or change any organization's self-declared Level.

G.7.7.1 The decision

The classification-ambiguity arbiter is the Steward (§11.2), a single named person, 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 change to the Standard is a routine versioning event (a patch or minor release) or a structural amendment (a major release, whose release notes list what it breaks). The declaration is itself a decision and produces a decision record per Section 6 §6.2,

dispatched under the Charter that governs authoring of the Standard (published in the repository's `governance/` folder).

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 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. Anyone who disputes an arbiter call may raise it as an issue or pull request under GOVERNANCE.md; the Steward decides the dispute and records the decision publicly.

G.7.7.2 Operating mechanics

The arbiter operates against a fixed scope: contested cases of the major-release rule as applied to specific releases, the reading 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, which are for the deployer to determine.

Anyone may raise a classification question as an issue. The Steward aims to answer within 30 days of the issue being raised. No answer is not a ruling either way. The Steward records the call as a decision record per Section 6 §6.2. Anyone who disputes the call may raise the dispute as an issue or pull request under GOVERNANCE.md.

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, the deployer's own determination, and

the implementer'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, industry, 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.

PROV-AGENT (Souza et al., 2025). A research paper that proposes a provenance model, built on W3C PROV, for tracking the interactions of AI agents in automated workflows (§12.3.3). This Standard's disclosure block (§4.6) records related information about AI-drafted content. The two work at different altitudes — PROV-AGENT traces what agents did inside a workflow; this Standard records human-judgment decisions — and the relationship is complementary, not derivative. Where a deployer's implementation has both (a Mode 2 record whose underlying agent activity was traced with PROV-AGENT), the trace can be referenced from the Decision Provenance Standard record under §6.

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; whether any obligation applies, and to whom, is for the deployer to determine, 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, which employment, data-protection and automated-decision-making laws apply to that use (for example NYC Local Law 144, EU GDPR Article 22, works-council consultation regimes, and state and federal employment-discrimination law) is for the deployer to determine (Companion A lists some citations and is not complete). The Standard does not authorize, validate, or recommend any such use; whether and how the records are used for employment purposes is for the deployer to determine. **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 consent gate is the Standard's own voluntary gate. It makes no claim about the legal basis for processing. 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, for example through a Data Protection Impact Assessment or equivalent assessment under whatever law applies to the deployer (Companion A lists some citations and is not complete). The Standard does not specify the floor; the deployer determines it and records it. The Standard recommends that the floor be determined by the deployer 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 may bring the records into scope of NYC Local Law 144 (Automated Employment Decision Tools — AEDT), whether that law applies, and what it requires, is for the deployer to determine (Companion A lists some citations and is not complete); 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 may involve decisions based solely on automated processing, whether GDPR Article 22 applies, and to whom, is for the deployer to determine (Companion A lists some citations and is not complete). The §5.2 affirmation requirement (an affirmation MUST be an affirmative human act, not a passive signal) records that a named human affirmed the record. Whether that takes a decision outside 'based solely on automated processing' is for the deployer to determine, and the requirement does not by itself discharge any Article 22 obligation." "Where individual-altitude records inform an employment-adverse decision, whether the Fair Credit Reporting Act (15 U.S.C. §1681 et seq.) or its state analogs (including California's ICRAA and CCRAA) apply, including whether the records constitute a 'consumer report' under §1681a, and what they require, is for the deployer to determine (Companion A lists some citations and is not complete). Any such obligations 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 deployer may use the §5.5 redaction-event record to document removal of the named fields; whether that meets the right, and what happens to the archival record, is for the deployer to determine; the erasure determination and the lawful-basis analysis are not made, discharged, or satisfied by the Standard's records (see Companion A §A.2.bis)."*

Named employment counsel of record. Section 3.1 recommends that a Charter whose records describe people at function-leader altitude or below name employment counsel of record for the Charter's scope. This Appendix adds no rule to that recommendation.

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 be sought in litigation, including U.S. employment litigation. Whether they are discoverable, and how, is not something the Standard decides; the Standard makes no claim that its seal (§5.1(3)), which is tamper-evident on a stored, access-controlled record per §2.2.18 and is not cryptographic immutability against all attack surfaces, affects that question. The seal preserves tamper-evidence of what the deployer recorded and when, and a deployer that installs the Standard expecting the seal to function as a litigation shield has misread §5. Where a deployer has a litigation hold or anticipates U.S. employment litigation, how the records are produced, redacted and used in litigation is for the deployer to determine. 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. The Standard makes no claim that following it, or declaring any Level, is a defense to any employment claim. Whether and how a deployer uses the records in litigation is for the deployer to determine.

"Deployers operating in jurisdictions with works-council co-determination or consultation regimes — including German Betriebsverfassungsgesetz §87, French Code du travail Article L. 2312-26, 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 any consultation obligation. Whether one applies, and who carries it, is for the deployer to determine. The Charter records the outcome in its `works_council_consultation_record` (§3.1)."

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 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 whatever law applies to the deployer (Companion A lists some citations and is not complete): 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 is the deployer's, and the two 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 departure from this section, 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`), as the deployer determines. 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, and the name is not used as a certification mark.**

Vendors do NOT certify against the Standard. A vendor MAY build "Standard-aware tooling" or "tooling that implements the Standard" 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 the Standard's requirements 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. 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.

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); traditional internal-control and assurance frameworks; and case law. 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 for that organization to determine; Companion A cites them and does not characterize them.

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 Intentionally removed

This subsection is intentionally removed; its number is kept so that later numbering stays stable.

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.

Souza et al. — "PROV-AGENT: Unified Provenance for Tracking AI Agent Interactions in Agentic Workflows." - Authors: Renan Souza, Amal Gueroudji, Stephen DeWitt, Daniel Rosendo, Tirthankar Ghosal, Robert Ross, Prasanna Balaprakash, Rafael Ferreira da Silva - Publication: 2025 IEEE 21st International Conference on e-Science (accepted paper) - Year: 2025 - Pointer URL: `https://arxiv.org/abs/2508.02866` (arXiv:2508.02866) - Neutral descriptor: Research paper proposing a provenance model, built on W3C PROV, for AI-agent interactions in workflows.

AGENTSAFE — framework for agentic AI safety. - Publication: arxiv preprint - Date: December 2025 - Pointer URL: `https://arxiv.org/abs/2512.03180` (arXiv:2512.03180, "AGENTSAFE: A Unified Framework for Ethical Assurance and Governance in Agentic AI"). - 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: none cited; the entry is identified by author, title, and year. - 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 Reference files

§12.4 records the machine-readable reference files published with the Standard. Reference files are not a conformance-certifying body, and using them does not by itself produce a Standard-conformant Charter. Reference files structure the inputs; whether anything conforms to any framework is for the deployer to determine. No public implementation of the Standard is cited here.

Reference files. - License: MIT License (the Standard's own text is separately licensed under Creative Commons Attribution 4.0 International (CC-BY 4.0)). - Neutral descriptor: Machine-readable reference files (schemas, state machines, the signal list, the reporter contract and a test

plan), published in `standard/v5.0/`. The reference files are not a substitute for the Standard's normative text and do not declare conformance with any external regulatory framework.

12.5 Versioning

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 files are updated. 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: 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.
