

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

The worked Charters below are self-contained illustrations of the Standard's Charter mechanism (Section 3) across common functional surfaces. They are illustrative, not normative, and not certifications. This Companion's cross-references to the Standard's core sections (§1–§7) resolve against the core Reading Edition (rev. 10).

Companion B – Worked Charter Library

Disclaimer pointer. See top of the core Standard.

B.1 Purpose of Worked Examples

This Companion makes the Standard's primitives concrete. Sections 3, 4, 5, 6, and 7 describe the Charter mechanism, the Mode taxonomy, the record lifecycle, the required artifact set, and the conformance levels at the altitude of normative text. Companion B takes those primitives and shows them operating against nine specific cases: seven function-specific Decision Interface Charters and two executive-altitude interfaces. Each worked example does six things. It states the recurring decision class the Charter binds. It declares the Mode the Charter authorizes per Section 4 §4.4. It names the schedule of records the Charter commits to per Section 6. It surfaces the conformance signals the Charter would emit at Levels 1, 2, and 3 per Section 7. It states whether the Section 4 §4.6 disclosure block is required, leaving whether Article 50 applies for the deployer to determine. And it notes how records under the Charter progress through the Section 5 lifecycle to **affirmed** -and-sealed status.

The worked examples are not certifications. They do not certify that the underlying Charter satisfies any regulatory requirement, that the Charter's decisions discharge any legal obligation, or that the schedule of records produced under the Charter constitutes evidence in any forum. They describe how the Standard's primitives operate against the recurring decision class each Charter binds — the audit-ready decision provenance the Charter would produce if operated conformantly. Counsel and auditors convert that provenance into evidence; the worked examples themselves do not.

Each worked example closes with an implementation-context note in the Standard's own voice, describing the organizational surface the Charter is filed against so the reader needs no outside reference to follow it.

A reader looking for the per-Section primitives reads Sections 3 through 7 of the core Standard and Companion A (Regulatory Cross-References). A reader looking for how those primitives operate against named decision classes at director, executive, and board altitude reads Companion B.

B.2 PMM Decision Interface Charter – Pre-Commitment Positioning Lock Against Build-Lock

Worked Example 9.1 — Product Marketing Decision Interface Charter.

Context

The Product Marketing (PMM) Decision Interface Charter binds the recurring decision class of **positioning lock before build-lock**. This is the surface at which the Director of Product Marketing commits the category frame, the three-tier message hierarchy, the named-alternative competitive frame, and the launch-tier classification before product scope is committed for the bet. The Charter's load-bearing wall protects against positioning elaboration after build-lock. That is the failure mode where the category story is retrofitted to scope rather than authored against the market the product will enter. Three archetypes shift Charter weight: category-creator, demand-gen-led, and brand-led. Inputs, cadence, and record hold across all three. The Charter operates at director altitude, paired against the Chief Product Officer (CPO). The Emotional Cost splits between the Director's "no" on a late scope change and the CPO's acceptance of positioning as a pre-commitment decision rather than post-scope copy.

Mode declaration

mode-1 — Mode 1, Human-Led, AI-Enforced. The Charter authorizes Mode 1 dispatch for the positioning-lock decision class. The Director of Product Marketing is the author of record for the positioning artifact, the message hierarchy, the launch-narrative brief, and the launch-tier classification. The AI worker participating in the decision flow operates as the enforcement mechanism per Section 4 §4.2. It reads Charter state. It validates that the Director's draft satisfies the Charter's structural requirements: five named input classes populated, T1/T2/T3 launch-tier

declared, named-alternative competitive frame present. And it flags conformance-signal gaps before the dispatch state machine permits the record to close.

Schedule of records

The Charter commits to maintaining the following record-types per Section 6 §6.3.1:

- **Decision records** — one per positioning-lock decision under the Charter, carrying decision statement, options considered, required inputs used, success criteria threaded to the T+2 evidence window, and accountable-owner sign-off at close.
- **Re-decision records** — fired on three evidence types. Outcome-evidence is a pipeline or activation miss against the T+2 window. Market-evidence is a category motion fired by Competitive Intelligence: a new entrant redefining the category, a competitor shipping a substitute, or a pricing move compressing the premium. Counterparty-specific evidence is scope drift against the positioning commit across two consecutive bet reviews.
- **Escalation records** — fired when the Charter's bounded-exit mechanism activates (the positioning lock is reopened after build-lock without a re-decision trigger firing); the withdrawal returns the decision to formulation.
- **Charter-amendment records** — fired when the Charter itself is amended (mode change, decision-class boundary change, accountable-owner change, schedule-of-records change). Per Section 3 §3.3 and Section 6 §6.3.1, amendments dispatch under the Charter's declared `mode_declaration` and produce records of their own.

Conformance signals

At Section 7 grading, a conformant PMM Charter would emit:

- **Level 1** — Charter-conformant signals: `mode_declaration` populated; `charter_id`, `accountable_owner`, `decision_class`, `inside_decisions`, `outside_decisions` populated at `mode-declared`; `cadence`, `record_location`, `re_decision_triggers`, `escalation_rule` populated at `fields-required`; `schedule_of_records` and `conformance_level_declared` populated at `fields-completed`.
- **Level 2** — Mode-disambiguated signals: every record under the Charter carries `dispatch_mode`; the Mode-Drift Composed Mitigation's Layer 2 Substantive-Authorship Challenge fires at record-close on every Mode-1-declared record; sampled records pass Layer 1 statistical detection at the 15% baseline rate.
- **Level 3** — Continuously auditable signals: the schedule of records resolves to a durable, searchable, linkable surface satisfying the Standard's findability rule (records reach a reader not in the room within 30 seconds); Product Operations quarterly findability audit confirms;

charter-amendment records preserve the Charter's pre-amendment Mode for the records dispatched under it.

Disclosure block (Section 4 §4.6)

Disclosure block required under §4.6: no for the canonical Mode 1 positioning-lock decision class. Whether Article 50 applies: for the deployer to determine. The artifact reaching the reader is human-authored (the positioning artifact, the launch-narrative brief, the message hierarchy), and the AI worker's contribution is internal to the production process rather than externalized in the output. The Mode 1 edge case at Section 4 §4.7 governs the narrow circumstance in which AI-generated draft messaging is embedded inside an otherwise human-authored positioning artifact; that case carries the disclosure block at the embed point per Section 4 §4.6.2 even though the surrounding container is Mode 1. A Charter whose accountable owner anticipates AI-drafted message variants embedded into positioning records declares `mode-1-with-embedded-mode-2-summary` instead.

Footer

Worked Example 9.1 — implementation context: This Charter governs the positioning-lock surface, where the category frame, message hierarchy, and launch-tier classification are committed before product scope is locked for the bet. Three archetypes shift its weight: category-creator (the category frame is itself the bet), demand-gen-led (the message hierarchy serves a pipeline machine), and brand-led (positioning is a brand-coherence commitment). Inputs, cadence, and record hold across all three. The Charter operates at director altitude, paired against the CPO. The cost of the Director's "no" on a late scope change is set against the CPO's acceptance of positioning as a pre-commitment decision rather than post-scope copy. The underlying discipline is that go-to-market is a strategic choice made at formulation, not a handoff after build.

B.3 Design Decision Interface Charter — Interaction-Model Lock and Design-System Stewardship

Worked Example 9.2 — Design Decision Interface Charter.

Context

The Design Decision Interface Charter binds two recurring decision classes from the Director of Design's seat as a peer inside the organizational leadership triad: **interaction-model lock before build-lock**, and **design-system stewardship as a portfolio-altitude asset with a named envelope**. The Charter's load-bearing wall protects against two paired failures. The first is the design system degrading faster than the product surface grows, which is design-system coherence debt. The second is UX research collapsing into post-ship usability testing rather than pre-commitment evidence. Three archetypes shift which failure dominates: embedded product designer, in-house design team, and enterprise design function. Inputs, cadence, and record hold across all three.

Mode declaration

mode-1 — Mode 1, Human-Led, AI-Enforced. The Charter authorizes Mode 1 dispatch for both the interaction-model-lock decision class and the design-system stewardship decision class. The Director of Design is the author of record for the interaction-model artifact, the design-system release and adoption record, the accessibility-posture call, and the design-debt retirement sequence against the CPO-CTO envelope. The AI worker enforces Charter discipline per Section 4 §4.2. It checks that interaction-model lock fires before build-lock, that design-system release notes attach to the record, and that envelope reconciliation is named. It does not author the substantive design decision content.

Schedule of records

Per Section 6 §6.3.1:

- **Decision records** — one per interaction-model lock, one per design-system evolution call, one per accessibility-posture commitment, one per design-debt retirement decision.
- **Re-decision records** — fired on three evidence types. Outcome-evidence is a usability signal or activation miss tied to an interaction-model breach. Market-evidence is a Competitive-Intelligence-fired category motion that invalidates a live interaction assumption, or a competitor that redefines the evaluator craft bar. Counterparty-specific evidence is a design-system adoption drop past a named floor, or an envelope cut without joint CPO-CTO review.
- **Escalation records** — fired on the bounded-exit conditions: interaction-model lock reopened after build-lock without a trigger firing, or envelope cut as an operations line rather than decided as a portfolio call. The withdrawal routes the decision back to the CPO-CTO-Dir Design forum.
- **Charter-amendment records** — fired on Charter-altitude amendments (cadence change, archetype reclassification on scale transition, etc.).

Conformance signals

- **Level 1** — Charter-conformant: all 16 Charter fields populated at the required-at-state per Section 3 §3.2; the Charter has reached `fields-completed` with mode declared, schedule enumerated, and conformance level declared.
- **Level 2** — Mode-disambiguated: every record carries `dispatch_mode` ; Layer 2 Substantive-Authorship Challenge fires at close on every Mode-1-declared record. The structural enforcement in Section 4 §4.3 (Mode 2 records require `disclosure_metadata_pointer` at `drafted`) is moot under Mode 1 but recorded as inapplicable.
- **Level 3** — Continuously auditable: the design-system evolution record resolves to a versioned, linkable surface; the Product Operations quarterly findability audit passes; design-debt retirement decisions thread to the bet-level Decision Records that consume them.

Disclosure block (Section 4 §4.6)

Disclosure block required under §4.6: no for the canonical Mode 1 dispatch. Whether Article 50 applies: for the deployer to determine. The interaction-model artifact, design-system release notes, accessibility-posture commitment, and design-debt retirement sequence are human-authored at director altitude. The Mode 1 edge case applies if a deployer at the enterprise design function archetype dispatches AI-generated interaction prototypes into the design-review forum as embedded sub-outputs; in that case, the Charter declares `mode-1-with-embedded-mode-2-summary` at the Charter altitude and per-record disclosure metadata blocks attach at the embed point per Section 4 §4.7.

Footer

Worked Example 9.2 — implementation context: This Charter governs two surfaces: interaction-model lock before build-lock, and design-system stewardship as a portfolio-altitude asset with a named envelope. Three design-team archetypes shift which failure dominates: the embedded product designer (interaction-model lock is the live risk), the in-house design team (design-system coherence debt accumulates fastest), and the enterprise design function (accessibility posture and envelope discipline dominate). The Charter protects against a dual failure mode: design-system coherence debt outpacing product growth, and UX research collapsing into post-ship usability testing rather than pre-commitment evidence. Design operates as a peer inside the organizational leadership triad, not as a downstream production function.

B.4 Engineering Decision Interface Charter – Three Engineering-Owned Decision Classes at the CPO-CTO Interface

Worked Example 9.3 — Engineering Decision Interface Charter.

Context

The Engineering Decision Interface Charter binds three engineering-owned decision classes dispatched against the CPO-CTO interface: **reliability posture as a named output** (SLOs, error budgets, incident-rate floor), **architectural-debt retirement as a decided sequence** with envelope discipline, and **build-buy-OSS at the platform layer** as a joint CPO-CTO call. The Charter's load-bearing wall protects against two failures. The first is tech debt being absorbed as operations rather than decided as a strategic bet. The second is architectural-debt retirement being silently deferred until a reliability breach forces it at cost. Three archetypes shift which decision drifts first: platform-led, product-led, and embedded-team-led.

Mode declaration

`mode-1-with-embedded-mode-2-summary` — Mode 1 with the embedded-Mode-2-summary edge case. The Charter authorizes Mode 1 dispatch for the canonical decision content: the Chief Architect is the author of record for reliability posture, architectural-debt retirement sequencing, and build-buy-OSS calls. The embedded-Mode-2-summary edge case is engaged for AI-assisted architecture-review outputs and AI-generated capacity-model reads dispatched into the architecture-review forum. The structural handling per Section 3 §3.4 and Section 4 §4.7 attaches in two places. First, the Charter's `disclosure_metadata_pointer` is required at `fields-completed`. Second, per-decision records carrying AI-generated sub-outputs (capacity reads, debt-classification synthesis, vendor-license summaries) attach disclosure metadata blocks at the embed point with `mode_1_edge_case_flag = true` per Section 4 §4.4.

Schedule of records

- **Architecture decision records** — one per architecture-review gate outcome on a committed bet.
- **Reliability-review records** — monthly platform review records and quarterly architecture review records, synchronized to the enterprise reliability calendar.
- **Build-buy-OSS evaluation records** — one per platform-layer call.

- **Re-decision records** — fired on three evidence types. Outcome-evidence is a reliability breach against the SLO floor, or a velocity collapse against the capacity model. Market-evidence is a build-buy-OSS inflection: a platform category commoditizes, a vendor changes license terms, or an OSS maintainer stance shifts. Counterparty-specific evidence is architectural debt compounding past a named floor, or an envelope cut that re-scopes retirement outside joint review.
- **Escalation records** — fired when reliability posture is reopened as a scope conversation without a trigger firing, or when architectural-debt retirement is re-scoped as operations.
- **Disclosure-review records** — required per Section 6 §6.3.1 because Mode 2 sub-outputs are in scope under the embedded-summary classification. The disclosure metadata block attached to each AI-generated capacity-model read or debt-classification synthesis is re-reviewed on the cadence the Charter declares (typically monthly synchronous with the platform review).
- **Charter-amendment records** — fired on Charter-altitude amendments.

Conformance signals

- **Level 1** — Charter-conformant: all 16 Charter fields populated; the Charter has reached `fields-completed` with `mode_declaration = mode-1-with-embedded-mode-2-summary` and `disclosure_metadata_pointer` populated.
- **Level 2** — Mode-disambiguated: every record carries `dispatch_mode` ; embedded AI-generated sub-outputs each carry their own disclosure metadata block satisfying Section 4 §4.6.2's five required fields (declaring-authority, ai-system-identity, jurisdictional-applicability-tag, content-type-tag, generation-timestamp); the Layer 1 statistical detection rate is elevated to 30% on this Charter because `mode-1-with-embedded-mode-2-summary` is a structural surface for drift.
- **Level 3** — Continuously auditable: architecture decision records, reliability-review records, and build-buy-OSS evaluation records all findable in 30 seconds; envelope-reconciliation telemetry threads from this Charter into the CPO-CTO interface; disclosure metadata blocks survive anonymization where the records circulate beyond the Charter (preserving the four-of-five fields per Section 4 §4.6.3's anonymization-protocol cross-reference).

Disclosure block (Section 4 §4.6)

Disclosure block required under §4.6: yes, conditionally, at the embed point. Whether Article 50 applies: for the deployer to determine. The canonical Mode 1 architecture decision carries no disclosure block in itself — the Chief Architect's substantive call is human-authored and the AI worker functions as Charter-conformance enforcement. But the AI-generated capacity-model reads, AI-assisted debt-classification syntheses, and AI-drafted vendor-license summaries

embedded into architecture decision records are Mode 2 content at the embed point under Section 4 §4.7's content-level analysis, and the disclosure block attaches to those embedded sub-outputs. The five required metadata fields per Section 4 §4.6.2 attach at the embed point. Unless the Charter declares the Section 4 §4.6.1 exclusion, the block attaches wherever the record is produced.

Footer

Worked Example 9.3 — implementation context: This Charter governs three engineering-owned decision classes operated against the CPO-CTO interface: architecture-review gates that set reliability posture, technical-debt retirement sequenced as a decided plan rather than absorbed as operations, and build-buy-OSS calls at the platform layer. Three archetypes shift which decision drifts first: platform-led (build-buy-OSS dominates), product-led (reliability posture against feature velocity), and embedded-team-led (architectural debt compounds silently). The load-bearing discipline is that the CPO and CTO decide these as peers at formulation. Tech debt is then decided as a strategic bet rather than deferred until a reliability breach forces it at cost.

B.5 Data and Analytics Decision Interface Charter — Metric-Integrity Contract Under the Five-Condition Discipline

Worked Example 9.4 — Data and Analytics Decision Interface Charter.

Context

The Data and Analytics Decision Interface Charter binds the recurring decision class of **metric-integrity contract operation** under a five-condition discipline: a single definition, a stable measurement window, a single owner, a single source of truth, and a documented exclusion list. The Charter's load-bearing wall protects against silent metric-integrity erosion. That is the failure mode where any of the five conditions drifts (a redefined activation event, a shifted window, a new owner who did not inherit the exclusion list) and the learning signal degrades without a failure message. Three archetypes shape contract ownership across the organization's scale forms: embedded product analyst, embedded data team, and enterprise data function.

Mode declaration

`mode-1-with-embedded-mode-2-summary` — Mode 1 with the embedded-Mode-2-summary edge case. The Data Science Lead is the author of record for the canonical metric definitions, the cohort floors and boundaries, the metric-integrity audit log against the five conditions, and the gating call on whether a measurement is powered for the decision it informs. The embedded-Mode-2-summary edge case is engaged for a reason. Most modern analytics environments produce AI-generated experiment summaries, AI-generated cohort-trajectory reads, and AI-generated metric-drift detections that dispatch into a consuming Charter as decision-aid content. Each embedded AI-generated artifact attaches a disclosure block at the embed point per Section 4 §4.7.

Schedule of records

- **Decision records** — one per metric-integrity contract instance (definition, owner, source of truth, exclusion list); one per cohort-floor commitment; one per experiment call (powered or not powered).
- **Canonical metric catalog** — versioned, linkable from every consuming Charter. Decision records reference metrics by version, not name.
- **Metric-integrity audit log** — quarterly drift-flag reads against the five conditions.
- **Re-decision records** — fired on three evidence types. Outcome-evidence is metric drift detected in the monthly audit, or a cohort floor breached below the statistical-power threshold. Market-evidence is an experiment assumption disproven by new data, or an instrumentation-coverage drop past a named floor. Counterparty-specific evidence is a metric-definition contract with Product Operations or a consuming Charter that drifts out of version sync.
- **Escalation records** — fired when the Data Lead withdraws the gating input because a decision depends on instrumentation that does not yet exist; the decision either waits or proceeds explicitly without the signal.
- **Disclosure-review records** — re-review of disclosure metadata blocks attached to AI-generated experiment summaries, cohort-trajectory reads, and drift-detection outputs.
- **Charter-amendment records** — fired on archetype reclassification (e.g., the embedded product analyst archetype scaling to embedded data team) or on cadence amendment.

Conformance signals

- **Level 1** — Charter-conformant: 16 fields populated; `mode_declaration = mode-1-with-embedded-mode-2-summary`; `disclosure_metadata_pointer` populated; six input classes

enumerated (canonical metric definitions, event-stream health, cohort floors, experiment results, data-product release notes, metric-integrity audit log).

- **Level 2** — Mode-disambiguated: every embedded AI-generated artifact carries the five required disclosure metadata fields per Section 4 §4.6.2; Layer 1 statistical detection runs at the 30% rate (the elevated rate for `mode-1-with-embedded-mode-2-summary` Charters); the metric-version-linkage check confirms decision records reference metrics by version not name.
- **Level 3** — Continuously auditable: the canonical metric catalog satisfies the Standard's findability rule with version pointers; the metric-integrity audit log is preserved across cadences; the disclosure metadata blocks survive anonymization per Section 4 §4.6.3 with the four-of-five rule (the declaring-authority field may transform to a controlled-vocabulary placeholder; the AI-system identity, the Mode declaration, the jurisdictional-applicability tag, the content-type tag, and the generation-timestamp pass through unchanged).

Disclosure block (Section 4 §4.6)

Disclosure block required under §4.6: yes for embedded AI-generated artifacts. Whether Article 50 applies: for the deployer to determine. This is the Charter with the most embedded AI-generated content among the seven Appendix C instances — most modern analytics environments rely on AI-generated content in the dispatch path. The five required metadata fields per Section 4 §4.6.2 attach at every embed point: `declaring-authority` (the deploying organization, not the AI vendor), `ai-system-identity` (vendor + model + version), `jurisdictional-applicability-tag` (controlled vocabulary including `eu`), `content-type-tag` (the pre-set vocabulary `{decision-summary, recommendation, decision-aid, draft, classification, synthetic-media, other:<type>}`), and `generation-timestamp` (ISO 8601). Unless the Charter declares the Section 4 §4.6.1 exclusion, the block attaches wherever the read is produced.

Footer

Worked Example 9.4 — implementation context: This Charter governs the metric-integrity contract operated against a five-condition discipline: a single definition, a stable measurement window, a single owner, a single source of truth, and a documented exclusion list. Three scale-form archetypes shape contract ownership: the embedded product analyst (the contract lives with one person), the embedded data team (the contract is a team artifact), and the enterprise data function (the contract is a governed standard). The underlying discipline is that an organization learns through outcomes only when the

measurement is integrity-protected. Silent drift in any of the five conditions degrades the learning signal without a visible failure.

B.6 Customer-Outcome Charter (CS / Value Realization Instance) – Promise-to-Outcome Gap as Gating Input to Portfolio Review

Worked Example 9.5 — Customer-Outcome Decision Interface Charter (Customer Success / Value Realization Instance).

Context

The Customer-Outcome Charter binds the recurring decision class of **promise-to-outcome gap surfacing by cohort as a gating input to the portfolio review**. The Charter is co-owned by the Director of Customer Success and the Value Realization lead. It is the only worked example in this library with dual director-altitude ownership, reflecting Customer Success and Value Realization as the function that makes the customer's verdict visible. The Charter's load-bearing wall protects against outcome absorption. That is the failure mode where the realized-value signal is absorbed into the adoption dashboard, the QBR narrative, or the expansion forecast, and the customer-outcome loop silently closes against self-report. Three archetypes shift the read mechanics: high-touch enterprise, tech-touch scale, and hybrid commercial. Naming the archetype is itself load-bearing, because an unnamed Charter defaults to high-touch and misreads the scale surface.

Mode declaration

mode-1-with-embedded-mode-2-summary — Mode 1 with the embedded-Mode-2-summary edge case. The Director of Customer Success and the Value Realization lead are co-authors of record for the substantive decision content (renew, expand, reference, escalate calls at account and cohort altitude; realized-value gap surfacing; cohort-outcome cadence definition). The embedded-summary edge case is engaged because the tech-touch scale archetype, in particular, relies on AI-generated content in the dispatch path. That includes AI-generated cohort-trajectory reads, AI-generated churn-risk signals against named thresholds, and AI-generated onboarding time-to-first-value classifications dispatched into the cohort-outcome record as decision-aid sub-outputs.

Schedule of records

- **Customer-outcome records** — one per cohort per cadence cycle, threaded to each bet's Decision Record by bet-ID. Each record carries seven input classes: account-health score with segment cut, cohort adoption-depth trajectory, realized-value signal as promise-to-outcome gap by cohort, expansion and contraction leading indicators, churn-risk signals, customer-promise register from sales handoff at the CPO-CRO interface, and onboarding time-to-first-value.
- **Re-decision records** — fired on three evidence types. Outcome-evidence is a realized-value gap that exceeds the Charter's threshold across two consecutive cohorts. Market-evidence is a competitor introducing a substitute outcome the promise cannot match. Counterparty-specific evidence is a sales-handoff promise register that drifts against the customer-outcome record across two cycles.
- **Escalation records** — fired in two cases. The first is when the realized-value signal confirms the continuation threshold across three cohort cycles (clean-completion bounded exit). The second is when a re-decision trigger fires and the bet returns to formulation. Horizon mis-specification escalates to the CPO-CFO interface for threshold recalibration.
- **Disclosure-review records** — required for the AI-generated cohort-trajectory and churn-risk artifacts dispatched into the customer-outcome record.
- **Charter-amendment records** — fired on archetype reclassification (high-touch → hybrid commercial as scale shifts) or on co-ownership protocol amendment between the Director of Customer Success and Value Realization.

Conformance signals

- **Level 1** — Charter-conformant: 16 fields populated; `accountable_owner` carries the Charter's dual-co-owner protocol named on a single record (one of the two co-owners is named per record per Section 3 §3.2's "one and only one" rule, and the co-ownership protocol is recorded in `inside_decisions`); `mode_declaration = mode-1-with-embedded-mode-2-summary`; `disclosure_metadata_pointer` populated at `fields-completed` per Section 3 §3.4 and Section 6 §6.2.
- **Level 2** — Mode-disambiguated: every record carries `dispatch_mode`; embedded AI-generated cohort-trajectory reads carry the five required disclosure-block fields; the customer-promise register threads from the CPO-CRO interface into this Charter without version drift.
- **Level 3** — Continuously auditable: customer-outcome records findable in 30 seconds; promise-to-outcome gap by cohort published into the portfolio review on cadence; T+6 and T+12 checkpoints threaded to each bet's Business-Case artifact per Section 6's Required Artifact Set treatment.

Disclosure block (Section 4 §4.6)

Disclosure block required under §4.6: yes, conditionally, for embedded AI-generated cohort-trajectory reads and churn-risk signals. Whether Article 50 applies: for the deployer to determine. The block attaches more often in the tech-touch scale archetype, where AI-generated reads dominate the dispatch path, and less often in the high-touch enterprise archetype, where named-account cadence relies more on human-authored quarterly reads. The five required metadata fields per Section 4 §4.6.2 attach at every embed point. The declaring authority names the person who prepares the disclosure and the organization they act for, not the AI-system vendor, per Section 4 §4.6.2.

Footer

Worked Example 9.5 — implementation context: This Charter governs the surfacing of the promise-to-outcome gap by cohort as a gating input to the portfolio review, co-owned by the Director of Customer Success and the Value Realization lead. Three archetypes shift the read mechanics: high-touch enterprise (named-account human reads dominate), tech-touch scale (AI-generated cohort reads dominate), and hybrid commercial (a blend). Naming the archetype is itself load-bearing, because an unnamed Charter defaults to high-touch and misreads the scale surface. The load-bearing wall protects against outcome absorption, where the realized-value signal is absorbed into adoption dashboards, QBR narratives, or expansion forecasts and the customer-outcome loop silently closes against self-report. Value Realization is the function that makes the customer's verdict visible, and the discipline is learning through outcomes.

B.7 Business Development Decision Interface Charter — Partnership Architecture as Pre-Commitment Artifact

Worked Example 9.6 — Business Development Decision Interface Charter.

Context

The Business Development Decision Interface Charter binds the recurring decision class of **partnership-architecture commitment before bet commits**. Business Development is the third strategic formulation surface alongside Product Management and Product Marketing. This Charter is its altitude-specific instance of the Charter template. Business Development and the Chief Financial Officer (CFO) co-decide partnership economics. Business Development is

accountable for the partnership-architecture spine, the CFO for the margin floor, with the Chief Revenue Officer (CRO) interface receiving the output as input. The Charter's load-bearing wall protects against the partnership-architecture decision being made inside a deal room rather than at formulation. Three archetypes shift input weighting: platform-dependencies-first, revenue-share-first, and co-sell-credit-first.

Mode declaration

mode-1 — Mode 1, Human-Led, AI-Enforced. The Business Development Leader is the author of record for the partnership-architecture map, the make-buy-partner TCO-over-three-years read, the negotiated economics (co-decided with the CFO), the platform-dependency risk register, and the partner-mediated GTM motion call. The AI worker enforces Charter discipline. It checks that the architecture artifact precedes the bet commit, that the six input classes are populated, and that the outside-decisions boundary is honored: deal-level pursuit stays with the CRO, product scope on committed bets stays with the Director of Product Management, and the margin floor stays with the CFO as a gating input. The AI worker does not author substantive partnership content.

Schedule of records

- **Partnership Architecture artifact** — the Charter's signature record, one per bet at formulation.
- **Per-bet Partner-Dependency Reads** — one per bet, gating Launch Readiness.
- **Re-decision records** — fired on three evidence types. Outcome-evidence is partner-sourced pipeline below hypothesis across the T+2 window. Market-evidence is a platform owner's pricing or API move that compresses category position, or a competitor that secures an exclusive with a partner. Counterparty-specific evidence is revenue-share compression, MDF withdrawal, or a channel-margin squeeze across two consecutive cycles.
- **Escalation records** — fired when negotiated economics are renegotiated post-commit without a re-decision trigger firing, or when partnership architecture is overwritten inside a deal room; the withdrawal returns the decision to formulation.
- **Charter-amendment records** — fired on archetype reclassification (platform-dependencies-first → revenue-share-first as the partnership thesis shifts) or on cadence amendment.

Conformance signals

- **Level 1** — Charter-conformant: 16 fields populated; the joint-decision pairing with the CFO recorded in `inside_decisions`; the Chief Revenue Officer and Director of Product

Management interfaces recorded in `outside_decisions` .

- **Level 2** — Mode-disambiguated: every record carries `dispatch_mode = mode-1` ; Layer 2 Substantive-Authorship Challenge fires at close; the negotiated-economics record threads from CFO sign-off without version drift.
- **Level 3** — Continuously auditable: the Partnership Architecture artifact resolves to a durable, linkable surface; per-bet Partner-Dependency Reads thread to Launch Readiness records; charter-amendment records preserve the Charter's pre-amendment archetype classification.

Disclosure block (Section 4 §4.6)

Disclosure block required under §4.6: no for the canonical Mode 1 dispatch. Whether Article 50 applies: for the deployer to determine. The partnership-architecture map, make-buy-partner TCO read, negotiated-economics record, and platform-dependency risk register are human-authored at director altitude. The Mode 1 edge case applies if a deployer dispatches AI-generated TCO models or AI-generated partner-tier economic projections into the negotiated-economics input; in that case, the Charter declares `mode-1-with-embedded-mode-2-summary` at the Charter altitude and per-record disclosure metadata attaches at the embed point per Section 4 §4.7.

Footer

Worked Example 9.6 — implementation context: This Charter governs partnership-architecture commitment at formulation, protecting it from being decided inside a deal room after the fact. Three archetypes shift input weighting: platform-dependencies-first (the architecture is a risk-management exercise), revenue-share-first (the economics dominate), and co-sell-credit-first (the GTM motion dominates). Partnership architecture is one of the pre-commitment go-to-market disciplines. Make-buy-partner is a single sourcing decision read across three executive lenses: Business Development for the architecture, the CFO for the margin floor, and the CRO for the channel motion.

B.8 Product Operations Decision Interface Charter — Cadence Integrity and the Live-or-Die Forum Call

Worked Example 9.7 — Product Operations Decision Interface Charter.

Context

The Product Operations Decision Interface Charter binds the recurring decision class of **operating-layer stewardship**: cadence integrity, ritual spec, decision-record findability, the metric-definition integrity contract with Data and Analytics, and the live-or-die call on which forums earn their place and which are retired. The Charter's load-bearing wall protects against governance theater. Rituals proliferate without gating power, records drift out of findability, and the cadence calendar becomes a meeting inventory rather than a decision architecture. Product Operations brings cadence integrity, ritual spec, record hygiene, and the telemetry that predicts outcome failure. It does not bring decision content. Three configuration archetypes shape the artifact's central versus distributed hold: embedded-per-product, centralized-ops-function, and federated hybrid.

Mode declaration

mode-1-with-embedded-mode-2-summary — Mode 1 with the embedded-Mode-2-summary edge case. The Director of Product Operations is the author of record for the canonical decisions: cadence and ritual spec, record hygiene standard, operating-calendar versioning, the metric-definition integrity contract, and the live-or-die forum call. The embedded-summary edge case is engaged for a reason. Telemetry pipelines increasingly include AI-generated cadence-adherence classifications, AI-generated charter-decay signal detections, and AI-generated re-decision-integrity reads dispatched into the quarterly charter-decay review at portfolio review.

Schedule of records

- **Operating calendar** — versioned, linkable from every Charter that schedules into it.
- **Canonical Rituals reference** — the named-spec record for each ritual, with the criteria under which the ritual is retired.
- **Charter archive** — the index of all Charters running across the organization.
- **Decision-record archive** — the index of all decision records produced under those Charters.
- **Machinery-signal telemetry dashboard** — the five-signal observability layer (decision latency, commitment drift, charter decay, re-decision integrity, cadence adherence).
- **Re-decision records** — fired on three evidence types. Outcome-evidence is a cadence-adherence or charter-decay signal that breaches its floor across two consecutive cycles. Market-evidence is a scale-transition that invalidates an existing configuration, or AI-acceleration that shortens the cycle inside which machinery failures compound. Counterparty-specific evidence is a Business Operations handoff that degrades, or the metric-definition contract drifting.

- **Escalation records** — fired when cadence and ritual spec are overridden by executive calendar accommodation without a re-decision trigger firing; the withdrawal returns the decision to the interface between the Chief Product Officer and the Director of Product Operations.
- **Disclosure-review records** — for embedded AI-generated telemetry classifications and charter-decay signal detections.
- **Charter-amendment records** — fired on configuration archetype reclassification (embedded-per-product → federated hybrid as scale shifts).

Conformance signals

This Charter is the meta-Charter: it is the Charter that audits the schedule of records for every other Charter. Its own conformance signals are:

- **Level 1** — Charter-conformant: 16 fields populated; `mode_declaration = mode-1-with-embedded-mode-2-summary`; `disclosure_metadata_pointer` populated at `fields-completed` per Section 3 §3.4 and Section 6 §6.2; the six input classes enumerated.
- **Level 2** — Mode-disambiguated: every record carries `dispatch_mode`; embedded AI-generated telemetry artifacts each carry the five required disclosure-block fields; the Layer 1 statistical detection rate is at the 30% rate for `mode-1-with-embedded-mode-2-summary`.
- **Level 3** — Continuously auditable: this is the Charter's home altitude. Product Operations is the function that operates the Standard's findability rule across every Charter. The 30-second reconstruction test is Product-Operations-audited quarterly per Section 6 §6.3.2. The Product Operations Charter's own records satisfy the rule by construction: the Charter cannot pass Level 3 without findability, and the function whose Charter this is is the function that audits findability.

Disclosure block (Section 4 §4.6)

Disclosure block required under §4.6: yes, conditionally, for embedded AI-generated cadence-adherence telemetry, AI-generated charter-decay signal detections, and AI-generated re-decision-integrity classifications. Whether Article 50 applies: for the deployer to determine. The block attaches at a moderate rate because telemetry pipelines increasingly include AI-generated signal classification, but the substantive decision content (which forums live, which die, what cadence the operating calendar carries) is human-authored. The five required metadata fields attach at the embed point.

Footer

Worked Example 9.7 — implementation context: This Charter governs operating-layer stewardship: cadence integrity, ritual specification, decision-record findability, the metric-definition integrity contract with Data and Analytics, and the live-or-die call on which forums earn their place and which are retired. It consumes a five-signal machinery telemetry layer: decision latency, commitment drift, charter decay, re-decision integrity, and cadence adherence. Three configuration archetypes shape its central-versus-distributed hold: embedded-per-product (stewardship is local), centralized-ops-function (stewardship is a shared service), and federated hybrid (a blend). The load-bearing wall protects against governance theater, where rituals proliferate without gating power and the cadence calendar becomes a meeting inventory rather than a decision architecture.

B.9 CEO-CPO Charter — The Charter Mechanism Applied at the One Interface the CPO Cannot Govern Unilaterally

Worked Example 9.8 — CEO-CPO Decision Interface Charter.

Context

The CEO-CPO Charter binds the recurring decision class of **strategic direction of the portfolio against the company's plan, plus tenure-defining organizational bets that require Chief Executive Officer (CEO) air cover**. The load-bearing structural fact is that the CEO is the one node the Chief Product Officer (CPO) cannot charter into existence or re-architect mid-tenure. What the CPO can design is the interface, and the interface is only as good as the honesty about which CEO is on the other side. Three CEO archetypes break CPO tenure if left un-managed: product-native, function-native, and outcome-native. The Charter's archetype-named field surfaces which CEO is on the other side and which risk pattern dominates. The Charter's bring/do-not-bring discipline is sharp. Brought: portfolio state, bet state, decision-system state, and organizational state. Not brought: elaboration on scope, pricing, or ownership re-assignments. The "out loud" pushback rule is the structural primitive that prevents Outside decisions from collapsing into the forum.

Mode declaration

mode-1-with-embedded-mode-2-summary — Mode 1 with the embedded-Mode-2-summary edge case. The CPO is the author of record for the canonical decisions reaching the CEO interface: portfolio state against the capital envelope, bet state against continuation thresholds, decision-

system state against health signals, and organizational state against commitments made. The embedded-summary edge case is engaged because the monthly CEO-CPO sync increasingly carries AI-generated portfolio-state summaries. These include signal classification across the bet portfolio, AI-generated continuation-threshold trajectory reads, and AI-drafted re-decision-trigger fire-event summaries dispatched into the sync as decision-aid sub-outputs.

This is one of the cases where the `mode-1-with-embedded-mode-2-summary` declaration is most load-bearing. A Mode 1 declaration alone would miss the AI-generated portfolio-state summaries embedded into the sync's outputs. It would then silently skip the disclosure block on those embedded sub-outputs, including where they reach downstream board materials, investor communications, or regulatory disclosures. A Mode 2 declaration would mis-classify the substantive decision content, which is human-authored. The third value is the structural handling Section 4 §4.4 names for precisely this case.

Schedule of records

- **CEO-CPO Charter** — the artifact itself, maintained by Product Operations and auditable by the General Counsel.
- **Decision records** — one per CEO-CPO sync, carrying the four bring-classes (portfolio, bet, decision-system, organizational state), the explicit Outside-decisions push-back log when applicable, and the accountable-owner sign-off.
- **Re-decision records** — fired on three evidence types. Outcome-evidence is a portfolio-wide metric miss against externally committed guardrails. Market-evidence is a category-motion signal reshaping the portfolio. Counterparty-specific evidence is CEO elaboration re-entry across two consecutive cycles.
- **Escalation records** — fired on Charter exit conditions (CEO change, fundamental restructure, scale inflection); the Charter is then closed and a successor Charter is filed with a new `charter_id` per Section 3 §3.3's terminal-state rule.
- **Disclosure-review records** — for the AI-generated portfolio-state summaries embedded into the sync's outputs.
- **Charter-amendment records** — fired on archetype reclassification (a function-native CEO transitioning to product-native as the org matures, etc.) or on cadence amendment.

Conformance signals

- **Level 1** — Charter-conformant: 16 fields populated; the CPO's `accountable_owner` assignment recorded; `mode_declaration = mode-1-with-embedded-mode-2-summary`; `disclosure_metadata_pointer` populated at `fields-completed` per Section 3 §3.4 and Section 6 §6.2; the CEO-co-signature recorded as a structural commitment without reaching for legal

force (the Charter binds the organization to a mechanism, not the CEO to an individual obligation, per Section 3 §3.5's "binds the organization" framing).

- **Level 2** — Mode-disambiguated: every record carries `dispatch_mode` ; embedded AI-generated portfolio-state summaries carry the five required disclosure-block fields; Layer 1 statistical detection runs at 30%; the Mode-Drift Composed Mitigation's Layer 4 Named Human-Attestation Fallback is triggered at every record close because of the Charter's high stakes (per Section 4 §4.8.1).
- **Level 3** — Continuously auditable: the CEO-CPO Charter and its records satisfy the Standard's findability rule; quarterly charter re-signing produces a Charter-amendment record per Section 6 §6.3.1; CEO archetype is named explicitly so the next CPO inheriting the seat reads which interface they inherit.

Disclosure block (Section 4 §4.6)

Disclosure block required under §4.6: yes, conditionally, for embedded AI-generated portfolio-state summaries and continuation-threshold trajectory reads. Whether Article 50 applies: for the deployer to determine. Unless the Charter declares the Section 4 §4.6.1 exclusion, the block attaches wherever the sync record is produced. The five required metadata fields attach at the embed point. The declaring authority is the deploying organization (the company whose CEO and CPO operate the Charter), not the AI-system vendor.

Footer

Worked Example 9.8 — implementation context: This Charter governs strategic direction of the portfolio against the company's plan, plus tenure-defining organizational bets that need CEO air cover, at the one interface the CPO cannot charter into existence unilaterally. Three CEO archetypes break CPO tenure if left un-managed: product-native (deep involvement risks elaboration into scope), function-native (the CEO treats product as a function to direct), and outcome-native (the CEO holds outcomes and grants latitude on means). The Charter's archetype-named field surfaces which CEO is on the other side and which risk pattern dominates. The bring/do-not-bring discipline is sharp: portfolio, bet, decision-system, and organizational state are brought; elaboration on scope, pricing, and ownership re-assignments is not. The "out loud" pushback rule is the structural primitive that prevents Outside decisions from collapsing into the forum.

B.10 CPO-GC Charter – Charter Scaffolding as Schedule-of-Records, Three GC Extensions

Worked Example 9.9 – CPO-General Counsel Decision Interface Charter.

Context

The CPO-GC Charter operates at the interface where the Charter scaffolding (Decision Records, Charters, re-decision triggers, state persistence) is read as a schedule-of-records readable to a legal eye. The interface carries three General Counsel extensions to the Charter scaffolding. The first is **regulatory-acceptance gates inside success criteria**: definition of done is regulatory acceptance, not shipping, for any bet whose completion depends on a regulator. The second is **IP and licensing as a required input class**: patent exposure, third-party OSS obligations, trademark scope, and data-license provenance. The third is **privacy as a metric-integrity constraint**: every metric-integrity choice (definition, window, owner, source of truth, exclusions) may be a processing decision under a privacy law, as the deployer determines. The Charter's load-bearing wall is that the evidentiary posture must hold under external scrutiny, not just internally.

This worked example carries a recursive cross-reference back into the Standard's own Section 6 (Required Artifact Set) and Companion A (Regulatory Cross-References). The reading of Charter scaffolding as a schedule-of-records is the conceptual frame the Standard's Section 6 operationalizes; the General Counsel extensions are the inputs Companion A's regulatory-cross-reference work surfaces against named frameworks (NIST AI RMF, ISO/IEC 42001, the EU AI Act, GDPR/CCPA/LGPD, Caremark-line oversight duties, SOX 404). Section 3's three explicit non-claims (the Charter is not a contract, not a regulatory filing, not a certification) bind here as elsewhere — the Charter structures the inputs counsel and auditors consume; it does not discharge their obligations.

Mode declaration

`mode-1-with-embedded-mode-2-summary` — Mode 1 with the embedded-Mode-2-summary edge case. The canonical decisions at the CPO-GC interface are Mode 1 by structural design. The General Counsel review register is the load-bearing mechanism by which substantive legal posture decisions are made, and counsel's substantive review is precisely the kind of authorship the Standard's not-legal-advice notice keeps with counsel. The CPO is the author of record for the bring-classes: regulatory-acceptance gate status on bets, IP and licensing input classes attached to Charters, and privacy-as-metric-integrity-constraint commitments threaded into the metric-integrity contract. The General Counsel is the reviewer-of-record for the legal posture each Charter takes.

The embedded-summary edge case is engaged for AI-generated regulatory-posture reads dispatched into the General Counsel interface. These include AI-generated jurisdictional-applicability classifications across a portfolio of bets, AI-generated IP-exposure scans against patent landscapes, and AI-generated privacy-impact synthesis from a metric-integrity contract draft. Each embedded sub-output attaches a disclosure block at the embed point per Section 4 §4.7. The structural handling is precisely what Section 4 §4.4 names: a Charter where most decisions are human-authored (and counsel-reviewed) but specific decisions incorporate AI-generated sub-outputs declares the third enumerated value at the Charter altitude.

Schedule of records

This is the Charter where the schedule-of-records concept is most tightly bound to the Standard's own framing.

- **CPO-GC Charter** — the artifact itself, with the three General Counsel extensions named in `inside_decisions`.
- **Decision records** — one per regulatory-acceptance-gate commitment on a bet, one per IP-and-licensing input-class attachment to a Charter, and one per privacy-as-metric-integrity-constraint commitment threaded into the metric-integrity contract.
- **Re-decision records** — fired on three evidence types. Outcome-evidence is a regulator rejecting an acceptance milestone, IP exposure surfacing post-commit, or privacy posture breaching a documented constraint. Market-evidence is a regulatory-regime shift in scope: a new EU AI Act delegated act, a state-level privacy regime change, or a Caremark-line precedent reshaping oversight expectations. Counterparty-specific evidence is a General Counsel review that identifies a posture gap counsel will not sign off on.
- **Escalation records** — fired on the Charter's Emotional Cost trigger: when either the CPO or the General Counsel flinches and the evidentiary posture becomes reactive. The escalation routes back to the CPO-GC interface for posture re-authoring.
- **Disclosure-review records** — required for the AI-generated regulatory-posture reads. The cadence is at minimum quarterly per the General Counsel review register, and synchronized to the regulatory regime's own update cadence where applicable.
- **Charter-amendment records** — fired on regulatory-regime-driven Charter amendments (e.g., the EU AI Act's phased applicability provisions becoming live on a named bet).

Conformance signals

- **Level 1** — Charter-conformant: 16 fields populated; the three General Counsel extensions named in `inside_decisions`; the Charter's audit assignment to the General Counsel ("auditable by the General Counsel for IP and licensing posture") recorded; `mode_declaration`

= mode-1-with-embedded-mode-2-summary ; disclosure_metadata_pointer populated at fields-completed per Section 3 §3.4 and Section 6 §6.2.

- **Level 2** — Mode-disambiguated: every record carries dispatch_mode ; embedded AI-generated regulatory-posture reads carry the five required disclosure-block fields; the Mode-Drift Composed Mitigation's Layer 4 Named Human-Attestation Fallback is triggered at every record close (per Section 4 §4.8.1); Layer 3 peer-review interrupt is also live because General Counsel interface records are the highest-stakes records in this slate.
- **Level 3** — Continuously auditable: the CPO-GC Charter and its records satisfy the Standard's findability rule; the regulatory-acceptance gate status thread runs from each bet's Business-Case artifact (per Section 6's Required Artifact Set treatment) into this Charter's records; the IP-and-licensing input-class records are linkable from the Charters that consume them; the privacy-as-metric-integrity-constraint records thread into the Data and Analytics Charter's metric-integrity audit log (Worked Example 9.4).

Disclosure block (Section 4 §4.6)

Disclosure block required under §4.6: yes, conditionally, for embedded AI-generated regulatory-posture reads, IP-exposure scans, and privacy-impact syntheses. Whether Article 50 applies: for the deployer to determine. The block matters most here, because the artifacts circulate to counsel and auditors who form views about the deploying organization's regulatory posture. The five required metadata fields per Section 4 §4.6.2 attach at every embed point; the declaring authority is the deploying organization; the AI-system identity, jurisdictional-applicability tag, content-type tag, and generation-timestamp pass through any anonymization protocol unchanged per Section 4 §4.6.3's four-of-five rule.

A note on the Layer 4 attestation per Section 4 §4.8.1. The CPO-GC interface is the highest-stakes Charter in this Companion's worked-example slate. The named human attestor at record close attests in their organizational capacity, and the record captures that capacity; the Standard makes no claim about how that capacity affects anyone's personal liability. The Charter's records, including the embedded AI-generated sub-outputs, are designed to support counsel and auditors in their substantive review work; they do not discharge counsel's substantive obligation.

Footer

Worked Example 9.9 — implementation context: This Charter governs the CPO-General Counsel interface, where the Charter / Decision Record / re-decision-trigger architecture is read as a schedule-of-records that must hold under a legal eye. It carries three General Counsel extensions to the Charter scaffolding: regulatory-acceptance gates inside success

criteria (definition of done is regulatory acceptance, not shipping), IP and licensing as a required input class (patent, OSS, trademark, data-license provenance), and privacy as a metric-integrity constraint (each integrity choice may be a processing decision under a privacy law, as the deployer determines). The General Counsel is the reviewer-of-record for the legal posture each Charter takes. The evidentiary posture must hold under external scrutiny, not just internally.

B.11 Section 9 Closure

The nine worked examples above are not exhaustive of the Charter mechanism's reach. The mechanism applies to any function-specific decision interface, where the seven director-altitude surfaces illustrated here are a representative set, not a closed list. It applies to any executive-altitude interface the CPO operates: the CEO-CPO and CPO-GC interfaces are illustrated here, and the CPO-CTO, CPO-CRO, CPO-CFO, CPO-Board, CPO-COO, CPO-CMO, and CPO-Chief People Officer interfaces follow the same pattern. It also applies to the install-order Charters a smaller organization files at sub-Series-C scale. Companion B surfaces nine of these as the worked-example slate that demonstrates the Standard's primitives operating across director, executive, and board altitudes. A deployer whose decision class falls outside the nine instances above operates the Charter mechanism per Section 3 directly. Section 3's primitives do not depend on Companion B instantiation.

The Mode declarations across the nine examples cluster into two groups. Three are pure Mode 1 Charters (PMM, Design, Business Development), where the canonical decision content is human-authored at director altitude with the AI worker as Charter-conformance enforcement. Six are `mode-1-with-embedded-mode-2-summary` Charters (Engineering, Data and Analytics, Customer Success and Value Realization, Product Operations, CEO-CPO, CPO-GC), where the canonical decision content is human-authored but specific sub-outputs are AI-generated and dispatched into the record at embed points carrying disclosure blocks. No worked example in this slate carries a pure Mode 2 declaration. This is consistent with the Charter mechanism being authored from director-and-above seats, where the substantive decision content is human-authored as a default. A future expansion of this Companion could surface decision classes that dispatch under pure Mode 2: for example, a deployer-authored Charter binding AI-generated draft messaging variants in the demand-gen-led PMM archetype, or an AI-authored synthesis dispatched into a sub-Charter in the embedded-data-team archetype. Such an expansion would extend the slate, and the Standard's primitives accommodate it without modification.

The records are input, not evidence. The Standard informs frameworks without satisfying them. Conformance is self-declared; no body certifies it. It is not legal advice and not a regulatory substitute.

Decision Provenance Standard™ v1.2 — Reading Edition (rev. 10). Open standard under CC-BY 4.0. Not a certified product. Founding Steward: Yohay Etsion; institutional Steward: Etsion Brands Ltd.