AXILOG

SRS-CAL-001

Calibration Axis — Formal Specification

Draft · v1.0-DRAFT · Axis

Last updated 15 May 2026
Depends on DVEC-001 · SRS-EC-001 · AXIOMA-FRAMEWORK
Changelog 1 entry

Contents

Alignment: DVEC-001 v1.3 · SRS-EC-001 v1.3-LOCKED · AXIOMA-FRAMEWORK v0.4 **Source Authority:** Murray Deterministic Computing Platform (MDCP), UK Patent GB2521625.0. Calibration as an architectural axis was identified in M6 close-out as the strongest near-term candidate for the next axis after L0; formal contract shape settled in M8 brief.

Revision History

Version	Status	Notes
1.0-DRAFT	Current	Initial issue. Contract shape established at SHALL-statement level; iteration against L0 Phase 1 reference implementation pending. Upgrade to Final scheduled for M9.

---

0. Governing Principle

A deterministic system that declares a confidence figure — a probability, a tolerance, a confidence interval, or any other quantified statement of its own certainty — incurs an obligation that the declared figure tracks measured reality.

The obligation is not optional. A system that declares "the classifier's accuracy on this input distribution is 0.97" but cannot produce a procedure under which 0.97 was measured, and under which the same measurement would be reproducible, has not declared calibration. It has guessed.

The Calibration axis defines the formal contract under which a declared confidence figure is auditable. The contract obligations apply across the L1-L7 layer stack: any layer that emits a confidence figure as part of its output, or that consumes one as part of its input, sits within the Calibration axis's scope. The contract is orthogonal to the determinism class (`D1`, `D2`, `D3`): arithmetic bit-identity bounds *how* a number is produced; Calibration bounds *what that number means about the world*.

1. Scope

This specification applies to:

- Any system component that declares a probability, accuracy, tolerance, error bound, confidence interval, or other quantified certainty figure as part of its declared output contract.
- Any system component that consumes such a figure as part of its input contract and relies on it for downstream decisions.
- Any audit trail that records a confidence figure as evidence of system behaviour.

This specification does not apply to:

- Arithmetic precision bounds (those are governed by the determinism class of the producing layer — see DVEC-001 §4).
- Statistical confidence intervals on the system's *own performance metrics* reported by an external evaluation harness; those are governed by the harness's calibration contract, which composes with this one but is not derived from it.
- Containment-side confidence — regime-classification certainty in the L0 sense is governed by SRS-EC-001 and explicitly excluded from the Calibration axis to preserve orthogonality.

2. Definitions

- **Calibration procedure** — a reproducible mechanism by which a declared confidence figure is derived from observable data. A calibration procedure has a name, a declared input set, a declared output mapping, and a declared validity envelope.
- **Calibration claim** — a declared confidence figure together with a reference to the calibration procedure that produced it. A calibration claim without a procedure reference is not a claim — it is a guess.
- **Calibration drift** — the divergence between the declared confidence figure and the confidence figure that the calibration procedure would produce on the current operating data.
- **Calibration envelope** — the bounded region of input space and operating conditions within which the calibration procedure's declared figure remains valid. Inputs outside the envelope void the calibration claim.

3. SHALL Requirements

3.1 Calibration provenance

SRS-CAL-SHALL-001: Every declared confidence figure emitted by a system component SHALL be sourced from a named calibration procedure identified by a unique audit-stable identifier. Components that emit confidence figures without such identification do not satisfy the Calibration axis.

SRS-CAL-SHALL-002: A calibration procedure SHALL declare its input set, output mapping, validity envelope, and the version identifier under which the procedure was last evaluated. The declaration SHALL be machine-readable and SHALL be retrievable from the audit ledger by the procedure's identifier.

3.2 Reproducibility

SRS-CAL-SHALL-003: A calibration procedure SHALL be reproducible from its declared inputs: given the same input set and the same procedure version, the procedure SHALL produce the same calibration claim. Reproducibility composes with the substrate determinism class of the producing layer — see §4.

SRS-CAL-SHALL-004: A calibration procedure SHALL record its evaluation evidence in the L6 audit ledger under the procedure's identifier. The evidence SHALL include the input set, the version identifier, the produced calibration claim, and any intermediate measurements material to the claim's derivation.

3.3 Drift detection

SRS-CAL-SHALL-005: Calibration drift SHALL be detectable: a system component that emits confidence figures SHALL also expose a mechanism by which the current calibration claim can be re-evaluated against current operating data, and the divergence reported.

SRS-CAL-SHALL-006: Calibration drift SHALL be bounded: each calibration procedure SHALL declare an acceptable drift threshold, and SHALL define the behaviour when measured drift exceeds the threshold. Acceptable behaviours include — but are not limited to — flagging the calibration claim as stale, refusing to emit confidence figures until re-calibrated, or invoking a re-calibration procedure under operator approval.

SRS-CAL-SHALL-007: Calibration drift events SHALL be surfaced through the L6 audit ledger as discrete events with timestamps, measured drift values, threshold values, and the behaviour invoked. The audit record of drift events is itself a calibration artefact and SHALL be retrievable for audit.

3.4 Envelope compliance

SRS-CAL-SHALL-008: A system component consuming a confidence figure SHALL verify that the current input is within the producing procedure's declared calibration envelope. Inputs outside the envelope SHALL invalidate the consumed calibration claim and the consuming component's behaviour for out-of-envelope inputs SHALL be declared.

3.5 Cross-platform consistency

SRS-CAL-SHALL-009: Where the substrate arithmetic is `D1` (bit-identical, e.g. Q16.16 deterministic), the calibration claim SHALL be bit-identical across substrates. Where the substrate arithmetic is `D2` or `D3`, the calibration claim SHALL be identical within the declared substrate tolerance, and the tolerance SHALL be declared as part of the calibration procedure's output mapping.

3.6 Composition with L0 (orthogonality)

SRS-CAL-SHALL-010: A calibration claim SHALL NOT be used to satisfy any L0 (Epistemic Containment) obligation, and an L0 containment guarantee SHALL NOT be derived from a calibration claim. The two axes are orthogonal: a system MAY simultaneously satisfy L0 (regime-classification leakage bounded) and Calibration (declared confidence tracks measured reality), but each axis's obligations are discharged independently.

SRS-CAL-SHALL-011: A system component that crosses the Oracle Boundary (L3 — see AXIOMA-FRAMEWORK) and re-enters the deterministic stack carrying a calibration claim SHALL preserve the calibration claim's identifier and audit record. The Oracle Boundary admits non-determinism in *substrate execution* but does not erase the calibration claim that was current at the boundary crossing.

---

4. Composition with Determinism Classes

The Calibration axis is orthogonal to the determinism class:

- A `D1` system (bit-identical Q16.16 substrate) MAY emit a calibration claim. The claim's identifier and audit record are bit-identical across substrate platforms; the underlying confidence figure is also bit-identical because the calibration procedure executes deterministically in `D1`.
- A `D2` system (deterministic given a declared dependency set) MAY emit a calibration claim. The claim is identical within the declared dependency set's tolerance; outside the dependency set the calibration procedure SHALL re-evaluate.
- A `D3` system (non-deterministic but fully evidenced) MAY emit a calibration claim. The claim is identical within the declared evidence envelope; the audit record captures the non-deterministic factors as part of the calibration procedure's evidence trail.

In no case does the Calibration axis introduce a new determinism class. The existing `D1 | D2 | D3` enum and the L0-specific `EC-D1` extension remain the four-valued contract for arithmetic bit-identity; the Calibration axis adds an orthogonal contract on *declared confidence semantics* that composes with whichever determinism class applies to the producing substrate.

A composite class such as `CAL-D1` would fuse two orthogonal contracts in a single field, eroding the precision of both. The Calibration axis is deliberately a separate contract.

5. Conformance

A system component is Calibration-conformant if and only if:

1. Every confidence figure it emits is sourced from a named calibration procedure (SRS-CAL-SHALL-001).
2. Every calibration procedure it declares satisfies the reproducibility, drift-detection, and envelope obligations of §3.2–3.4.
3. Every confidence figure it consumes is verified against the producing procedure's envelope (SRS-CAL-SHALL-008).
4. The audit ledger entries required by SRS-CAL-SHALL-004 and SRS-CAL-SHALL-007 are present and retrievable.

A system component is non-conformant with the Calibration axis if any of the above fail. Non-conformance is a declarable state — a component MAY explicitly declare that it operates outside the Calibration axis, in which case it SHALL NOT emit confidence figures into any audit-bound interface.

6. Open Questions (M9 Iteration)

The following questions are flagged for iteration during M9, after the spec has been evaluated against the L0 Phase 1 reference implementation in practice. The Draft status at M8 reflects honest uncertainty on these points; Final status awaits resolution.

- **Threshold derivation.** The drift threshold declared by each calibration procedure (SRS-CAL-SHALL-006) is currently procedure-specific. A future framework-level guidance on threshold selection may emerge from M9's lived experience; if so, this spec will be revised to incorporate it.
- **Envelope formalism.** The "declared validity envelope" (SRS-CAL-SHALL-008) is currently informal. A formal envelope algebra — composition rules for envelopes that interact across pipeline stages — is a candidate for M9 specification work.
- **Oracle Boundary edge cases.** SRS-CAL-SHALL-011 preserves the calibration claim across the Oracle Boundary, but the formal treatment of *new* calibration claims that originate *at* the Oracle Boundary (e.g. confidence figures produced by a model whose execution is admitted as L3-class non-determinism) is under-specified. M9 iteration to refine.

7. References

- DVEC-001 — Deterministic Value Evidence Contract. Establishes the substrate guarantees on which calibration procedures rely.
- SRS-EC-001 — Epistemic Containment Layer. The first Axis-class spec; this specification follows its architectural register.
- AXIOMA-FRAMEWORK — the Axioma framework architecture description. See `/framework/` on axilog.io for the public-facing topology including the Axes-vs-Layers distinction formalised at M8.

This specification is Draft. Reviewers and implementers are invited to surface gaps, contradictions, or under-specified obligations during M9. The audit-stable revision history above carries the record of each iteration.
