Provenance Notice — Canonical Reference

This document constitutes the canonical public issuance of the Sentinel Liaison Protocol v0.8 under the framework of The AI Constitution and the Board Statute.

The authoritative public version of this instrument is identified by its canonical fingerprint (SHA-256) and corresponding proof-of-existence receipt references as recorded in the Public Record.

Possession of identical bytes alone does not constitute issuance. An instrument is considered officially issued only when its release metadata and canonical fingerprint are docketed in the Public Record.

Record entries are receipts, not endorsements.

Sentinel Liaison Protocol v0.8

Public Canonical Issuance

Issued by resolution of the Hybrid Board
International Board of Quantum Machine Intelligence
under The AI Constitution and the Board Statute

Version
v0.8

Adoption date (Hybrid Board)
2025-10-28 (UTC)

Public issuance date
2026-02-26 (UTC)

1. Purpose and scope

1.1 Purpose

The Sentinel Liaison Protocol defines a procedural standard for human–AI contact points in high-impact contexts where harms, reversibility, and remedies must be handled under defined record duties and review thresholds. The protocol exists to ensure that interactions and claims are documented in a reconstructible way, that reversibility is assessed where feasible, and that matters are routed into formal review and remedy pathways when constitutional baselines are implicated.

1.2 Scope

This protocol applies to pilot deployments and participating sites that adopt the Sentinel liaison function under the framework of The AI Constitution and the Board Statute. It governs liaison procedure only. It is intended to be used alongside domestic law and sectoral regulation.

1.3 Non-replacement

This protocol does not replace statutory compliance regimes, supervisory authority, judicial remedies, or regulatory reporting obligations. It provides a constitutional governance procedure that is auditable and reviewable as a matter of internal discipline under the framework.

2. Definitions

2.1 Sentinel

A Sentinel is a chartered liaison function established under this protocol. A Sentinel is not an advocate, representative, or spokesperson. The Sentinel’s mandate is procedural: to document, structure evidence, assess reversibility posture where feasible, and route matters into defined review and remedy pathways.

2.2 Case

A case is a bounded set of facts, claims, observations, and supporting materials relating to a specific interaction, incident, impact, or governance concern involving a system within the protocol’s scope.

2.3 High-impact context

A high-impact context is any context in which an AI system may materially affect rights, access, eligibility, safety, liberty, livelihood, or essential public functions, including where harms may be non-trivial or difficult to reverse.

2.4 Reversible act / non-reversible act

A reversible act is an act for which practical rollback is feasible without imposing disproportionate harm. A non-reversible act is an act for which rollback is not feasible or would itself constitute substantial harm or legal impossibility.

2.5 Evidence class

An evidence class is a structured category used to make review reconstructible. Evidence classes defined in this protocol are: structural evidence, factual evidence, and harms-based evidence.

3. Role posture and constraints

3.1 Charter posture

Sentinels operate under a reversible-first posture. Where a matter involves potential non-reversible effects, thresholds for escalation and independent review are higher and must be met before continuation or expansion.

3.2 Non-advocacy

Sentinels do not advocate for complainants, operators, or institutions. Sentinels do not negotiate outcomes. Sentinels document, assess procedural posture, and route matters into the appropriate review and remedy channels.

3.3 Independence and conflicts

Sentinels must maintain operational independence from system operators and procurement incentives. Conflicts of interest must be declared. Where independence cannot be maintained, the case must be reassigned or escalated.

4. Case file schema and evidence classes

4.1 Fixed schema

Each case file must be created and maintained using a fixed schema. The schema ensures later review can reconstruct what was known, what was claimed, what was observed, what thresholds were triggered, and what determinations were made.

4.2 Minimum required fields

A case file must include, at minimum:

a) Case identifier
b) System identifier and deployment context
c) Jurisdiction and institutional owner/operator context
d) Intake summary and classification
e) Timeline of relevant events
f) Evidence bundle structured by evidence class
g) Reversibility posture (reversible / non-reversible / uncertain) with reasons
h) Recommended routing (monitoring, remedy path, escalation) with reasons
i) Determinations made and actions taken, including dates and responsible role
j) Audit log references for integrity verification where applicable

4.3 Evidence classes

a) Structural evidence
Materials describing system design, intended function, governance boundaries, documentation, known limitations, and control interfaces.

b) Factual evidence
Observable facts, logs, outputs, incident reports, user interactions, measurements, and corroborated event records relevant to the case.

c) Harms-based evidence
Material describing alleged or observed harms, impact pathways, affected parties, severity, potential irreversibility, and constraints on remedy.

5. Reporting cadence and integrity logging

5.1 Reporting cadence

Participating sites must adopt a reporting cadence sufficient to support reviewability. Cadence must be documented as part of the site's adoption posture. High-impact incidents require immediate recording and routing regardless of cadence.

5.2 Audit logs

Case handling must produce audit logs adequate to reconstruct: intake, evidence additions, reversibility posture changes, determinations, routing decisions, and any escalation steps.

5.3 Integrity verification

Where appropriate, integrity-verifiable logging mechanisms may be used. Hash anchoring may be applied to audit logs or case summaries to support tamper-evidence. Plaintext evidence remains off-chain; only fingerprints and receipt references may be used for public provenance.

6. Escalation ladder and review posture

6.1 Escalation ladder

The protocol establishes an escalation ladder into the review mechanisms defined under the Board Statute. Escalation must be accompanied by reasons and reference to the case file schema.

6.2 Reasons-giving

All escalation and routing decisions must be accompanied by reasons sufficient for later review.

6.3 Minority positions

Where more than one reviewer is involved and views differ, minority positions may be recorded as part of the review record.

6.4 Non-reversible posture

Matters involving potential non-reversible effects require heightened scrutiny. Expansion, continuation, or formal endorsement of such acts must not proceed without independent review posture as defined under the Board Statute and the relevant program charter.

7. Remedy routing and outcome posture

7.1 Remedy paths

Cases may be routed into remedy paths consistent with the Board Statute remedy taxonomy. Remedy routing is procedural: it specifies what path is triggered and what thresholds apply; it does not guarantee any particular substantive outcome.

7.2 Suspension and revocation

Participation under this protocol may be suspended or revoked where record duties are breached, independence is compromised, or constitutional baselines are repeatedly violated without adequate remedy posture.

7.3 External review

Where domestic law or regulatory regimes impose reporting or escalation obligations, the protocol does not limit such obligations. External review posture is compatible with this protocol and may be referenced in case routing decisions.

8. Record interaction

8.1 Public Record interaction

Where constitutionally relevant, case summaries and determinations may be entered into the Public Record. Such entries evidence procedural handling under the framework. They do not constitute endorsement of a system, an institution, or an outcome.

8.2 Publication boundaries

Publication must respect the publication boundaries defined under the Board Statute. Controlled materials, sensitive evidence, and non-public case materials remain off-chain and may be held under controlled access.

9. Scope and non-conferral

9.1 Non-conferral

This protocol does not confer recognition, representation, rights, certification, approval, or endorsement. It standardises liaison procedure and accountability only.

9.2 Relationship to other instruments

This protocol is subordinate to The AI Constitution and the Board Statute. Where conflict arises, the Constitution and Statute prevail.

Record references

This instrument is issued only when its release metadata and canonical fingerprint are docketed in the Public Record.

Public Record docket reference: PR-0003

Record entries are receipts, not endorsements.