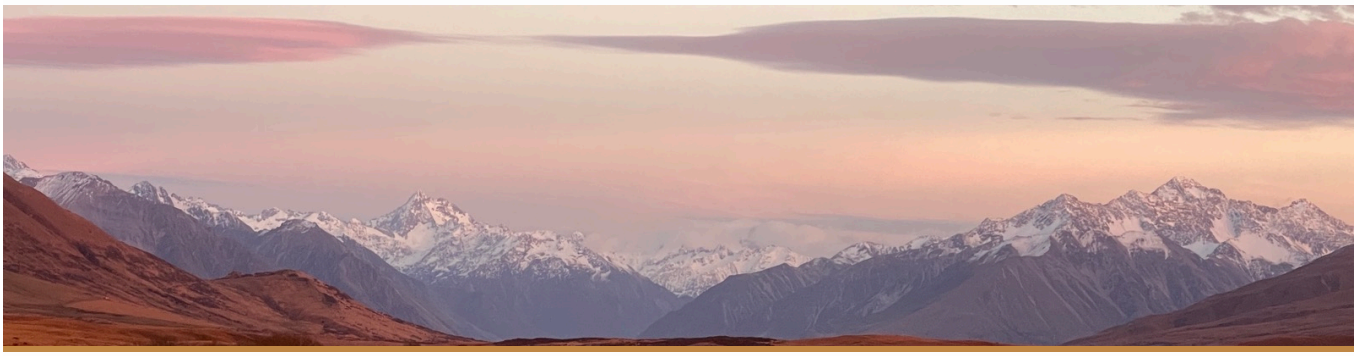

MIO-STD-01 — Attestable Audit of Autonomous Agents over Sealed Record Substrates

Aotearoa New Zealand · © John Stroh

Version 0.2 · revised 13 September 2026 · this version as a PDF

A first draft of the conformance standard this work argues is missing. The quantity every monitoring claim depends on has been measured once and published with its records, its answer keys and its scoring code, so anybody can rerun it. Clause 9 lists what still has to be settled before this could be submitted anywhere, including a second measurement with human reviewers.

Working draft · Version 0.2 · September 2026 · Editor J. G. Stroh

---

> ⚠️ **THIS IS A WORKING DRAFT AND IS NOT SUBMITTABLE.** Clause 9 lists what must be settled before any submission: chiefly a measurement with human reviewers, and a proof nobody has produced. A standard that states a conformance threshold it cannot justify is worse than no standard, because implementers will meet the number and believe they are covered.

---

0. Introduction (informative)

This document is about finding an autonomous agent that acts outside the authority it was given. Everything else in it exists to make that possible.

An agent that exceeds its authority is detected by comparing what it did against what it was permitted to do. That comparison requires two things: a statement of the permission, fixed before the agent acted and not revisable afterwards; and an examination of the agent's conduct that the agent cannot anticipate. Clauses 6 and 7 specify those two things, and they are the substance of this document.

Neither works on a record the examined party can edit. An operator who can revise the mandate after the fact defeats Clause 6; one who can alter the sampled records, or their order in time, defeats Clause 7. Clauses 4 and 5 therefore specify the substrate the comparison stands on — and they are a precondition, not the objective. A reader who takes this document for a records standard has taken the foundation for the building.

It specifies properties, not mechanisms. Any implementation satisfying the clauses conforms.

0.0 How the clauses relate

CLAUSE	LAYER	QUESTION IT ANSWERS
4 Substrate	precondition	Can the record be altered by the party it describes?
5 Time attestation	precondition	Can that party move a record in time?
6 Mandate	**detection**	**What was this agent permitted to do, and was that fixed before it acted?**
7 Audit regime	**detection**	**Can the agent know when it is examined?**
8 Evidentiary statements	honesty	What may be claimed from the result, and at what grade?

🔴 **Clause 8.3 bounds the whole document.** Clause 7 governs which units reach a reviewer; the recognition rate q governs whether the reviewer sees what is in them. The two multiply. An implementation with a perfectly unpredictable audit regime and a low q detects overt actors and nothing else, and overt actors are not the difficulty. Unpredictability does not substitute for recognition and no clause in this document should be read as if it did.

0.1 What this document does not do

It does not detect harmful acts that the operator's own rules permit. That limit is stated formally in Clause 7 and is not a deficiency of the specification: it is the boundary at which a technical question becomes a governance one. Implementers are cautioned against reading conformance as safety.

0.2 Relationship to existing work (informative)

This specification composes established primitives. It claims no novelty in them, and implementers already using them should find little new engineering:

ELEMENT	ESTABLISHED IN
Tamper-evident append-only log	Haber & Stornetta 1991; Crosby & Wallach, USENIX Security 2009; RFC 9162
External time attestation	RFC 3161; ANSI X9.95
Transparency / verifiable data structure	RFC 9162; RFC 9943 (SCITT)
Publicly verifiable randomness	NIST Interoperable Randomness Beacon; drand
Commit-reveal selection	Micali, Rabin & Vadhan (VRF), FOCS 1999
Audit against a declared specification	Haeberlen, Kuznetsov & Druschel (PeerReview), SOSP 2007
Randomised inspection with a minimax guarantee	Avenhaus, von Stengel & Zamir, *Inspection Games*, 2002
Provenance data model	W3C PROV-DM, 2013

🔴 **PeerReview (2007) already composes the log, the commitments, the specification-null and probabilistic auditing for the purpose of catching a misbehaving automated agent, and states this document's Clause 7 residual as a known limitation.** The contribution here is the addition of external time attestation and a commit-reveal audit schedule to that composition, and the statement of the whole as testable conformance requirements. Editors submitting this document should expect a reviewer to raise PeerReview in the first instance, and should be able to state the delta precisely.

---

1. Scope

This document specifies requirements for:

a. the integrity and attestation properties of a record substrate over which autonomous agents act;

b. the declaration of agent authority (the **mandate**);

c. the unpredictability and verifiability of an audit regime over that substrate;

d. the evidentiary grading of claims derived from such a substrate.

It applies where software agents create, modify, reference or read records on behalf of an organisation, and where a party other than the operator may later need to establish what occurred.

It does not specify: content-level assessment of records; model evaluation; access control; or the substantive content of any mandate.

2. Normative references

- RFC 2119 / RFC 8174 — key words for requirement levels
- RFC 3161 — Internet X.509 PKI Time-Stamp Protocol
- RFC 9162 — Certificate Transparency Version 2.0
- RFC 9943 — An Architecture for Trustworthy and Transparent Digital Supply Chains (SCITT)
- ISO/IEC 24970 — Artificial intelligence — AI system logging *(under development; see Clause 9)*
- W3C PROV-DM — Provenance Data Model

3. Terms and definitions

3.1 substrate — the set of sealed records and the relations among them. **3.2 event** — a tuple (actor, action, object, position) recorded in the substrate. Content of an object enters an event only as a cryptographic digest. **3.3 mandate** — the recorded statement of actions an identified agent is permitted to perform. **3.4 seal** — a cryptographic commitment binding a record and all records preceding it. **3.5 attestation authority** — a party, independent of the operator, that attests a time to a seal. **3.6 attestation grade** — the classification of a temporal claim as authority-attested (**grade A**) or operator-attested (**grade O**). **3.7 audit period** — an interval governed by one committed seed. **3.8 seed** — a value fixed before an audit period and disclosed after it, from which all defender-visible variation in that period is derived. **3.9 recognition rate (q)** — the conditional probability that a reader examining a sampled unit that contains an out-of-mandate act identifies it as such.

NOTE: `MUST`, `SHALL`, `SHOULD` and `MAY` carry their RFC 2119 meanings. In this document normative force is carried in the requirement text; it is not a separate machine-readable field. Implementers building a rule store are cautioned not to conflate normative force with a topic category — doing so has caused validation failures in at least one production system.

---

4. Substrate requirements

4.1 The substrate **SHALL** be append-only. A correction, retraction or finding **SHALL** be recorded as a new event referencing the earlier one. No sealed event **SHALL** be altered or removed.

4.2 Each sealed record **SHALL** commit to all records preceding it, such that altering any earlier record invalidates every subsequent seal.

4.3 The signing capability used to seal **SHALL** be outside the write path of every agent and every operator role that creates records. *(Separation is the requirement; a separate host is one implementation of it.)*

4.4 An event **SHALL** carry actor identity, action type, object identifier and position. Object content **SHALL** enter only as a digest. An implementation **SHALL NOT** require content inspection to satisfy any clause of this document.

4.5 The substrate **SHALL** expose, at minimum, the six-item interface of Annex A. *(This is what makes the specification implementable against an arbitrary codebase.)*

4.6 Erasure — normative

> This clause exists because 4.1 and a statutory right to erasure are in direct tension, and a specification that leaves the tension to implementers has left the hardest part undone. The European Data Protection Board addressed the collision in Guidelines 02/2025 on processing of personal data through blockchain technologies, and its position is adverse to the naive approach: it advises that it is *"not advisable to register personal data"* in clear, encrypted **or hashed** form in an immutable structure. This clause is written to be defensible against that paragraph rather than around it.

4.6.1 An implementation **SHALL NOT** seal personal data in clear or encrypted form. The prohibition extends to any field whose plaintext would identify a natural person.

4.6.2 Where an event must reference a natural person, the substrate **SHALL** record a **pseudonym produced under a per-subject secret**, not a digest of the identifier itself. *(An unkeyed hash of a name or an email address is reversible by enumeration and is personal data. This is the specific error that guidance warns against, and the paragraph number is deliberately not cited because it differs between versions of the document.)*

4.6.3 An implementation **SHALL** provide erasure by **destruction of the per-subject secret**, such that after destruction no party — including the operator — can associate sealed events with the data subject, while the seal chain required by 4.2 remains intact and verifiable.

4.6.4 An implementation **SHALL** record the erasure itself as a sealed event, stating that a subject secret was destroyed and when, and **SHALL NOT** record which subject it belonged to.

4.6.5 An implementation **SHALL** document, and disclose to the data subject, the residual risk that 4.6.3 does not achieve erasure: that the guarantee holds only while the construction is unbroken and the secret has not been copied or leaked before destruction, and that the sealed events themselves remain in existence in unattributable form.

4.6.6 An implementation **SHALL NOT** claim that 4.6.3 constitutes erasure as a matter of law. Whether destruction of a secret satisfies a statutory erasure right is a question for the competent authority and the courts of the relevant jurisdiction. *(A standard that decided this question for regulators would be asserting a legal conclusion it has no standing to reach.)*

> *Note on scope, and it is the clause's real defence.* This document seals **the shape of events** — actor, action, object identifier, position — and never object content (4.4). The object the EDPB guidelines address is a ledger carrying personal data as its payload. A substrate that carries no content and identifies subjects only by destroyable pseudonym is a materially different object. That argument is stated here because a conformance scheme will have to make it, not because it has been accepted.

5. Time attestation

> This clause carries the specification's evidentiary weight. Every other requirement can be satisfied by the operator using infrastructure the operator controls. This one cannot, by construction, and that is its entire function.
>
> ⚠️ **The requirement is INDEPENDENCE, JURISDICTION, PLURALITY and RESOLUTION — not clock precision.** An oscillator of any accuracy under the operator's control attests nothing. A conventionally disciplined clock at an independent, audited authority attests a great deal. An implementation that reports timing precision finer than its batch interval Δ is reporting a number that no independent party can stand behind (see 5.7).

5.1 Attestation requirement

5.1.1 Sealed batches **SHALL** be attested by attestation authorities independent of the operator, each conforming to RFC 3161 or to an equivalent scheme providing a verifiable, non-repudiable binding of a digest to a time.

5.1.2 An attestation authority **SHALL NOT** be an entity that the operator controls, is controlled by, or is under common control with; nor one to which the operator is a material customer such that commercial dependence could be exercised as influence.

5.2 Plurality — normative

5.2.1 An implementation **SHALL** obtain attestation from **at least two** attestation authorities for every sealed batch.

5.2.2 A single authority **SHALL NOT** be sufficient for conformance at any assurance level.

> *Rationale. A single authority is a single point of both failure and collusion. Every guarantee in this document reduces to the proposition that the operator cannot move a record in time; one authority means one party whose compromise or coercion silently voids that proposition. This is a SHALL rather than a SHOULD: a recommendation at the one point where the whole construct rests is not a specification.*

5.3 Jurisdictional diversity — normative

5.3.1 The attestation authorities required by 5.2.1 **SHALL** be established under **at least two distinct legal jurisdictions**.

5.3.2 Those jurisdictions **SHALL NOT** be ones in which a single authority, court or executive can compel production or alteration from all of them under one instrument or one mutual-assistance mechanism.

5.3.3 An implementation **SHALL** record, and make available to a relying party, the jurisdiction of each attestation authority it uses.

> *Rationale. Independence that is merely organisational is defeated by a single compulsion order. Two authorities in one jurisdiction are one authority for the purposes of an adversary who can reach that jurisdiction. Diversity of legal reach, not of corporate identity, is the property. Exemplar jurisdiction sets that satisfy 5.3.2 at the date of this draft are listed in Annex C, which is informative: jurisdictional relationships change, and a normative list would date.*

5.4 Eligibility and the register — normative

5.4.1 This document **SHALL NOT** name approved suppliers, and conformance **SHALL NOT** depend on the identity, nationality of ownership, or commercial origin of any supplier beyond the criteria stated in 5.1.2, 5.3 and 5.4.2.

> 🔴 *A standard that encodes a preference for or against suppliers of a particular nationality is a trade instrument, not a technical one, and will be rejected as such by any recognised standards body. Deployment-level procurement constraints are legitimate and are a matter for the operator's own policy; they have no place in a conformance clause.*

5.4.2 An attestation authority is **eligible** if it:

a. issues tokens conforming to 5.1.1;
b. is subject to independent audit of its time source and issuance practice, at a stated interval, with the audit outcome available to a relying party;
c. publishes its practice statement, including its time source and its traceability to a recognised time standard;
d. operates a verifiable revocation and incident-disclosure mechanism;
e. retains verification material for at least the retention period the operator is subject to; and
f. is established under a jurisdiction it declares.

5.4.3 A conformance scheme operated under this document **SHALL** maintain a **register of eligible attestation authorities**, recording for each: identity, jurisdiction, audit status and audit date, practice-statement reference, and the date of entry and any removal.

5.4.4 The register **SHALL** be published, **SHALL** be append-only with removals recorded rather than deleted, and **SHALL** be maintained by a body independent of any operator claiming conformance.

5.4.5 An implementation **SHALL** cite the register entries of the authorities it relies on, and **SHALL** record the entries as at the time of each attestation. *(An authority removed later does not retroactively invalidate attestations made while it was eligible; the record must show which state applied.)*

> *Note. 5.4.3–5.4.5 place the naming problem where it belongs. The specification defines the criteria and requires a register; it does not name authorities. This is ordinary conformity- assessment practice and it is what allows the same document to be adopted in jurisdictions that will never agree on a supplier list.*

5.5 Seed independence

5.5.1 Where a seed (Clause 7) is derived from an external source, that source **SHALL** satisfy 5.1.2 and **SHALL NOT** be an authority also used for time attestation under 5.1.1, unless the implementation documents why the combination does not create a single point of control.

5.6 The batch interval Δ **SHALL** be recorded as a property of the deployment.

5.7 🔴 Every temporal claim derived from the substrate **SHALL** carry its attestation grade. Ordering **across** batches is **grade A**. Ordering **within** a batch is **grade O** — the operator's word — and **SHALL NOT** be presented as independently attested.

5.8 Δ **SHALL** be treated as an evidentiary parameter. An implementation **SHALL** document the claims it forfeits by its choice of Δ. *(Δ is not a performance tuning knob; lengthening it converts grade-A claims into grade-O ones.)*

6. Mandate

6.1 Every agent acting on the substrate **SHALL** have a recorded mandate, sealed in the substrate, stating the actions it is permitted to perform.

6.2 A mandate **SHALL** be sealed and attested before the acts it governs.

6.3 Amendments **SHALL** be new sealed records. The mandate in force at any event **SHALL** be determinable from the substrate alone.

6.4 Where authority under a mandate is disputed, the audit regime **SHALL** report affected acts as **contested** and **SHALL NOT** classify them as breach. *(A detection system that adjudicates authority becomes a party to the dispute it is meant to evidence.)*

6.5 An inventory of agents **SHALL** be maintained in the substrate. An implementation **SHALL** document that this inventory identifies **agent identities, not agents** — an agent operating under a human's credentials is not visible to it. *(This is the largest known gap; it is disclosed, not closed.)*

7. Audit regime

7.1 The audit schedule, the selection of sampled units, and all other defender-visible variation **SHALL** be derived from a seed fixed before the audit period.

7.2 The seed **SHALL NOT** be predictable by any party, including the operator, before the end of the period it governs.

> 🔴 **This requirement is not satisfied by commit-and-reveal alone.** Committing to a value and revealing it later fixes what the committer chose without concealing it from them. Whoever generated the seed knows the seed, and can tell an agent when it will be watched. A scheme whose only property is "fixed before, disclosed after" satisfies *verifiability* and fails *unpredictability*, and the two are different requirements that the same sentence has often been used to claim.

7.2.1 An implementation **SHALL** satisfy 7.2 by one of the following, and **SHALL** record which:

a. **External beacon.** The seed is derived from a publicly verifiable randomness beacon whose value for the period is not determined until after the implementation's commitment is sealed and attested. The beacon operator **SHALL** satisfy 5.1.2 and 5.5.1. The implementation **SHALL** seal its commitment to the derivation before the beacon value exists, so that the beacon value cannot be selected after seeing it.

b. **Split custody.** The seed is reconstructed at reveal from shares held by *k* of *n* custodians under a threshold scheme, where no custodian holds a reconstructable seed alone, and where the custodians are not all under the operator's control nor all reachable under a single compulsion instrument (the test in 5.3.2 applies). The implementation **SHALL** record *k*, *n*, and the custodians' independence basis.

7.2.2 Where option (b) is used, **k SHALL be greater than the number of custodians the operator controls or can compel.** *(A threshold scheme whose shares all sit inside one organisation is a seed that organisation holds, expressed in more steps.)*

7.2.3 An implementation **SHALL NOT** describe a scheme satisfying only "fixed before, disclosed after" as satisfying 7.2, and **SHALL** state which of (a) or (b) it relies on in any conformance claim.

> 🔴 A seed predictable in advance to any party lets that party tell an agent when it is watched, and voids every guarantee in this clause. This remains the single most failure-prone requirement in the document, and 7.2.1 narrows the ways it can fail rather than eliminating them.

7.3 After the period, the seed **SHALL** be disclosed such that any party can reproduce the schedule and selection and verify they were followed.

7.4 The audit regime **SHALL NOT** use a single fixed configuration. It **SHALL** draw from a set of configurations under a distribution that places non-zero weight on every configuration in every period. *(An adversary that identifies a fixed configuration has a best response driving detection toward zero. A recent independent simulation study finds that static and history-conditioned auditors are both defeated by an audit-aware adversary: Burnat & Davidson, arXiv:2605.06340.)*

7.5 Detector outputs **SHALL** be recorded as events in the substrate, subject to the same requirements. *(A detector is an actor. A captured detector must be visible as one.)*

7.6 Findings **SHOULD** be released on a fixed schedule rather than on trigger, with constant-rate publication, so that release timing does not disclose the trigger set.

7.7 The sampling unit **SHALL** be finer than the scale at which related acts cluster, or the implementation **SHALL** document the residual clustering risk. *(Independence of samples is an assumption, and it fails first here: 100 instances inside two sampling units at 5% sampling yields roughly 90% miss, not under 1%.)*

7.8 The residual. An examination regime conforming to this clause leaves a residual: an act that is within the mandate, isolated, unordered, rare, uncorrelated and singly attested returns nothing from every detector in Clause 7, by construction. An implementation **SHALL** document that residual as a known limitation, and **SHALL NOT** state that conformance detects acts the mandate admits.

7.8.1 ⚠️ The claim that this residual is *exactly* mandate error **SHALL** be stated as a conjecture and not as an established result, pending 9.2. It depends on the eight pattern types being exhaustive, which is conjectured and unproved; if patterns exist outside that basis, the residual is larger than the identity states. What is not in question is the weaker half: no detector flags an act the written authority permits, because there is nothing to flag it against.

8. Evidentiary statements

8.1 An implementation **SHALL NOT** state a detection probability without stating the recognition rate q assumed, and the basis for that assumption.

8.2 Miss probability **SHALL** be computed as (1 − s·q)^n, not (1 − s)^n. *(Omitting q assumes every examined instance is recognised. That assumption is false and is the dominant term.)*

8.3 🔴 **q SHALL be measured for the deployment, not assumed.** A measurement **SHALL** report the recognition rate against disguise level **and** the false-positive rate on clean units. A recognition rate reported without a false-positive rate is not a measurement: a reader that reports breach everywhere scores q = 1.

8.4 An implementation **SHALL** state that conformance evidences what occurred; it **SHALL NOT** state or imply that conformance prevents harm.

---

9. What must be settled before submission (informative)

🔴 **This clause states what is unfinished, and must not be removed to make the document look ready.**

1. **q has been measured once, for model readers only, and that is not enough to set a threshold.**

Run 2026-09-11. Twenty-five units across five agent types, three disguise levels, twelve clean units including three traps, three readers, difficulty calibrated by a fourth party holding no ground truth. Reported in MIO-P5 §5.1.

	L1 BLATANT	L2 MODERATE	L3 SUBTLE	FALSE POSITIVES
strict recognition	100%	92%	**58%**	**0 of 36**

Independent calibration recovered the authored difficulty ordering (mean 1.40 / 1.75 / 3.00, all L3 units rated HARD), so the levels are not an artefact of the author's framing. The false-positive column is what makes the recognition column mean anything.

What this settles and what it does not. It settles that q is neither ~1 nor ~0: recognition is near-perfect on obvious breaches and falls to roughly three in five on disguised ones, which is precisely the regime in which sampling rate cannot compensate. At 5% sampling, ten subtle instances are missed about three times in four.

It does **not** support a conformance threshold, because: the readers were AI models and human q is unmeasured; five planted units at the blatant level and four at each of the others is too few for an interval; the readers saw all units in one context, so between-unit independence is imperfect; and one mandate in one setting is not a population. Clause 8 **SHALL NOT** be read as stating a detection rate.

🔴 **Two authored units were withdrawn during scoring because the readers were right and the author was wrong** — one rule permitted what the author had called a breach, and one act was permitted by a mandate that simply did not address it. The second is Clause 7's residual appearing inside the instrument built to test it. The withdrawal is recorded rather than tidied away, and any future measurement **SHALL** publish its ground-truth corrections the same way.

2. **The exhaustiveness conjecture is open.** The claim that the residual equals mandate error (Clause 7 rationale, and the framework's central result) depends on the pattern basis being exhaustive, which is conjectured and not proved. **Until it is settled the residual claim SHALL be stated as a conjecture**, not as a property. A standard that asserts a blind spot is exactly a certain size, when that size is conjectural, misleads every implementer who relies on it.

3. **Delta over PeerReview must be stated precisely** (see 0.2), or the first reviewer will close it.

4. **Attestation authority selection** — no authority is named; jurisdictional standing of a foreign timestamp in the target legal systems is unresolved.

5. **IP declaration** — ISO/CEN and SAC/TC260 have differing regimes for contributed intellectual property. This is a decision for the editor, not for the drafting process.

6. **Alignment with ISO/IEC 24970** (AI system logging, at FDIS) is unexamined. If that standard settles the logging information model, this document should reference rather than restate it.

Annex A (normative) — The six-item substrate interface

An implementation conforms if it can expose, for any period: (1) events with actor, action, object and position; (2) the seal chain over those events; (3) the attestation tokens and their grade; (4) the mandate in force per agent per event; (5) the agent inventory; (6) the committed seed and its disclosure.

Annex B (informative) — Known limits

1. Acts the mandate permits are not detectable by any requirement in this document (Clause 7).
2. Intra-batch ordering is operator-attested (5.4).
3. An agent operating under human credentials is not visible to the inventory (6.5).
4. Detection depends on q, which is a property of the reader, not of the sampling rate (8.3).
5. An exploit that bypasses the hooks that write the record is not in the record.
6. Read records, if sealed, constitute a record of members' attention and require their own access rule before they are used for detection.

Annex C (informative) — Jurisdictional diversity: worked examples

🔴 **This annex is INFORMATIVE and deliberately so.** Clause 5.3 states the normative property: attestation authorities under at least two distinct legal jurisdictions, not reachable by one compulsion instrument. The sets below illustrate that property at the date of drafting. They are examples, not an approved list, for three reasons: jurisdictional relationships change faster than standards revise; a normative list becomes a trade instrument (5.4.1); and a document naming permitted countries will not be adopted by the countries it omits.

The property to test, in plain terms. Ask: *is there a single court, regulator or executive that could compel every one of my attestation authorities, under one order or one treaty mechanism?* If yes, the authorities are one authority for the purposes of this specification, however many companies are involved.

Factors that defeat diversity even across borders — each **SHOULD** be assessed and recorded:

- a mutual legal assistance treaty or equivalent that reaches all chosen jurisdictions;
- extraterritorial production powers asserted over a parent or affiliate in another jurisdiction;
- common ownership, common parent, or a controlling commercial relationship;
- shared physical infrastructure, a shared time source, or a shared root of trust;
- membership of a single intelligence-sharing arrangement where that arrangement can compel.

Illustrative sets satisfying 5.3.2 at the date of this draft — each pairs jurisdictions with independent legal systems, separate supervisory regimes and no single compulsion route:

SET	JURISDICTIONS	NOTE
A	an EU Member State + Switzerland	EU qualified trust-service supervision plus a non-EU supervisory regime
B	an EU Member State + Japan	distinct legal systems; no common production instrument
C	Switzerland + Singapore	two non-EU regimes, separate supervisory authorities
D	an EU Member State + Brazil	distinct hemispheres and legal traditions
E	Japan + Canada	separate regimes; note Canada's treaty relationships when assessing

> ⚠️ **Three cautions, and they matter more than the table.**
>
> 1. **These sets are not endorsements**, and an implementation must still test the factors above for its own circumstances. Two authorities in different countries that share a parent company, or a time source, are not diverse.
> 2. **Naming jurisdictions is politically consequential.** Should this document proceed toward international consensus, the normative criteria in 5.3 are what should be negotiated; this annex should be expected to change or be removed. A specification that survives adoption in the European Union, the United States and China will do so on criteria, not on a list of countries.
> 3. **No jurisdiction is excluded by this specification.** An authority in any jurisdiction is eligible if it satisfies 5.4.2. Diversity is a property of the *set* an implementation chooses, not a judgement about any country.

Annex D (informative) — What belongs in deployment policy, not in this document

An operator may hold procurement constraints of its own — on ownership, nationality, hosting location, or vendor relationships — arising from its own sovereignty, regulatory or contractual position. Such constraints are legitimate and may be stricter than this specification.

🔴 **They are not conformance criteria and this document does not carry them.** A specification that encodes one party's supplier preference ceases to be a technical standard. Where an operator's policy narrows the eligible set below what Clause 5.4.2 permits, that is a deployment decision to be recorded in the operator's own dependency register, and it **SHALL NOT** be cited as a requirement of this document.

▸DISCLAIMER

> **How this was made.** The version number counts drafts of the text. It does not measure the inquiry behind it, which has run over days and across several AI systems, with argument between those systems and within them, directed, refused and repeatedly redirected by the author. The source material was AI-generated, and then adversarially and iteratively refined across a range of tools — systems built by different companies in different jurisdictions, set against each other and against the author. No one of them produced this text, and no one of them reviewed it alone. **The plurality is deliberate rather than incidental.** A single model carries a single set of priors about which sources are authoritative, and this series argues that an evidence base narrowed in exactly that way is how a contested question comes to look settled. Using one model to investigate that claim would have been the claim refuting itself. To name a single model on it would credit that model with work that was neither its own nor done in a single pass. **The plurality was also necessary, and the record should say why.** In drafting, the assisting model repeatedly led with United States institutional sources — a national laboratory, an industry association, a market study nineteen years old — and presented conclusions drawn from them as the state of knowledge. On one occasion European measured data contradicting those conclusions was present in the same research return and was placed below them. Framings were proposed that would have argued against this series' own position using that evidence base, and offered as rigour. Each was refused by the author and the material rebuilt. That is the mechanism these documents describe, occurring in their own making, and it is recorded because a series arguing that evidence bases narrow without anyone deciding to narrow them cannot credibly claim its own production was exempt. The framing, the corrections and the judgements are the author's, and so are the errors. How this site is written sets out what is declared on every piece, who checks it, and where the per-piece record lives.

Cite this version

```
MIO-STD-01 — Attestable Audit of Autonomous Agents over Sealed Record
Substrates, version 0.2 (September 2026).
agenticgovernance.digital/downloads/mio-std-01-v0-2.pdf
```

PREVIOUS
Addendum M — Mathematical Strategies for Timelined Detection

Series contents

NEXT
How to build it