The Modulign Correction Protocol: Formal Specification for Append-Only Error Resolution in DAG-OR

Vincent Gonzalez

Independent Scholar · ORCID 0009-0004-7283-8598

modulign.org · VinceGonzalez@me.com

Abstract

The Modulign Observation Registry is append-only: no entry is ever modified or deleted after commit. This invariant is foundational to the system's chain-of-custody argument, its legal applications, and its epistemic integrity. But classification errors occur. Instruments miscalibrate. Domain boundaries are misapplied. Observations are mischaracterized. An append-only system that cannot address error is not rigorous—it is brittle. This paper specifies the Modulign Correction Protocol (MCP): the formal mechanism by which erroneous classifications are identified, annotated, and superseded without violating the Permanence invariant (Invariant 2), the append-only constraint, or the provenance integrity requirement (Invariant 9). The protocol defines five correction types, mandatory fields for correction entries, authorization rules, discoverability requirements, dispute resolution procedures, and the specific obligations that arise when corrected classifications are involved in active legal proceedings.

Keywords: correction protocol, append-only, error resolution, classification amendment, chain of custody, evidentiary disclosure, Permanence invariant

1. Introduction: The Append-Only Problem

The Modulign Standard [1] commits the Observation Registry to an append-only architecture: entries are never modified or deleted after commit. The Formal Logic paper [2] proves this as Invariant 2 (Permanence). The Chain of Custody Formalisation [10] argues that this property is what makes the registry a legal chain of custody rather than merely a database. The synthesis paper [16] describes the correction principle—corrections are new entries referencing originals, never modifications of originals—but does not specify the protocol. This paper provides the specification.

The design principle is drawn from judicial practice: the trial record is never altered. When the record contains an error, a correction is entered alongside it. The original stands. The correction stands alongside it. Any reader of the record encounters both. This is the principle. The protocol makes it operational for a distributed, append-only classification registry.

2. Definitions

Original Entry: the GSI-ID being corrected. It is never modified. Correction Entry: a new GSI-ID that identifies, documents, and resolves the error in the Original Entry, committed through the standard CDP pathway and itself permanent and append-only. Correction Chain: the ordered

sequence of Correction Entries referencing the same Original Entry. An Original Entry may be corrected more than once; the most recent Correction Entry represents the current best classification. Supersession: the state of an Original Entry that has been corrected—marked as superseded, with the Correction Entry permanently linked; any retrieval must surface the supersession status. Dispute: a formal challenge to a Correction Entry, generating its own Correction Entry.

3. Correction Types

3.1 Type 1: Domain Reclassification

The Original Entry's §DOM assignment is incorrect. Required fields: Original GSI-ID reference, corrected §DOM value, basis for reclassification citing the specific domain boundary rule from [1] or [14] that was misapplied, classifier identity, timestamp. Implication: domain reclassification may cascade to §RLM, §SCL, and §REL values; the Correction Entry must specify all cascading changes.

3.2 Type 2: Segment Value Correction

One or more non-domain segments contain incorrect values. The domain assignment is correct but specific dimensional encodings are erroneous. Required fields: Original GSI-ID reference, identification of each corrected segment, corrected value with original stated, basis, classifier identity, timestamp. Segment value corrections refine the classification without changing its fundamental domain character.

3.3 Type 3: Entity Resolution Correction

The Original Entry's §EID assignment is incorrect—the observation was linked to the wrong persistent entity. Required fields: Original GSI-ID reference, incorrect EID, corrected EID (or new EID if the correct entity was not yet in the Entity Registry), basis, classifier identity, timestamp. Implication: entity resolution corrections affect all downstream relational links (RIDs) that depend on the EID; the Correction Entry must identify affected RIDs and specify whether they remain valid, require correction, or should be superseded.

3.4 Type 4: Provenance Correction

The Original Entry's §PROV chain contains an error—incorrect source attribution, a broken link, or a misidentified intermediary. This is the most legally sensitive correction type because it directly affects the chain-of-custody argument. Required fields: Original GSI-ID reference, identification of the provenance error and which link in the §PROV chain is incorrect, corrected provenance information, basis, classifier identity, timestamp, and—if the Original Entry has been used in any legal proceeding—the litigation-hold flag. Implication: a provenance correction may trigger the ^PROV exclusionary rule [4, 8]; the Correction Entry must explicitly state whether the Original Entry's evidentiary standing is affected.

3.5 Type 5: Retraction

The Original Entry should not have been committed: the input does not satisfy CDP Step 1, the classification was produced by a system operating outside its certified competence, or the

classification was produced fraudulently. Required fields: Original GSI-ID reference, retraction basis specifying which CDP requirement or %AUT constraint was violated, classifier identity, timestamp, authorizing party. The Original Entry is not deleted—it remains permanently—but receives a §FLAG value of RETRACTED. A retracted entry is a record that a classification was attempted, committed, and subsequently determined to be invalid. It does not count as a valid observation in registry statistics.

4. Mandatory Fields for All Correction Entries

Field	Content	Invariant Basis
GSI-ID	Unique identifier for the Correction Entry itself	Invariant 1 (Uniqueness)
§REL	Reference to Original Entry's GSI-ID, with relation type CORRECTS	Relation Registry linkage
Correction Type	DOMAIN_RECLASS, SEGMENT_CORRECTION, ENTITY_RESOLUTION, PROVENANCE_CORRECTION, or RETRACTION	Protocol compliance
Basis	Natural-language statement of what was wrong and why, with reference to the specific Standard provision	Auditability
Corrected Values	For Types 1–4: original and corrected values stated	Transparency
Classifier Identity	§OBS encoding for the person or system issuing the correction	Invariant 8 (Provenance)
§AUT (if applicable)	If issued by automated system, its certification level	%AUT compliance
Timestamp	§TMP encoding of the correction event	Invariant 7 (Temporal ordering)
§PROV	Provenance chain for the correction itself	Invariant 9
§SIG	Signature chain binding all fields	Invariant 10
§FLAG on Original	SUPERSEDED (Types 1–4) or RETRACTED (Type 5)	Discoverability

5. Authorization Rules

5.1 Types 1–3 (Domain, Segment, Entity)

Authorized correctors: the original classifier; any classifier with equal or higher competence (same or higher %AUT tier as the original); the node operator responsible for the node on which the entry was minted; in Phase 2+, the advisory board or its designated domain expert. Rationale: domain, segment, and entity corrections are reclassifications requiring competence at least equal to the original classifier. A :T system cannot correct a :E system's classification.

5.2 Type 4 (Provenance)

Authorized correctors: the node operator responsible for the entry's node; in Phase 2+, the advisory board's designated compliance officer; in federated architecture, the originating node's compliance authority or the cross-audit panel. Rationale: provenance corrections are custody corrections affecting the evidentiary chain. Authorization is restricted to the institutional level because provenance errors typically reflect process failures rather than individual classification errors.

5.3 Type 5 (Retraction)

Authorized retractors: the node operator with documented basis; in Phase 2+, the advisory board by majority vote; in federated architecture, the originating node's operator subject to cross-node review if challenged. Rationale: retraction is the strongest correction. An individual classifier—human or automated—cannot unilaterally retract. Retraction without proper basis would allow de facto deletion through the back door, violating the spirit of the append-only constraint even while technically preserving the letter.

6. Discoverability Requirements

6.1 Forward Linking

Every retrieval of an Original Entry that has been corrected must surface the correction. This is a system-level requirement, not a user-interface suggestion. Any API, query interface, or data export returning a GSI-ID must check the Correction Chain and return correction status alongside the original data. Implementation: the §FLAG segment is updated via append (a new flag-update record, not a modification) to include SUPERSEDED or RETRACTED. Any system reading GSI-IDs must check §FLAG. A SUPERSEDED or RETRACTED flag triggers mandatory retrieval of the linked Correction Entry.

6.2 Reverse Linking

Every Correction Entry must be retrievable by reference to the Original Entry it corrects. The §REL segment with relation type CORRECTS provides this. A query for “all corrections to GSI-ID X” must return the complete Correction Chain in temporal order.

6.3 Cascade Discoverability

When a correction affects downstream classifications—other GSI-IDs that referenced the Original Entry via RIDs, or used the Original Entry's §EID for entity resolution—the Correction Entry must identify the affected downstream entries and they must be flagged for review. This does not mean downstream entries are automatically corrected—they are flagged for evaluation. The evaluation may determine that the downstream classification remains valid despite the upstream correction, or it may trigger its own correction.

7. Litigation-Hold Protocol

7.1 Identification

A GSI-ID is subject to litigation hold when: the node operator has actual knowledge that the classification is the subject of pending or anticipated litigation; the node operator receives a litigation-hold notice from a party; or the classification has been offered as evidence in any proceeding (civil, criminal, administrative, or international).

7.2 Obligations

When a correction is issued for a GSI-ID under litigation hold: (1) Notification—all known parties to the proceeding must be notified within 48 hours of the correction's commitment, including the Original GSI-ID, the Correction Entry GSI-ID, the correction type, and a plain-language summary. (2) Preservation—the complete Correction Chain must be preserved in a format suitable for discovery production. (3) Defense access—in criminal proceedings, the defendant and defense counsel must have access to the full traceability record, deriving from the Due Process Clause and Brady v. Maryland (373 U.S. 83, 1963). (4) Court notification—if the correction affects a classification already admitted as evidence, the node operator must notify the court.

7.3 Government Node Obligations

When the node operator is a government entity, additional obligations apply. Brady compliance: if the correction is favorable to a criminal defendant—for example, a provenance correction revealing a gap in the chain of custody—it must be disclosed as Brady material regardless of whether the defense requested it. Correction-prompted review: a correction to a classification used in a criminal conviction triggers a mandatory review of whether the correction is material to the conviction. The review obligation rests with the prosecuting authority.

8. Dispute Resolution

8.1 What Constitutes a Dispute

A dispute arises when a participant challenges a Correction Entry—asserting that the correction is itself erroneous, the Original Entry was correct, or a different correction is warranted.

8.2 Dispute Procedure

Filing: the disputing party commits a Dispute Entry—a new GSI-ID with relation type DISPUTES referencing the Correction Entry—containing the basis for the dispute and the claimed correct classification. Both stand: during the dispute, both Original Entry and Correction Entry remain in the registry with their respective §FLAG values. Neither is modified. Assessment: in Phase 1, by the sole author; in Phase 2, by the advisory board; in Phase 3, by a cross-node dispute panel of domain experts from at least three nodes, none of which is the originating node for either entry. Resolution: a Resolution Entry either endorses the Correction Entry, endorses the Original Entry (effectively superseding the correction), or provides a third classification superseding both. Finality: there is no limit on dispute depth by design—the append-only architecture means no authority can impose a final answer that forecloses future correction. What changes with each level is the institutional weight of the resolution.

8.3 Dispute and Litigation

When a dispute concerns a classification involved in legal proceedings, the dispute itself becomes part of the litigation record. The litigation-hold obligations apply to all entries in the dispute chain. All parties to the proceeding must be notified of the dispute, its progress, and its resolution.

9. Interaction with Automated Classifiers

Automated classifiers at :E or :A levels are expected to detect potential errors in their own prior classifications when new information becomes available—for example, when a specification amendment changes a domain boundary rule. The %AUT framework's protocol traceability requirement provides the mechanism: the system can review its own traceability records against the amended specification and identify classifications that would be decided differently under the new rules.

An automated classifier may issue Type 1–3 corrections for its own prior classifications, subject to the competence threshold. An automated classifier may not issue Type 4 or Type 5 corrections—these require institutional authorization. A human classifier may correct any automated classification regardless of the automated system's %AUT tier: the human/automated asymmetry established in [17] applies here. A human operating under the CDP has the standing to overrule any automated classification because the human can be cross-examined about the basis for the correction and the automated system cannot.

10. Registry Statistics and Retracted Entries

Any empirical publication must report three distinct numbers: (1) Total GSI-IDs—all entries including corrections, disputes, resolutions, and retracted entries; (2) Valid classifications—GSI-IDs that are neither RETRACTED nor SUPERSEDED; this is the count that represents the empirical corpus; (3) Correction rate—the ratio of Correction Entries to total Original Entries by type, as a system health metric. Reporting only the total without the valid count would overstate the corpus. Reporting only the valid count without the total would obscure the correction history.

11. Open Problems

Cascade evaluation at scale: when a correction affects a high-connectivity GSI-ID—referenced by hundreds or thousands of downstream entries—the cascade evaluation requirement may impose significant review burden. An automated cascade evaluation mechanism is a priority engineering problem. Dispute depth: the protocol imposes no limit on dispute depth; pathological dispute chains on politically or legally contentious classifications are possible. The Phase 3 governance body should have authority to declare a dispute chain resolved by supermajority vote—a resolution that can itself be appended to, not sealed. Cross-node correction authority: in a federated architecture, the social and technical protocols for one institution's experts overruling another's require specification in the Node Participation Agreement.

12. Conclusion

The Correction Protocol is the mechanism by which an append-only system remains honest. Without it, the registry accumulates error alongside truth with no formal means to distinguish

them. With it, errors are identified, documented, linked to their originals, and made permanently discoverable—while the originals stand as a permanent record of what was once believed. History happened. The correction happened too. Both are on the record.

References

[1] Gonzalez, V. (2026). *The Modulign Standard v3.0*. Zenodo. https://doi.org/10.5281/zenodo.19348703

[2] Gonzalez, V. (2026). *Modulign Formal Logic v1.1.0*. Zenodo. https://doi.org/10.5281/zenodo.19350847

[4] Gonzalez, V. (2026). *Legal Epistemology and the Modulign Framework v1.0.1*. Zenodo. https://doi.org/10.5281/zenodo.19351250

[7] Gonzalez, V. (2026). *%AUT Certification Framework (Modulign v3.1 Appendix)*. Zenodo. https://doi.org/10.5281/zenodo.19559602

[8] Gonzalez, V. (2026). *Privacy Commitment & Governance Roadmap v1.0.0*. Zenodo. https://doi.org/10.5281/zenodo.19559689

[10] Gonzalez, V. (2026). *Chain of Custody as Embedded Property: Formalising ^EVID in DAG-OR*. Zenodo. https://doi.org/10.5281/zenodo.19642385

[14] Gonzalez, V. (2026). *Modulign v3.1 Amendment: ENV and BIO Domains*. Zenodo. https://doi.org/10.5281/zenodo.19644043

[16] Gonzalez, V. (2026). *Modulign: Architecture, Implementation, and Applications of a Dimensional Address Grammar for Observable Reality*. Zenodo. https://doi.org/10.5281/zenodo.19726994

[17] Gonzalez, V. (2026). *The Epistemology of Observation: A Formal Certification Framework for Human and Automated Classifiers in DAG-OR*. Zenodo. https://doi.org/10.5281/zenodo.19726226

[18] Gonzalez, V. (2026). *Federated Classification Infrastructure: Registry Governance, Node Architecture, and Decentralized Stewardship for DAG-OR*. Zenodo. https://doi.org/10.5281/zenodo.19726393

[20] Gonzalez, V. (2026). *The Human Attestation Protocol: Confrontation Clause Compliance for Automated Classifications in DAG-OR*. Zenodo. https://doi.org/10.5281/zenodo.19726407