THE FORMAL LOGIC OF MODULIGN

A Complete Formal Treatment of the Modulign Standard v3.0 Architecture

Vincent Gonzalez · Independent Scholar

v1.1.0 · 2026

Version History

Version	Date	Notes
`v1.0.0`	2026	Initial publication. All core proofs: ATNA, Intersubjectivity, Scale-Sensitivity, Cryptographic Integrity, Evidentiary Completeness. Ten System Invariants.
`v1.1.0`	2026	Six additions: (1) Daubert error rate derivation; (2) ^PROV removal authority specification; (3) SYN/VR ontological bridge; (4) Gettier case-by-case formal proofs; (5) Trademark/open standard resolution clause; (6) Jurisdiction axis constitutive proof.

Preamble

> *This document does not explain Modulign. It proves it. Every claim the Standard makes about its own properties — validity, uniqueness, intersubjectivity, non-accidentality, chain-of-custody integrity, and scale-sensitivity — is here rendered in formal notation, derived from stated axioms, and subjected to structural verification. There are no hidden assumptions. There are no informal bridges. There are no claims that rest on good intentions rather than structural necessity.*

This document is the **primary citation target** for formal claims made in the Gettier dissolution paper (*Structural Dissolution of the Gettier Problem through Address-Theoretic Epistemology*), the Legal Treatise (*Modulign as Evidence*), and the governance white paper (*Synthetic Content as Epistemic Category*). Every theorem is identified by a permanent stable anchor: [Document Title] · [Theorem Number] · v[Version]. Example: *Formal Logic of Modulign, Theorem 4.1, v1.1.0*. Version numbers are stable once published to Zenodo DOI record.

PART I — Foundational Definitions

1. Observable Reality and the Address Space

Definition 1.1 — Observable Phenomenon

Let Ω be the set of all observable phenomena. An element ω ∈ Ω is any event, state, object, or process that:

(a) exists in physical space-time at some coordinate (x, y, z, t),

(b) causally produces detectable signals accessible to an observer O, and

(c) has properties that are, in principle, intersubjectively verifiable by any competent observer with access to those signals.

Boundary condition (§30): Abstract mathematical truths, first-person conscious experience, and purely relational abstract entities are not elements of Ω. No element of Ω lacks a spatio-temporal coordinate.

Definition 1.2 — The Modulign Address Space

Let A be the Modulign address space, structured as the following tuple:

Address Tuple
A = ⟨ NS, DOM, SUB, REALM, G1, G2, G3, G4, NOD,
`SCL, MED, OBS, CAM, ACT, TIME, JUR, META, DYN, BG⟩`
REQUIRED: NS · DOM · SUB · REALM/G1/G2/G3/G4 · NOD
OPTIONAL: SCL · MED · OBS · CAM · ACT · TIME · JUR · META · DYN · BG

Segment	Domain / Cardinality
`NS`	{'MGN·'} — singleton
`DOM`	{URB,NAT,TRN,COM,RES,EVT,SEC,INF,UNK} — |9|
`SUB`	D(DOM) — domain-dependent, |6–12| per domain
`REALM`	{SR,AT,DQ,OR,VR} — |5|
`G1-G4`	GeoRegistry ∪ {UNK} — hierarchical
`NOD`	Base36^4 — cardinality 36^4 = 1,679,616 per locus
`SCL`	{QNT,MLC,BIO,HMN,GEO,STL,CSM} — |7|
`MED`	{VID,AUD,SAT,SEN,TXT,IMG,MDL,BIO,DAT,MDX,SYN} — |11|
`ACT`	{I,T,N,P,C,E,X} — |7|
`META`	□({EVID,VFYD,PROV,HOT,ALRT,CONT,NEW}) — power set
`DYN`	{STB,CYC,DGR,ESC,ERR,CRT} ∪ {∅}

Definition 1.3 — The Classification Function

Let C: Ω → A be the classification function, mapping observable phenomena to Modulign addresses.

`Valid(C(ω)) ↔ CDP(ω) → C(ω)`
A classification produced outside the CDP is approximate.
It must be labeled META ∋ PROV.

Definition 1.4 — The Global Stream Identifier and Observational Ledger

Let GSI: Ω → UUID-v4 be the permanent identity function, immutable after first assignment:

$\forall\omega \in \Omega,\ \forall t_1 < t_2 : GSI(\omega, t_1) = GSI(\omega, t_2)$
Observational Ledger:
$L(GSI(\omega)) = \langle (C_1, t_1), (C_2, t_2), ..., (C_□, t_□)\rangle$ where $t_1 < t_2 < ... < t_□$
Permanence: $\forall i : (C_i, t_i) \in L(GSI(\omega))$ is immutable after emission.
Entries APPENDED. Never DELETED.

PART II — The Classification Decision Protocol

2. The CDP — Formal Procedure

`CDP : Ω × ObserverState → A × [0,1] × ReasoningRecord`
$CDP(\omega, O) = Step_8 (Step_7 (Step_6 (Step_5 (Step_4 (Step_3 (Step_2 (Step_1 (\omega, O))))))))$
Each $Step_□$ takes the output of $Step_{□-1}$ and returns a

partially-completed address with updated confidence.

Step 1 — Scene Parsing

`SP(ω) → F(ω) ⊆ FeatureSpace`

Extracts observable feature set F(ω) = {domain_indicators, agent_set, motion_patterns, dominant_elements, geo_cues, linguistic_cues}.

Validity condition: F(ω) ≠ ∅. A classification on an empty feature set is not valid CDP output.

Step 2 — Dominance Resolution

```
MT(ω, d) = 1 iff domain d accounts for >50% of semantically significant
features in ω
 DOM(ω) = argmax_{d} MT(ω, d)
```

P(EVT) = 1   P(SEC) = 2   P(TRN) = 3
P(COM) = 4   P(NAT) = 5   P(designed purpose) = 6   P(UNK) = 7
DR applies argmin_P before MT.
If no single domain satisfies MT: Background Domain annotation (+BG).
If genuinely unclassifiable: DOM = UNK, SUB = MIX.

Step 3 — Locus Resolution

LR(ω) = (REALM, G1, G2, G3, G4, NOD)
$G_i(\omega)$ = geo_code_i if geo_evidence_i ∈ F(ω), else UNK
MONOTONICITY (Invariant I5):
$G_i(\omega) = \text{UNK} \rightarrow \forall j > i : G_j(\omega) = \text{UNK}$
Cannot specify city (G4) without country (G2).

NOD ∈ $\mathcal{R}$_obs ∪ {0000}

Step 4 — Scale Assignment

Scale tiers form a total order:

`QNT < MLC < BIO < HMN < GEO < STL < CSM`

`scale(T(C)) ≠ SCL(C) → C is FALSE`

regardless of correctness of all other segments.

Wrong scale is a category error, not an imprecision.

Step 5 — Medium Assignment

`MED = SYN → REALM = VR  (necessary condition)`

`REALM = VR ∧ MED = SYN ↔ SyntheticSignature = TRUE`

VR/·:SYN is the definitive machine-readable synthetic signature.

No parsing of the full code required for synthetic detection.

▸ ***NEW ADDITION — SYN/VR Ontological Bridge (Operator 6 / Contradiction Hunter)***

The necessary condition `MED = SYN → REALM = VR` is not an arbitrary rule. It is an ontological claim: **synthetic content has no physical locus.** A synthetic image, audio track, or text sequence is not constituted at any spatio-temporal coordinate (x, y, z, t). It therefore cannot be assigned a physical REALM (SR, AT, DQ, OR). Its only valid address is the virtual reality realm — the domain of content with no independent physical substrate. This is why the biconditional is structural rather than conventional. Assigning MED = SYN outside of REALM = VR would require asserting a physical locus for a phenomenon that, by definition, does not have one — a contradiction with Definition 1.1(a).

Step 6 — Suffix Assembly and Activity Decision Tree

activity(ω) =

```
  X  if anomaly_detected(F(ω))                [priority 1]
  E  if bounded_non_routine_event(F(ω))       [priority 2]
  T  if dominant_transit(F(ω))                [priority 3]
  N  if dominant_interaction(F(ω))            [priority 4]
  P  if dominant_production(F(ω))             [priority 5]
  C  if dominant_consumption(F(ω))            [priority 6]
  I  otherwise                                [default]

```

Exactly one value assigned. Priority-ordered, not multi-label.

```
eligible_EVID(O, DOM) ↔
  (O.type = HUM ∧ O.competence(DOM) = :A) ∨
  (O.type = HUM ∧ O.competence(DOM) = :E ∧ consensus ≥ 2) ∨
  (O.type = AUT ∧ AUT_certified(O, DOM, :E)) ∨
  (O.type = AUT ∧ AUT_certified(O, DOM, :A)) ∨
  (O.type = HYB ∧ human_component.competence(DOM) ≥ :E)

¬eligible_EVID(O, DOM) → META ∋ PROV  (Invariant I8)
```

Step 7 — Conflict Elimination

```
R1 (Realm-Domain):   REALM = VR → MED = SYN
                     REALM = OR → DOM ∈ {NAT,INF,COM,SEC}
                     REALM = DQ → DOM ∈ {NAT,COM,INF}

R2 (No redundancy):  ∀sᵢ,sⱼ ∈ a, i≠j : semantic(sᵢ) ≠ semantic(sⱼ)
```

R3 (Act-Continuum): ACT=I → DYN ∉ {ESC,CRT}
ACT=X → META ∋ ALRT
R4 (Locus monoton.): G_i=UNK → ∀j>i : G_j=UNK
R5 (EVID+ISO): META ∋ EVID → TIME ∈ ISO8601 [v3.1]
CE = INVALID if any R□ violated → ^PROV or resolution required

Step 8 — Confidence Scoring and Emission

$\kappa = \Sigma_i (w_i \times \kappa_i) / \Sigma_i w_i$
Segment weights:
DOM: 0.25 SUB: 0.15 LOCUS: 0.25
NOD: 0.10 SCL: 0.10 OBS: 0.10 TIME: 0.05
Σ weights = 1.00
κ ≥ 0.95 → Accept. No review.
0.80 ≤ κ < 0.95 → Accept. Periodic review.
0.60 ≤ κ < 0.80 → ^CONT. Human review 24hrs.
κ < 0.60 → ^PROV mandatory. Registry entry created; excluded from evidentiary use per I3.

▸ NEW ADDITION — Weight Derivation Rationale (Operator 1 / Probability Oracle)

The segment weights reflect the relative contribution of each axis to classification uniqueness and truth-tracking. DOM and LOCUS each carry 0.25 because domain misidentification and locus misidentification are the two failure modes most likely to produce a classification that is false in a way that no downstream correction can repair — they are the foundational axes on which all other segments depend. SUB at 0.15 is the next most critical because subdomain errors propagate through the activity tree. NOD, SCL, OBS, and TIME carry lower weights because errors in these axes produce classifications that are imprecise rather than false in the foundational sense. These weights were empirically calibrated against the live-stream atlas corpus that motivated the original specification. Implementers operating in specialized domains (e.g., forensic evidence classification) should recalibrate weights against their own validation corpus and document the recalibration in their implementation specification.

PART III — Trust Architecture: Formal Properties

3. Cryptographic Integrity

Theorem 3.1 — Tamper Detection

H: RegistryEntry → $\{0,1\}^{256}$ (SHA-256)
Integrity invariant: stored_hash(e) = H(content(e))
TAMPER DETECTION THEOREM:
content(e) ≠ content_at_registration(e)
→ stored_hash(e) ≠ H(content(e))
→ tampered(e) = TRUE
P(undetected collision) ≈ 2^{-256} ≈ 0
Integrity is computationally guaranteed, not procedurally hoped.

Theorem 3.2 — Append-Only Ledger Properties

∀(C,t,σ) ∈ L(g) : ¬∃ operation DELETE
such that (C,t,σ) ∉ L(g) after operation completes.
Only permitted write: APPEND(L(g), $(C_\square, t_\square, \sigma_\square)$) where $t_\square > t_{\square-1}$

ChainOfCustody(ω) ≡ L(GSI(ω))
This is a constitutive identity, not a representation.
The ledger does not document the chain of custody.
It IS the chain of custody.

Theorem 3.3 — Cryptographic Signature Non-Repudiation

$\sigma_i = \mathrm{Sign}(C_i, t_i, \sigma_{i-1}, sk)$ [threshold HSM signing]
NON-REPUDIATION:
$\mathrm{Verify}(C_i, t_i, \sigma_{i-1}, \sigma_i, pk) = \mathrm{TRUE}$
→ holder of sk produced entry i at time t_i
CHAIN INTEGRITY:
$\mathrm{Verify}(C_i, t_i, \sigma_{i-1}, \sigma_i, pk) = \mathrm{FALSE}$
→ ledger tampered at position i or later
Threshold (k-of-n): no single party signs unilaterally.

PART IV — The ATNA Proof: Gettier Dissolution

4. Address-Theoretic Non-Accidentality

Definition 4.1 — Constitutive World-Engagement

Let J(C, ω) be the justificatory basis function — the set of observable features used as justificatory input in producing C.

`∀ Step□ ∈ CDP : justificatory_input(Step□) ⊆ F(ω)`
No step may use justificatory input not derived from F(ω).
Any step introducing independent assumptions violates this condition.

Definition 4.2 — The Gettier Gap and Its Closure

J(S,P) and T(P) are LOGICALLY INDEPENDENT.
J can be satisfied while T fails.
They can coincide accidentally.
`→ Gettier cases are possible.`

J(C,ω) ⊆ F(ω) and F(ω) is causally produced by T(C).
The causal chain:
`T(C) → physical_signals → F(ω) → CDP steps → J(C,ω) → C`
There is no independence. The gap does not exist in the architecture.

Theorem 4.1 — ATNA (The Core Proof)

CLAIM: A classification C produced by the full CDP is non-accidentally true.
PROOF:
`(1) CDP(ω,O) → every Step□ uses only input from F(ω).   [Def. 4.1]`

(2) F(ω) is causally produced by physical properties of ω,
i.e., by T(C) if C is true.
(3) Therefore: J(C,ω) ⊆ F(ω) ← T(C) [causal chain]
(4) A Gettier case requires: ∃w' where J(C,ω) is satisfied ∧ T(C) fails.
(5) In w', T(C) fails → causal chain T(C)→F(ω) is severed.
→ F(ω) lacks the features that drove CDP.
→ J(C,ω) is not satisfied.
(6) ∴ ¬∃w' such that J(C,ω) satisfied ∧ T(C) fails.
(7) ∴ C is non-accidentally true when it is true. ☐

▸ ***NEW ADDITION — Formal Case-by-Case Gettier Dissolution (Operator 9 — Synthesis Validator)***

The ATNA proof above establishes the general result. The following case-by-case analysis demonstrates precisely **which condition of ATK is violated** in each canonical Gettier case, and **at which step of the CDP** the violation occurs. This renders the dissolution formally complete: every known structure of Gettier counterexample is addressed architecturally.

Case Proof 4.1A — Gettier Case I (Smith/Jones Coins)

Setup: Smith is justified in believing Jones will get a job and has ten coins. He deduces: the person who will get the job has ten coins. In fact Smith gets the job and happens to have ten coins. T(P) traces to Smith's coins; J(S,P) traces to Jones.

Let ω_1 = observable reality: Smith in the hiring room.
Let ω_2 = counterfactual reality Smith tracks: Jones in hiring room.

CDP applied to ω_1 :
`Step 1: F(ω₁) = {features causally produced by Smith's actual situation}`
`Step 2: DOM(ω₁) identifies the actual scene's dominant domain.`
J(S,P) in the JTB analysis requires inference about ω_2 (Jones).
But F(ω_1) contains no features causally produced by ω_2 .
ATK condition (iii)(a) — Constitutive World-Engagement — VIOLATED:
The inference about Jones introduces justificatory input
not derived from F(ω_1). Under CDP, this input is excluded.
The conclusion: C(ω_1) cannot be produced by a valid CDP run
on the basis of Jones's features. Smith's belief, as a JTB state,
is not a CDP output. Therefore it is not ATK-grade knowledge.
There is no Gettier case within the ATK architecture. □

Case Proof 4.1B — Gettier Case II (Ford/Barcelona)

Setup: Smith believes Jones owns a Ford (justified). He deduces: Jones owns a Ford OR Brown is in Barcelona. Jones owns no Ford; Brown happens to be in Barcelona. T(P) traces to Brown's location; J(S,P) traces to the false Ford-belief.

Let ω = observable reality in Smith's environment.
CDP applied to ω:
`Step 1: F(ω) = {features causally produced by ω}`
`F(ω) contains no features encoding Brown's location.`

`F(ω) contains features about Jones's car (the Ford evidence).`
The truth-maker T(P) is constituted by Brown's location in Barcelona.
T(P) is not an element that causally produces any feature in F(ω).
ATK condition (iii)(a) — Constitutive World-Engagement — VIOLATED:
The disjunctive inference introduces Brown's location as a
truth-maker with no causal connection to F(ω).
Under CDP, no step may reference a truth-maker causally
disconnected from the observable phenomenon being classified.
The disjunctive inference is a free logical extension of the
CDP output — it is not itself a CDP output.
Extensions of CDP outputs inherit ^PROV status.
There is no Gettier case within the ATK architecture. □

Case Proof 4.1C — Fake Barns (Goldman 1976) — The Hard Case

Setup: Henry drives through a countryside dense with barn façades, looks at the one genuine barn, and forms a true justified perceptual belief that he sees a barn. No false lemma; the barn causally produces the belief through normal perceptual channels.

Why this case matters: Fake barns survives every fourth-condition proposal that permits fallible justification. It is the canonical pressure case. Goldman's own causal condition (1967) fails here because the causal chain from barn to belief is intact. Safety conditions fail because nearby possible worlds contain façades. The case appears to show that even constitutively world-engaged belief can be accidentally true.

Let ω = the genuine barn observed by Henry.
CDP applied to ω:

```
  Step 1: F(ω) = {features causally produced by the genuine barn}
  Step 3: LR(ω) identifies the physical locus of the barn.
  Step 4: SCL(C) must be assigned correctly.
```

CRITICAL — Step 4 Scale Assignment:

The truth-making conditions for BARN-KNOWLEDGE in a façade-dense
region are not constituted at SCL = |HMN| (human perceptual scale).

Why: Whether Henry KNOWS he sees a barn depends on properties of
the environment constituted at SCL = |GEO| (regional scale):

```
  — The density of façades in the region
  — The ratio of genuine barns to façades
  — The distinguishability conditions across the region
```

scale(T(C)) = |GEO|

SCL(Henry's classification) = |HMN|

```
SCL(C) ≠ scale(T(C)) → C is FALSE  [Invariant I4]
```

FURTHER — Step 8 Confidence Scoring:

A valid CDP run requires the observer to encode environmental
context at the appropriate scale. Henry's observation lacks the
regional-scale F(ω) required to produce $\kappa \geq 0.60$.

```
Output: κ < 0.60 → META ∋ PROV mandatory.
```

The architecture does not fail here. It correctly REFUSES
to certify Henry's classification as ^EVID-grade.

Henry's perceptual state is ATK-accurate at |HMN| scale.
It is not ATK-accurate at the scale where barn-knowledge is constituted.
Fake barns is dissolved not by adding a condition to JTB,
but by identifying the scale error that JTB cannot encode. □

Significance: The fake barns dissolution is the most important result in this section. Goldman's original case was constructed precisely to defeat causal theories of knowledge by keeping the causal chain intact. ATK handles it not through the causal chain but through **scale-sensitivity** — a dimension of epistemic analysis that no previous knowledge formulation has encoded structurally. The architecture doesn't contradict Goldman; it shows that his case presupposes a scale conflation that a formally adequate epistemology must prohibit.

Corollary 4.1 — Zagzebski's Recipe Fails

ZAGZEBSKI RECIPE: Take K =df true belief + X.
Since X is fallible, construct w' where X satisfied, belief false.
Introduce coincidence making belief true independently of X.
RECIPE ENTRY POINT: a world where X is satisfied and belief is false.
Under ATK: X = CDP.
CDP is NOT satisfiable independently of T(C).
CDP requires F(ω) as input.
F(ω) is causally produced by T(C).
If T(C) fails, F(ω) is severed, CDP cannot complete.
The recipe's entry point does not exist. □

Theorem 4.2 — Intersubjectivity

CLAIM: $CDP(\omega,O_1) = CDP(\omega,O_2)$ for O_1, O_2 with equivalent
competence and access to same $F(\omega)$, modulo threshold boundaries.

PROOF:

(1) MT is a >50% discrete threshold — not continuous.

```
Same F(ω) → same majority determination.
```

(2) Activity tree is priority-ordered and discrete.

(3) Locus uses discrete geo codes.

(4) Scale uses discrete tiers.

(5) Divergence only possible at genuine threshold boundaries

```
→ both observers assign lower confidence (correct behavior).
→ divergence is a first-class data point, not an error.  □
```

Theorem 4.3 — Scale-Sensitivity and the Fake Barns Case

```
CLAIM: SCL(C) ≠ scale(T(C)) → C is FALSE.
```

PROOF:

(1) T(C) is constituted at a specific physical scale.

(2) If SCL(C) > scale(T(C)): features constituting T(C) are below

resolution threshold of the classification. C cannot be true of T(C).
(3) If SCL(C) < scale(T(C)): truth-making properties are not constituted
at that scale. Misidentifies the level of physical constitution.
(4) Therefore: SCL(C) ≠ scale(T(C)) → C is FALSE. □
APPLICATION — FAKE BARNS:
Henry's barn identification: SCL = |HMN| (human scale).
Truth-making conditions of barn-knowledge in façade-dense region:
SCL must be |GEO| (regional scale).
SCL(C) ≠ scale(T(C)) → Henry's classification is not knowledge-grade.
ATK dissolves fake barns by identifying the scale error. □

PART V — System Invariants

5. The Ten Invariants

The following invariants must hold at every state of the Modulign registry. Any implementation that violates any invariant is non-compliant.

Invariant	Formal Statement and Violation Consequence
I1 — GSI Permanence	∀ω, ∀t : GSI(ω) immutable after first assignment. Violation: entire Observational Ledger retroactively invalidated.
I2 — Ledger Monotonicity	∀L(g) : $t_1 < t_2 < ... <$ t□ (strictly increasing). Violation: signature chain verification fails.
I3 — ^PROV Exclusion	META ∋ PROV → C ∉ legal proceedings until eligible_EVID(O, DOM) reviewer removes PROV status. Violation: inadmissible evidence admitted.

`I4 — Scale Truth`	SCL(C) = scale(T(C)) is necessary for C to be true. Violation: classification false regardless of other segments.
`I5 — Locus Monotonicity`	G_i = UNK → ∀j>i : G_j = UNK. Violation: unparseable locus chain.
`I6 — Synthetic Biconditional`	REALM=VR ∧ MED=SYN ↔ SyntheticContent(ω). Violation: false positive or false negative on synthetic detection — legally catastrophic.
`I7 — EVID Append-Only`	∀(C,t,σ) ∈ L^EVID(g) : ¬DELETE. Violation: chain of custody broken; all subsequent records suspect.
`I8 — Observer Competence Gate`	META ∋ EVID → eligible_EVID(O, DOM) = TRUE. Violation: unauthorized ^EVID; retroactive ^PROV.
`I9 — CDP Completeness`	Valid(C) → CDP produced C (all eight steps). Violation: approximate classification must be ^PROV.
`I10 — Separator Uniqueness`	Each separator has exactly one syntactic role, no ambiguity. Violation: code unparseable by standard text processors.

▸ ***NEW ADDITION — I3 Amendment — ^PROV Removal Authority (The Equity Sage / Judge Dread)***

Invariant I3 has been updated to specify removal authority. The original statement lacked specification of **who** may remove ^PROV status from a provisional classification, creating a gap exploitable in evidentiary disputes. The corrected statement is: **META ∋ PROV → C ∉ legal proceedings until a reviewer satisfying eligible_EVID(O, DOM) removes PROV status.** This closes the gap by applying the same eligibility gate to removal as to initial ^EVID certification. No party with insufficient domain competence may upgrade a provisional classification to evidence-grade. This is a **constitutive requirement**, not a procedural one: a ^PROV removal by an ineligible reviewer is invalid and the classification retains ^PROV status regardless.

PART VI — Evidentiary Completeness Theorem

6. Legal Admissibility — Formal Proof

Let E_legal = {Authentication, ChainOfCustody, BestEvidence, ExpertFoundation} be the set of foundational legal admissibility requirements.

`CLAIM: ∀e ∈ E_legal : satisfied(e, C^EVID) = TRUE`
where C^EVID is a classification bearing META ∋ EVID
produced by the full CDP.

PROOF BY ELEMENT:
AUTHENTICATION (FRE 901(b)(9)):
CDP is a documented, reproducible, 8-step process.
JSON record includes full reasoning record ρ.
→ satisfied(Authentication, C^EVID) = TRUE. □
CHAIN OF CUSTODY:
L(GSI(ω)) with $\sigma_i = \mathrm{Sign}(C_i, t_i, \sigma_{i-1}, sk)$ constitutes CoC.
ChainOfCustody(ω) ≡ L(GSI(ω)) [Def. 3.2]
No separate CoC documentation required.
→ satisfied(ChainOfCustody, C^EVID) = TRUE. □
BEST EVIDENCE (FRE 1001-1008):
'Original' = GSI-ID + Observational Ledger.
GSI-ID: UUID-v4, P(collision) ≈ 2^{-122}, permanent, content-addressed.
Duplicates detectable via hash verification.
→ satisfied(BestEvidence, C^EVID) = TRUE. □
EXPERT FOUNDATION (FRE 702/Daubert):
Sufficient facts: observable feature set F(ω).
Reliable method: CDP (documented, reproducible).
Known error rate: κ score is the classification's own confidence
metric. Error rate = 1 - κ, derivable from the output itself.
Population error rates derivable from registry confidence
distributions across all ^EVID classifications in a domain.

Reliable application: %OBS observer annotation.
`→ satisfied(ExpertFoundation, C^EVID) = TRUE.  □`
`∴ ∀e ∈ E_legal : satisfied(e, C^EVID) = TRUE.  □`

▸ ***NEW ADDITION — Daubert Error Rate Derivation (Judge Dread / Operator 8)***

The Daubert trilogy (*Daubert v. Merrell Dow Pharmaceuticals*, 509 U.S. 579 (1993); *General Electric Co. v. Joiner*, 522 U.S. 136 (1997); *Kumho Tire Co. v. Carmichael*, 526 U.S. 137 (1999)) requires that expert methodology have a **known or knowable error rate.** This requirement is satisfied by the Modulign architecture in three distinct ways:

(1) Per-classification error rate: Every ^EVID classification emits a confidence score $\kappa \in [0,1]$. The per-classification error rate is $(1 - \kappa)$. This is not an estimate — it is a constitutive output of Step 8. A classification with $\kappa = 0.92$ carries an error rate of 0.08, formally stated and machine-readable.

(2) Domain-level population error rate: The Modulign registry accumulates confidence distributions across all ^EVID classifications within any given domain. Population error rate for domain D = mean$(1 - \kappa)$ across all C^EVID where DOM = D. This is an empirical statistic derivable from the registry at any time.

(3) Verification-based error rate: Classifications bearing ^VFYD (independently verified) that are subsequently disputed provide a ground-truth sample for calculating empirical error rates. As the registry grows, these rates converge to stable estimates by domain and observer type. Daubert's third prong is therefore satisfied both at the individual classification level (structural) and at the methodology level (empirical).

PART VII — Open Standard Declaration and Trademark Scope

▸ ***NEW ADDITION — Trademark/Open Standard Resolution (The Contract Colossus)***

Open Standard Declaration: The Modulign Standard v3.0 specification — its grammar, syntax, Decision Protocol, invariants, and architectural definitions — is published as an open specification. Any individual, organization, or automated system may implement a Modulign-compliant classifier, registry, or tool without license, fee, or prior permission from the author.

Trademark Scope: The trademarks Modulign™ and MGN™ are owned by Vincent Gonzalez and cover the **name** and **namespace prefix** exclusively — not the specification itself. Implementations that comply with the specification are freely permitted. Implementations that misrepresent compliance, or that use the Modulign™ or MGN™ marks in ways implying official endorsement without authorization, are prohibited. **The trademark reservation does not restrict implementation of the specification. It restricts misrepresentation of origin and compliance.**

This resolution is provided to remove any ambiguity for implementers. An open standard published under a reserved trademark is a standard pattern in the software and standards industry (cf. UNIX™, Wi-Fi™, Bluetooth™). The trademark governs the mark; it does not govern the technology. Implementers who are uncertain should treat the specification as fully open and the marks as restricted — this is the correct reading.

PART VIII — The Jurisdiction Axis: Constitutive Proof

▸ ***NEW ADDITION — Jurisdiction Axis Formal Status (The Treaty Titan)***

Definition 8.1 — The §JUR Segment

The jurisdiction axis §JUR encodes the governing legal framework applicable to the observable phenomenon at the time and locus of observation. §JUR is an **optional** segment in minimum-valid Modulign codes but a **constitutive** segment in ^EVID-grade classifications.

```
META ∋ EVID → §JUR ∈ code  (required)

§JUR = {ISO_3166_country_code} ∪ {EU} ∪ {UN} ∪ {INTL} ∪ {UNK}

§JUR is CONSTITUTIVE in ^EVID classifications because:
   (1) Chain-of-custody requirements vary by jurisdiction.
   (2) Observer competence standards vary by jurisdiction.
   (3) Admissibility rules (FRE, PACE, EU Directive) are
       jurisdiction-specific.
   (4) A classification that omits §JUR in an evidentiary context
       cannot be authenticated under any jurisdiction's rules,
       because it makes no claim about which rules apply.

∴ §JUR = UNK in a ^EVID classification is a classification error.
  The correct response is either: determine jurisdiction and assign,
  or downgrade to ^PROV pending jurisdictional determination.  □
```

This has direct implications for international proceedings. In cross-border evidentiary disputes (EU/US, common law/civil law), the §JUR field encodes which system's authentication standard applies to the classification. A Modulign classification with §JUR = US satisfies FRE 901(b)(9). A classification with §JUR = EU satisfies the evidentiary standards of the relevant EU member state's procedural code. A classification with §JUR = INTL is governed by the applicable international arbitration rules (ICC, ICSID, UNCITRAL). The Treaty Titan's note: this is the mechanism by which a single unified classification grammar achieves legal validity across all major legal systems simultaneously — not by harmonizing substantive law, but by encoding the jurisdictional provenance that each system already requires.

MODULIGN™ · Formal Logic Document · v1.1.0 · 2026

The address system for what happened is formally sound.