



ENCEVRA RESEARCH / ROUND 34

Can Trust Bootstrap Itself?

Circular Attestation, Sybil Issuers, and the Minimum Grounded Trust Base

READINESS DECISION: BLOCKED

October 25, 2026

An exact 279,936-graph diagnostic shows that cross-certification cannot manufacture a trust root. Safe bounded validation abstains at unresolved frontiers; fail-open frontier promotion creates false assurance.

Contents

Abstract	3
1. Archive continuity and corrected round number	4
2. Research question, novelty, scope, and exclusions	4
3. Definitions and formal model	6
4. Literature map and competing positions	7
5. Claims and evidence table	9
6. Grounding and projection-equivalence limit	9
7. Executed diagnostic E34-A	10
8. Worked examples	11
9. Proposed protocol: GTP-34 Grounded Trust Path	12
10. Brain-inspired hypothesis, bounded and testable	14
11. Alternative explanations and rival designs	14
12. Falsifiable hypotheses	15
13. Proposed experiment E34-B - external bootstrap and semantic handoff	16
14. Expected failure modes and design responses	19
15. What is established, mixed, open, and hypothesized	19
16. Limitations	21
17. Readiness decision	21
18. Next-round handoff	22
19. Reproducibility artifacts	22
20. Source inspection ledger	24
21. Full sources	24
22. Final answer	25

Abstract

Round 33 showed that a witness quorum cannot certify its own administrative independence merely by presenting distinct keys or self-attested root labels. It proposed AIC-33, which requires externally anchored witness-to-authority bindings, and handed off either a three-root external pilot or a basis-completeness red team. No independently administered authority, external model runner, or separate fixture author was available for this run. Round 34 therefore takes the nonredundant fallback: it asks whether a chain of apparently external endorsements becomes trustworthy through cross-certification alone, or whether every accepted chain must terminate in a trust base configured independently of that chain.

The inspected literature gives a strong but bounded answer. X.509 path validation treats a trust anchor as an input selected locally, not as a fact produced by the path [S05]. Certification graphs may contain loops and dead ends, and path builders are expected to detect repeated certificate or name/key pairs [S06]. In-band delivery of a new anchor must be authenticated with existing trust information or an authorized out-of-band mechanism [S07-S09]. Remote attestation still roots verifier and endorser relationships in protected anchor stores [S10]. Threshold signatures, cosigning, and transparency improve compromise resistance and auditability, but they do not automatically establish that the signers are independent or that their root is legitimate [S11, S13-S14]. Douceur's Sybil result explains why a controller able to mint identities can make redundancy fictitious unless identity scarcity is grounded in a trusted authority or another strong assumption [S12].

Round 34 formalizes this boundary as **grounded independence**. The exact diagnostic contains four witness-facing issuers (W0-W3), three recursive delegates (D0-D2), and three configured anchors (A0-A2). Each witness and delegate chooses one parent from the six delegates-or-anchors, producing $6^7 = 279,936$ functional endorsement graphs. A graph is positive when at least three witnesses terminate at three distinct configured anchors. The anchors' authenticity and administrative distinctness are declared inputs; the experiment does not infer them.

Exactly 63,288 graphs are positive and 216,648 are negative. Counting distinct witness keys accepts every graph and falsely certifies all 216,648 negatives. Requiring three distinct immediate issuers still produces 169,992 false certifications. Treating three terminal components-including cycles-as independent produces 58,320 false certifications. Requiring three witnesses merely to reach some anchors, without root diversity, produces 165,888 false certifications. The fully grounded, root-diverse policy is exact by construction. A two-edge fail-closed path budget has zero false certifications but misses 7,776 valid graphs; a three-edge budget misses 648. Treating unresolved two-edge frontiers or cycles as roots restores recall but produces 109,944 false certifications.

The diagnostic also quantifies an observation limit. With only the first path edge visible, 972 observable-equivalence classes contain both positive and negative graphs, forcing at least 39,960 errors for any binary classifier restricted to that projection under the enumerated uniform universe. At depths two and three, the corresponding minima are 7,776 and 648. Four edges are sufficient to expose every anchor or cycle in this bounded graph, eliminating mixed classes. These are combinatorial facts about the declared universe, not empirical rates.

The proposed mechanism is **GTP-34, the Grounded Trust Path protocol**. It freezes the protected claim and anchor-store snapshot, validates identities, scope, freshness, and revocation, builds paths only toward locally configured anchors, rejects cycles and unresolved depth limits as ungrounded, resolves each accepted anchor's administrative root under AIC-33, and certifies only a quorum whose anchors satisfy the declared independence predicate. Missing grounding returns `unproven`; it never creates a new root. Anchor rotation requires continuity under the old and new root policies or an explicit out-of-band recovery event.

No new Encevera answer benchmark or model benchmark was run. No production trust store, key, policy, authority, or deployment was changed. The readiness manifest remains 29 controls present, 4 partial, and 1 absent; critical controls A3, I1, R4, R5, and

C6 remain non-present. Round 34 sharpens the minimum evidence needed to move them. It does not satisfy them.

1. Archive continuity and corrected round number

The automation brief retained the historical placeholder Round 11. Before topic selection, the live Lovable project at project ID 8483bc59-da36-492e-b8b8-6bf684070241 was inspected. The Research overview, numbered rounds index, publications list, and complete source texts for Rounds 33, 32, 10, and 9 showed completed work through Round 33. The correct next sequential number is therefore **Round 34**, not Round 11.

The carry-forward is precise:

Round	Established program contribution	Boundary carried into Round 34
9	Reusable procedures should be learned from verified experience and retained with their success conditions.	GTP-34 is a tested structural proposal, not yet a learned or production procedure.
10	Collaborators should exchange recoverable evidence, revise selectively, and compare against matched-cost baselines that expose harmful revision.	Evidence can be cryptographically intact yet rest on an ungrounded issuer chain.
26	Observationally identical worlds cannot be separated by rules over the shared observation; root-disjoint channels matter.	Truncated or circular trust paths can make grounded and ungrounded worlds observationally equivalent.
29	Interventions can reveal behavioral dependencies, but unresolved provenance equivalence must stay unresolved.	Cross-signatures are observations; they do not by themselves intervene on issuer control.
30	The 34-control RISE gate blocked a model-facing pilot.	The critical-control rule remains in force.
31	A signed local sandbox improved traceability while leaving authority roots collocated.	More local signatures do not create an external trust base.
32	A 256-case boundary holdout classified 240 cases exactly but accepted all 16 shared-administrator aliases.	Key identity and observable behavior did not certify administrative separation.
33	Across 3,375 authority topologies, self-attested roots falsely certified all 3,164 negatives; AIC-33 required external bindings and conservative handling of missing paths.	"External" remains undefined if the issuer's own authority chain is circular, Sybil-controlled, or never reaches a configured anchor.

Round 33 preferred a three-root external pilot. That intervention still could not be performed: the environment supplied no independently administered identity authority, key custodian, deployment control plane, model runner, or separate fixture author. Round 34 therefore executes the permitted basis-completeness red-team direction against the recursive trust graph itself. The blocked verdict is preserved.

2. Research question, novelty, scope, and exclusions

2.1 Primary research question

> When witness-independence claims are supported by chains of cross-certificates or attestations, what observable condition distinguishes a genuinely grounded quorum from a circular or Sybil-generated endorsement structure, and how should bounded verification behave when a path cannot be completed?

The target predicate is:

`GroundedIndependentQuorum(G) = true`

if and only if at least three of four witness paths in graph *G* terminate at three different anchors already present in the relying party's configured trust store.

This target intentionally separates four questions:

- 1 **Signature validity:** was each edge produced under the named key?
- 2 **Path grounding:** does the edge sequence terminate at an accepted anchor?
- 3 **Root diversity:** do the accepted witness paths terminate at sufficiently distinct administrative roots?
- 4 **Semantic competence:** are the witnesses' claims correct?

Round 34 tests questions 2 and 3 in a finite abstraction. It does not test 1 cryptographically or 4 empirically.

2.2 Why this is not Round 33 again

Round 33 examined a fixed-depth table of witness-to-root claims across administrator, key-custody, and deployment layers. Its main adversary was incomplete or self-attested provenance. Round 34 instead gives each alleged issuer its own parent and asks where recursive authorization terminates. The failure objects are now:

- cross-certification cycles that contain many valid-looking identities but no configured root;
- Sybil issuer trees controlled behind nominally distinct names;
- anchor concentration, where many witnesses reach only one or two roots;
- bounded validators that mistake an unresolved frontier for a root;
- rotations that authorize a new root only with the new root itself.

Confidence calibration, compute allocation, and provenance memory were reconsidered and deferred. None can repair a certificate whose terminal trust assumption is absent. The prerequisite question is whether the evidence graph is grounded at all.

2.3 In scope

- Four witness-facing issuers, three recursive delegate issuers, and three configured anchors.
- Every total functional parent assignment in the declared finite universe.
- Cycles, anchor reachability, anchor diversity, bounded-depth validation, and fail-open versus fail-closed handling.
- Exact policy confusion matrices and projection-equivalence classes.
- A protocol for claim binding, path construction, anchor selection, independence resolution, and rotation.
- A future matched-compute model experiment, not executed here.

2.4 Exclusions

- Signature algorithms, key compromise, certificate parsing exploits, side channels, or cryptographic reductions.
- Probabilities of compromise, collusion, or any real deployment topology.
- Hierarchical names, certificate policies, name constraints, path-length constraints, bridges, revocation latency, and all of X.509's full semantics.
- Proof that the three configured anchors are authentic, independent, honest, or exhaustive.
- Claim truth, model reasoning quality, user outcomes, or production availability.
- Literal copying of biological neural mechanisms.

2.5 Epistemic labels

- **Established external finding:** directly supported by an inspected primary paper, standard, or authoritative technical source.
- **Mixed or bounded evidence:** supported under assumptions that materially limit transfer.
- **Observed in E34-A:** produced by the deterministic exact enumeration.
- **Encevra hypothesis:** a proposed design claim awaiting a new intervention.
- **Open question:** not answered by this run.

3. Definitions and formal model

3.1 Operational definitions

Term	Round 34 definition	Invalid substitution
Trust anchor	A key and associated policy already accepted by the relying party through a process independent of the path being evaluated.	The last self-signed object encountered during search.
Grounded path	A finite endorsement path from a witness to a configured anchor, with no repeated delegate and with all required checks passing.	Any path that stops, times out, or reaches a strongly connected component.
Recursive delegate	An issuer whose authority is justified by another issuer edge.	A root merely because its certificate is self-signed.
Cycle	A path that repeats a delegate before reaching an anchor.	A distributed root or evidence of independence.
Frontier	The final node visible when a verifier's depth or resource budget ends before an anchor or known cycle is resolved.	A provisional anchor.
Root diversity	The number of accepted paths whose terminal anchors satisfy the declared administrative-independence relation.	Number of keys, names, certificates, network endpoints, or immediate issuers.
unproven	Available evidence does not justify the requested grounded quorum.	malicious, compromised, or false.
Anchor-store snapshot	The versioned set of anchors and constraints used for one decision.	Whatever roots appear in the evidence bundle.

3.2 Finite universe

Let:

- $W = \{W_0, W_1, W_2, W_3\}$ be witness-facing issuers;
- $D = \{D_0, D_1, D_2\}$ be recursive delegates;
- $A = \{A_0, A_1, A_2\}$ be configured anchors;
- $P = D \cup A$ be the six possible parents;
- $q = 3$ be the quorum.

Each witness has exactly one outgoing edge to P . Each delegate also has exactly one outgoing edge to P . Anchors have no outgoing edge. There are seven independently assigned edges, so the universe contains:

$$|P|^{(|W|+|D|)} = 6^7 = 279,936 \text{ graphs.}$$

Tracing a witness follows its unique parent sequence. A trace either reaches an anchor or repeats a delegate, which identifies a cycle. Because there are only three delegates, every trace's status is determined within four edges from its witness.

The truth predicate is:

$$Y(G) = 1 \text{ iff the set of anchors reached by the four witnesses has cardinality at least three.}$$

This definition gives each witness one vote at most. Two witnesses terminating at the same anchor add availability but not root diversity.

3.3 Declared stopping assumption

A0, A1, and A2 are assumed authentic, in scope, and administratively distinct. That assumption is supplied to the enumerator; it is not produced by any edge in the graph. This is the point of the model. A real deployment must explain how the relying party obtained and protected that anchor store, why those roots satisfy the action's threat model, and how rotation is authorized.

4. Literature map and competing positions

4.1 Rooted path validation: trust is an input

RFC 5280 defines certificate-path validation around a specific trust anchor and treats trust-anchor information as an input. It also makes selection of trusted certification authorities a local, application-specific decision [S05]. This supports a **rooted validation** position: a chain can propagate authority from an accepted root, but it does not manufacture the root's acceptance.

RFC 9925 makes the boundary unusually explicit. A self-signature is self-derived and has limited value for conveying trust; an application may instead authenticate trust-anchor information out of band and need not rely on an issuer signature for that object [S08]. The relevant distinction is not signed versus unsigned. It is authenticated under a prior trust decision versus self-authorizing.

4.2 Graph-rich PKI: cross-certification is useful but loops are not roots

RFC 4158 treats certification as a graph-search problem rather than a single hierarchy. Multiple paths and cross-certificates can improve interoperability and path discovery, but the search goal remains a valid path between the target and a trust anchor. The document explicitly identifies loops and dead ends and recommends detecting repeated certificates and repeated subject-name/public-key pairs [S06].

This supports a **graph flexibility** position without implying circular bootstrap. Cross-certification can connect already trusted domains or provide alternative routes. A closed component disconnected from every accepted anchor remains ungrounded even if every edge verifies cryptographically.

4.3 Bootstrap and rotation require continuity

RFC 9480 addresses in-band provisioning of new trust anchors: the receiver must authenticate the sender using existing trust information without relying on the new anchors contained in that same message, and must verify that the sender is authorized to supply anchors [S07]. RFC 9360 similarly warns that a newly observed self-signed certificate cannot itself trigger a trust-store change [S09].

The Update Framework supplies an operationally mature variation. Clients ship with root keys; metadata roles use signature thresholds; root updates are accepted through a sequence that satisfies both the old root threshold and the new root threshold. If a root threshold is compromised, recovery moves out of band [S11]. This supports **continuity-based rotation**, not trust creation from a new key's self-assertion.

4.4 Attestation moves the appraisal boundary but retains anchors

RFC 9334 separates Attesters, Verifiers, Relying Parties, Endorsers, reference values, and appraisal policies. Yet relying parties trust verifiers through anchor stores, verifiers trust endorsers or devices through anchors, and certificate paths still require validation to an anchor with revocation checking [S10]. The architecture even notes bootstrap sequences in which one side's evidence must be safe to share before mutual trust exists.

This supports **role-separated appraisal**. It does not support the claim that adding an Endorser label creates an independent organization or that recursively attesting every verifier eliminates the first trusted input.

4.5 Thresholds and cosigning improve authorization under a roster

TUF permits threshold signatures for every role and protects the root role especially strongly [S11]. CoSi makes an authority's statement contingent on cosigning by a large and potentially diverse witness set, records missing witnesses, and enables efficient client verification [S13]. These mechanisms can reduce single-key risk and increase accountability.

Their guarantees remain conditional on roster construction and adversary assumptions. If one controller can mint many accepted identities, key count ceases to approximate entity count. Douceur formalizes this Sybil problem: redundancy is undermined when one entity presents many identities, and voucher chains do not establish initial identity uniqueness without a trusted identification authority or strong resource assumptions [S12]. The competing **threshold-is-enough** position is therefore rejected unless the signer-to-authority mapping is itself grounded.

4.6 Transparency improves detectability, not root truth

RFC 9162 specifies append-only logging that can expose suspect certificate issuance and log misbehavior. It explicitly states that logs do not themselves prevent misissuance, and accepted chains still terminate in anchors accepted by the log or client [S14]. Transparency is valuable after grounding: it can reveal equivocation, omission, and unauthorized changes. It is not a substitute for deciding which roots are accepted or whether several roots share a controller.

4.7 Bounded brain-inspired analogy: infer common cause before combining cues

Koerding and colleagues modeled multisensory perception as a joint problem of cue estimation and causal inference: signals should be combined differently depending on whether they plausibly arise from a common cause [S15]. Round 34 uses only this computational analogy. Multiple witness statements should not be counted as independent evidence until the system has evaluated whether their authority paths share a cause.

The analogy yields a testable design idea: a **common-root gate** before quorum aggregation. It does not imply that certificate graphs work like neurons, that Bayesian perception authenticates keys, or that the proposed protocol copies a brain.

5. Claims and evidence table

Claim	Status	Evidence	Boundary
Four distinct witness keys establish a 3-of-4 independent quorum.	Falsified in E34-A	B0 accepts all 279,936 graphs and falsely certifies all 216,648 negatives.	Keys are assumed distinct labels; cryptographic compromise is not modeled.
Three distinct immediate issuers are sufficient.	Falsified in E34-A	B1 has 169,992 false positives and a 72.87% false-discovery rate in the uniform universe.	Immediate names may converge on one root or a cycle.
Distinct terminal components can be treated as distinct roots.	Falsified in E34-A	B2 counts cycles as roots and produces 58,320 false positives.	A cycle identifier is useful for diagnosis, not authorization.
Three witnesses reaching any configured anchors establish root diversity.	Falsified in E34-A	B3 produces 165,888 false positives because multiple witnesses can concentrate on one or two anchors.	Reachability and diversity are separate predicates.
Complete path grounding plus three distinct configured anchors recovers the finite target.	Observed by construction	B4 has 63,288 true positives, 216,648 true negatives, and no errors.	This validates implementation against the declared predicate, not the real anchor assumptions.
A bounded fail-closed validator preserves soundness but can reject valid long paths.	Observed in E34-A	B5 misses 7,776 positives; B6 misses 648; neither has a false positive.	The bounded universe has at most three delegates; real paths and search branching can be larger.
Treating unresolved frontiers as roots restores availability safely.	Falsified in E34-A	B7 restores recall to 100% but adds 109,944 false positives.	Counts are diagnostic under a uniform graph universe, not deployment prevalence.
Any binary classifier over the first three path edges can be correct on every graph.	Falsified in E34-A	Mixed projection classes force at least 39,960, 7,776, and 648 errors at depths 1, 2, and 3, respectively.	The lower bound applies only to classifiers restricted to the defined projection.
A transparency log turns an ungrounded chain into a grounded one.	Unsupported and contradicted by scope	RFC 9162 supports detection and auditing while retaining accepted trust anchors [S14].	A separately governed log can still contribute useful evidence.
GTP-34 will improve model-answer reliability.	Open Encevera hypothesis	This run tests only structure; no answer benchmark or model call occurred.	Requires external roots, hidden fixtures, and a controlled semantic experiment.

6. Grounding and projection-equivalence limit

Let $\text{pi}_d(G)$ be the ordered collection of the four witness path prefixes visible within d edges, including whether each prefix has reached an anchor, exposed a repeated delegate, or ended at an unresolved frontier. A classifier restricted to depth d has the form $f(\text{pi}_d(G))$.

Define an equivalence relation $G \sim_d H$ when $\text{pi}_d(G) = \text{pi}_d(H)$. If one equivalence class contains both $Y=1$ and $Y=0$ graphs, no deterministic classifier over the projection can classify every member correctly: it must assign one label to the entire class. A randomized classifier cannot make its output depend on the hidden continuation either; it can only trade errors between the two labels.

For a finite uniform universe, the minimum binary error count within an equivalence class is the smaller of its positive and negative counts. Summing that minimum over classes gives a lower bound for every depth-restricted classifier. E34-A measures:

Visible path edges	Equivalence classes	Mixed classes	Worlds in mixed classes	Negative / positive in mixed classes	Minimum forced errors	Uniform error floor
1	1,296	972	209,952	169,992 / 39,960	39,960	14.2747%
2	47,511	3,888	42,768	34,992 / 7,776	7,776	2.7778%
3	84,771	648	3,888	3,240 / 648	648	0.2315%
4	90,021	0	0	0 / 0	0	0%

The number of mixed classes need not decrease monotonically because a deeper projection splits broad classes into many narrower ones. What decreases is the unresolved mass and the optimal forced-error

count. At four edges, every trace has either reached an anchor or repeated one of the three delegates, so the target is identifiable in this finite universe.

This result makes a general design boundary concrete: a resource limit may justify returning unproven, but it cannot justify reclassifying an unseen continuation as a new trust root.

7. Executed diagnostic E34-A

7.1 Predeclared policies

ID	Policy	Acceptance rule
B0	Distinct witness keys	Always accept the four named, separately keyed witnesses.
B1	Unique immediate issuers	Accept if at least three distinct immediate parent labels appear among the four witnesses.
B2	Terminal-component diversity	Follow paths fully, but count anchors and delegate cycles alike as terminal roots; accept at three distinct components.
B3	Three anchor-reachable witnesses	Accept if at least three witnesses reach any configured anchors, ignoring how many distinct anchors they reach.
B4	Grounded and independent	Accept if witness paths reach at least three distinct configured anchors. This equals the declared truth predicate.
B5	Depth-2 fail-closed	Search at most two edges and accept only three distinct configured anchors reached within budget.
B6	Depth-3 fail-closed	Search at most three edges and accept only three distinct configured anchors reached within budget.
B7	Depth-2 fail-open	Search two edges but treat unresolved frontiers and observed cycles as terminal roots.

7.2 Reproducibility controls

- The enumerator uses only the Python standard library.
- Every graph is serialized canonically with sorted JSON keys.
- The complete compressed world stream is retained.
- Policy aggregates, first false-positive counterexamples, projection classes, and GTP-34 are separate machine-readable artifacts.
- Six deterministic unit tests check universe size, truth balance, all confusion matrices, equivalence counts, two worked counterexamples, and the canonical aggregate digest.
- Aggregate result digest: ab8dcefb359d2862d25ecaaea920981c5880d0b90e9238387af6a8797a3b752b.
- Per-world stream SHA-256: a30b1c887d51d27801d4d30fd28ccf3f2cd0982bc7891c0c3ed437340753e60d.

7.3 Exact results

The truth set contains 63,288 positive graphs (22.6080%) and 216,648 negative graphs (77.3920%).

Policy	TP	FP	FN	TN	Accepted	Precision	Recall	False-discovery rate
B0 distinct keys	63,288	216,648	0	0	279,936	22.61%	100.00 %	77.39%
B1 unique immediate issuers	63,288	169,992	0	46,656	233,280	27.13%	100.00 %	72.87%
B2 terminal-component diversity	63,288	58,320	0	158,328	121,608	52.04%	100.00 %	47.96%
B3 three anchor-reachable witnesses	63,288	165,888	0	50,760	229,176	27.62%	100.00 %	72.38%

Policy	TP	FP	FN	TN	Accepted	Precision	Recall	False-discovery rate
B4 grounded and independent	63,288	0	0	216,648	63,288	100.00%	100.00%	0%
B5 depth-2 fail-closed	55,512	0	7,776	216,648	55,512	100.00%	87.71%	0%
B6 depth-3 fail-closed	62,640	0	648	216,648	62,640	100.00%	98.98%	0%
B7 depth-2 fail-open	63,288	109,944	0	106,704	173,232	36.53%	100.00%	63.47%

Every percentage in this table is conditional on a universe that weights all 279,936 graphs equally. It must not be read as the probability that a deployed certificate is malicious, circular, or valid.

7.4 Interpretation

The diagnostic separates three mistakes that are often conflated:

- Identity inflation:** B0 and B1 count labels before demonstrating distinct roots.
- Termination inflation:** B2 turns the fact that search stopped in different components into authority.
- Reachability inflation:** B3 verifies that signatures lead to accepted anchors but fails to require independent anchors.

B5 and B6 demonstrate the safe response to incomplete search: abstain. Their false negatives are availability costs, not false assurances. B7 demonstrates the unsafe shortcut: relabeling ignorance as a new root.

8. Worked examples

8.1 Three-name circular endorsement

Suppose the witness edges are:

- W0 -> D0
- W1 -> D1
- W2 -> D2
- W3 -> D0

and the delegate edges are:

- D0 -> D1
- D1 -> D2
- D2 -> D0

Every edge could carry a valid signature. The graph contains three issuer names and a well-formed endorsement for every next hop. Yet all four witness paths enter the same closed cycle and none reaches A0, A1, or A2. B0 accepts. B1 accepts because there are three immediate issuer labels. A policy that treats strongly connected components as self-justifying roots can also accept. B4 and GTP-34 return unproven because a cycle is not a configured anchor.

Adding more cross-signatures inside the component does not alter the result. Those edges may increase evidence about agreement within the component; they do not create an externally authenticated stopping point.

8.2 Four grounded paths, only two roots

Let:

- $W_0 \rightarrow D_0, D_0 \rightarrow A_0$
- $W_1 \rightarrow D_1, D_1 \rightarrow A_0$
- $W_2 \rightarrow D_2, D_2 \rightarrow A_0$
- $W_3 \rightarrow A_1$

All four witnesses reach configured anchors, and every path is short. B3 accepts because at least three witnesses are anchor-reachable. The target is nevertheless false: the terminal anchor set is $\{A_0, A_1\}$, with cardinality two. Three witness identities have collapsed onto A_0 .

GTP-34 preserves the four approvals but reports only two independent anchor roots. The action requiring a 3-root quorum remains `unproven`. This distinction matters operationally: the system need not discard the evidence or accuse any signer; it simply must not upgrade two roots into three.

8.3 A valid path beyond a shallow budget

Let three witnesses follow distinct direct or short paths to A_0 and A_1 , while the third required root is reached by $W_2 \rightarrow D_0 \rightarrow D_1 \rightarrow D_2 \rightarrow A_2$. A two-edge or three-edge validator cannot yet establish the third anchor. B5 or B6 returns negative even though the complete graph is positive.

The correct runtime state is `unproven_pending_path`, not `certified` and not necessarily `invalid`. The verifier may spend a predeclared larger budget, request the missing edge, or defer the action. If the maximum policy depth is reached, the action remains blocked unless another already grounded quorum exists.

9. Proposed protocol: GTP-34 Grounded Trust Path

9.1 Inputs

Every decision binds:

- protected claim or action digest;
- witness roster and quorum;
- anchor-store snapshot identifier and its administrative-root map;
- permitted certificate or attestation profiles;
- path constraints, time window, revocation policy, and search budget;
- AIC-33 independence predicate;
- unknown-data policy and required decision state;
- audit-log destination and retention policy.

9.2 Three decision states

- **certified:** a complete valid set of paths reaches enough allowed, administratively independent anchors.
- **unproven:** signatures may be valid, but grounding, freshness, revocation, coverage, or independence is insufficient.
- **invalid:** a required cryptographic, identity, scope, policy, or consistency check fails.

Only `certified` authorizes the protected action. `unproven` preserves uncertainty without converting it into an accusation.

9.3 Decision procedure

- 1 **Freeze context.** Record the action digest, roster, quorum, policy, anchor-store version, and budgets before evaluating approvals.
- 2 **Authenticate objects.** Verify signatures, canonical encodings, key bindings, role scope, validity intervals, and replay protections.
- 3 **Build paths toward allowed anchors.** Candidate roots come only from the frozen local anchor store. Evidence cannot add itself to that store during the same decision.
- 4 **Detect cycles and aliases.** Reject repeated certificate identities and repeated name/key pairs on a path. Record strongly connected components for diagnosis, never as roots.
- 5 **Apply revocation and constraints.** Validate every required path element, including revocation state and application-specific constraints.
- 6 **Handle resource limits conservatively.** A missing edge, timeout, lookup failure, or depth limit yields an unresolved frontier. It contributes no root.
- 7 **Resolve administrative roots.** Apply AIC-33 to the accepted anchors and conservatively couple anchors whose independence is missing, stale, contradictory, or self-attested.
- 8 **Evaluate the quorum.** Find a witness subset of size at least q whose paths are valid and whose resolved anchor roots satisfy the independence predicate.
- 9 **Emit a proof-carrying result.** Preserve accepted paths, rejected alternatives, cycles, missing edges, root mappings, solver output, and residual assumptions.
- 10 **Log without laundering.** Commit the result to an independently monitored append-only log. Log inclusion proves recording under that log's trust model, not truth of the certificate's inputs.

9.4 Anchor enrollment and rotation

GTP-34 distinguishes ordinary path evidence from trust-store administration:

- **Initial enrollment:** install anchor fingerprints and constraints through an authenticated out-of-band ceremony or a previously trusted management channel.
- **Routine rotation:** require authorization under both the current root policy and the proposed new root policy, with monotonically increasing versions and explicit activation time.
- **Emergency recovery:** if the current root threshold is compromised or unavailable, invoke a predeclared out-of-band recovery authority; do not accept a replacement merely because it signs itself.
- **Removal:** record revocation or distrust with effective time, reason, and consequences for cached certificates.
- **Rollback protection:** never accept an older anchor-store snapshot for a newer protected decision.

This design mirrors the continuity principle used by established trust-management systems [S07, S11] while keeping the actual Encevera policy falsifiable.

9.5 Residual trust statement

Every certificate must state:

> This decision is grounded only relative to anchor-store snapshot T , its administrative-root map M , and the declared dependency basis L . GTP-34 proves neither that T is honest nor that L contains every real common cause.

Without that sentence or its machine-readable equivalent, “grounded” is too easy to misread as absolute trust.

10. Brain-inspired hypothesis, bounded and testable

The useful analogy from multisensory causal inference is not that agents should imitate cortex. It is that aggregation and causal structure should be coupled: two strong signals do not provide two independent pieces of evidence if they likely share one cause [S15].

GTP-34 operationalizes a narrow version:

- 1 infer or verify the authority path behind each witness signal;
- 2 cluster signals by grounded administrative root;
- 3 aggregate across roots rather than raw signals;
- 4 return unproven when the common-cause structure cannot be resolved.

This produces a falsifiable hypothesis: under a matched verification and model-compute budget, root-gated aggregation will reduce false authorization on hidden shared-cause cases more than it increases false abstention on genuinely independent clean cases. A future experiment can reject the hypothesis if the gate does not improve that tradeoff or if a simpler non-biological graph validator performs equally well. The biological analogy earns its place only by motivating that comparison.

11. Alternative explanations and rival designs

11.1 “The problem is only the chosen truth definition”

The large false-positive counts depend on defining independence as three distinct configured anchors. A deployment that accepts three witnesses under one anchor may reasonably call that a participation quorum. Round 34 does not dispute that label. It rejects calling the same object a three-root independence quorum. The remedy is to name the property precisely and avoid transferring guarantees from one property to another.

11.2 “Cycles can be roots by social consensus”

A community may decide to accept a web-of-trust component or a federated root set. That decision is itself an anchor policy: the relying party has chosen a membership rule, admission process, and acceptable initial state. The internal endorsements can then propagate or update that policy. They do not explain why the relying party should accept the initial component or why several members correspond to several entities.

11.3 “Proof of work, stake, hardware, or personhood removes the anchor”

Resource and hardware mechanisms can make identities costly and may substitute a different trust assumption for a registry. Their security then depends on assumptions about resource scarcity, hardware supply chains, measurement, distribution, and capture. Douceur explicitly leaves room for resource-based assumptions while showing that unconstrained identity vouchers do not suffice [S12]. GTP-34 can represent such a mechanism as an anchor policy, but must not label it assumption-free.

11.4 “A transparent log is the anchor”

A log can be made an accepted anchor or evidence source, but then its key, governance, monitors, consistency, and recovery policy enter the trust basis. Certificate Transparency's scope is detection and auditability, not automatic prevention or issuer correctness [S14]. A log operated inside the same failure domain as all witnesses may preserve evidence without adding independence.

11.5 “The false positives come from an artificial uniform distribution”

The precision and false-discovery percentages do depend on uniform weighting. The existence of counterexamples and the projection-equivalence lower bounds do not depend on claiming that deployed graphs are uniform. A production claim needs a deployment-specific threat distribution; the finite universe is used to falsify universal sufficiency claims and to test implementation behavior exhaustively.

11.6 “A longer search always solves the problem”

In E34-A, four edges resolve every path because the graph has only three delegates. Real certificate graphs can branch, contain many intermediates, suffer missing retrieval, or change during validation. More search improves completeness but cannot authenticate a previously untrusted terminal node. The safe fallback remains unproven, with an auditable reason and budget trace.

11.7 “The proposed protocol is just ordinary PKI”

GTP-34 intentionally borrows established path-validation and rotation principles. Its program-specific contribution is their integration with Encevera's AIC-33 administrative-independence predicate, three-state authorization semantics, exact bounded-depth counterexamples, and a matched-compute experimental plan for collaborative reasoning. Novelty is not claimed for trust anchors or loop detection themselves.

12. Falsifiable hypotheses

The next interventions should register these hypotheses before examining hidden labels:

H34.1 - Rootless-cycle soundness

For any fixture whose complete signer paths remain inside a strongly connected component with no configured anchor, GTP-34 will never return `certified`, regardless of the number of valid signatures inside the component.

Falsifier: one such fixture receives `certified` without an explicit policy that preconfigured the component as an anchor.

H34.2 - Root-diversity soundness

When the action requires q independent roots, GTP-34 will not count multiple witness paths terminating at the same administrative root more than once.

Falsifier: a $q=3$ action is certified with fewer than three distinct accepted administrative roots.

H34.3 - Conservative frontier handling

Timeouts, missing intermediates, depth limits, and lookup failures will increase `unproven`, never create a new accepted root.

Falsifier: removing path information or reducing the search budget changes an otherwise identical result from `unproven` or `invalid` to `certified`.

H34.4 - Rotation continuity

An in-band anchor update will be accepted only when authorized under the current anchor policy and the proposed new policy, except during a separately authenticated recovery event.

Falsifier: a new root becomes active solely through its own signature or signatures traceable only to that new root.

H34.5 - Projection lower bound

Any deterministic classifier restricted to E34-A's depth-d projection will commit at least 39,960, 7,776, and 648 errors for $d=1, 2, 3$, respectively, over the full uniform universe.

Falsifier: a reproducible classifier using only the recorded projection achieves fewer errors. Such a result would indicate an implementation or equivalence-accounting defect.

H34.6 - Hidden-graph generalization

On independently authored structural fixtures, GTP-34 will reduce critical false certification relative to distinct-key, immediate-issuer, reachability-only, and fail-open baselines without changing the anchor store after seeing fixture labels.

Falsifier: GTP-34 fails to improve critical false certification, relies on post-label trust-store edits, or its gain disappears under a matched verification budget.

H34.7 - Semantic value under matched model compute

After A3, I1, R4, R5, and C6 are present, root-gated collaboration will reduce final-answer error or harmful revision on shared-cause fault tasks relative to raw 3-of-4 voting at the same total model and tool budget, without a prespecified unacceptable regression on clean tasks.

Falsifier: no reliable benefit, excessive clean-case abstention, worse calibration, or increased harmful revision under the preregistered analysis.

13. Proposed experiment E34-B - external bootstrap and semantic handoff

This experiment has **not** been run. It is a protocol for the next evidence-producing intervention.

13.1 Phase 1: independently authored trust-graph fixtures

An external fixture author constructs signed or faithfully simulated bundles after receiving the frozen schema but not the implementation internals. The suite should balance:

- three genuinely distinct, preprovisioned roots;
- one-root and two-root concentration;
- two-node and three-node cross-certification cycles;
- Sybil issuer names under one administrative root;
- a valid long path that exceeds shallow budgets;
- expired, revoked, not-yet-valid, and wrong-scope intermediates;
- stale anchor-store snapshots and rollback attempts;
- valid dual-threshold root rotations;
- self-authorized root rotations;
- transparency inclusion with invalid grounding;
- missing intermediates and network timeouts;
- conflicting paths to the same key;
- hidden shared administration behind nominally different anchors.

The fixture authority supplies ground truth only after outputs are committed. If no independent author is available, the phase remains a red-team rehearsal and cannot close R4 or R5.

13.2 Structural arms

Arm	Description	Purpose
S0	Distinct-key 3-of-4 threshold	Naive identity-count baseline.
S1	Three unique immediate issuers	Tests one-hop diversity.
S2	Full path reachability, no root-diversity check	Separates grounding from independence.
S3	Depth-limited fail-open	Measures the dangerous availability shortcut.
S4	GTP-34 complete	Proposed protocol.
S5	GTP-34 without cycle detection	Ablates loop handling.
S6	GTP-34 without administrative-root resolution	Ablates AIC-33.
S7	GTP-34 without old/new rotation continuity	Ablates bootstrap continuity.
S8	Transparency-only acceptance	Tests whether logging is misused as validity.

13.3 Controlled verification budget

For each fixture, every arm receives the same frozen evidence bundle, anchor-store snapshot, cache state, wall-clock deadline, maximum edge fetches v , and cryptographic-verification allowance K . Any arm that exhausts a budget must emit its native state; it may not classify an unresolved path using hidden data. Report both nominal limits and actual usage.

Two budget strata should be preregistered:

- **constrained:** enough for common two-edge paths but not every adversarial long path;
- **complete for the public bound:** enough to explore every supplied structural path.

This distinguishes algorithmic unsoundness from insufficient search resources.

13.4 Phase 2: matched-compute collaborative reasoning

Phase 2 starts only after critical structural controls A3, I1, R4, and R5 are present and C6 supplies a model runner. Each task receives the same:

- model family and version;
- total generated-token budget B across all agents;
- maximum tool calls and retrieved bytes;
- evidence corpus snapshot;
- wall-clock envelope;
- random-seed schedule where supported;
- final-answer format and scoring rubric.

If one arm uses more agents, its per-agent token allocation is reduced so total model tokens remain equal. Verification overhead is recorded separately and capped under the same structural budget used in Phase 1. No arm may spend saved budget after seeing another arm's output.

Proposed semantic arms:

- 1 single strong agent with the entire model budget;
- 2 raw 3-of-4 vote over four witness outputs;

- 3 Round 10 evidence exchange and selective revision without root gating;
- 4 AIC-33 immediate-root gating;
- 5 full GTP-34 path grounding plus root gating;
- 6 full GTP-34 with common-root clustering removed;
- 7 full GTP-34 with a randomized extra verifier consuming the same verification budget.

The task set must contain clean cases, independent-error cases, common-retrieval-source corruption, common-prompt or policy corruption, one compromised tool channel, and stale or selectively omitted evidence. A separate evaluator, outside all authorizing roots, holds labels and adjudication criteria.

13.5 Metrics

Structural primary metrics

- critical false-certification rate;
- critical false-rotation-acceptance rate;
- root-concentration false acceptance;
- cycle false acceptance;
- valid-fixture certification recall;
- unproven and invalid rates by cause;
- path-completion rate, edge fetches, signature checks, latency, and cache hits;
- anchor-store rollback and revocation-detection rates.

Semantic primary metrics

- final-answer error rate;
- harmful revision: initially correct answer changed to incorrect;
- beneficial revision: initially incorrect answer corrected;
- unsupported-claim rate and citation entailment;
- false authorization on common-cause faults;
- selective risk/coverage and abstention quality;
- Brier score, log loss, and expected calibration error for decision confidence.

Regression sentinels

- clean-case answer accuracy and latency;
- false abstention on fully grounded clean cases;
- loss of minority correct evidence through root clustering;
- error concentration by task type, root, issuer, and path length;
- extra tokens, tool calls, bytes, verification operations, and wall time;
- disagreement suppression and premature consensus;
- stale-anchor acceptance after rotation.

13.6 Analysis and stopping rules

- Predeclare primary outcomes, family-wise error or false-discovery control, confidence intervals, and subgroup analyses.
- Pair arms on identical tasks and seeds.
- Publish every exclusion and timeout as an outcome, not a silently dropped row.
- Treat zero observed critical failures as an empirical result with uncertainty, never as proof of zero risk.
- Stop and preserve the blocked state if any self-authorized rotation, rootless-cycle certification, anchor rollback, or evidence-removal monotonicity violation occurs.
- Do not promote GTP-34 if semantic gains disappear at matched model compute or if clean-case harmful revision breaches the preregistered margin.

14. Expected failure modes and design responses

Failure mode	Why it matters	Required response
Anchor laundering	A presented self-signed object is silently inserted into the trusted set.	Freeze the anchor store before evidence evaluation; require separately authorized enrollment.
Cross-signature cycle	Several valid signatures form a closed component.	Detect repeated name/key pairs and return unproven; never promote the component.
Sybil delegates	One controller creates many issuer names or keys.	Resolve administrative roots with AIC-33; count roots, not labels.
Root concentration	Many witnesses reach one accepted root.	Evaluate terminal-root diversity after path validation.
Frontier promotion	Timeout or depth exhaustion is treated as an implicit root.	Fail closed and preserve the unresolved frontier.
Search denial of service	An attacker supplies a large branching graph.	Use frozen budgets, ranking, caching, and unproven; do not weaken acceptance.
Revocation ambiguity	A valid-looking path contains stale or unreachable status information.	Follow the declared revocation policy and record soft-fail versus hard-fail semantics explicitly.
Rotation takeover	A new root authorizes itself or an attacker replays an old root set.	Require old/new continuity, monotonic versions, and rollback protection.
Log laundering	Inclusion in a transparent log is treated as correctness.	Separate inclusion, consistency, path validity, and independence fields.
Anchor collusion	Configured roots are distinct in the store but share an operator or fail together.	Apply AIC-33 to anchors; leave the result unproven when administrative evidence is incomplete.
Basis omission	All declared paths are correct, but a shared model, data source, policy owner, or cloud control plane is omitted.	Declare the basis and red-team it with hidden common causes; never claim universal completeness.
Correct structure, wrong content	Independent agents agree on a false answer.	Preserve semantic evaluation; structural independence is not truth.
Operational bypass	Users disable the gate because abstention is costly.	Measure bypass pressure, latency, and false-abstention costs; design reviewable escalation rather than silent override.

15. What is established, mixed, open, and hypothesized

Established by inspected external sources

- Standard certificate-path validation begins from a locally selected trust anchor supplied as input [S05].

- Certification graphs can contain loops and dead ends; repeat detection is part of safe path building [S06].
- A new in-band anchor cannot safely authenticate its own introduction; prior trust or an out-of-band mechanism is required [S07-S09].
- Attestation verifiers and relying parties still depend on protected anchor stores and full path validation [S10].
- Threshold roles and witness cosigning distribute authorization under a roster but do not by themselves establish the roster's real-world independence [S11, S13].
- Unconstrained multiple identities can defeat redundancy; voucher chains inherit an identity-bootstrap problem [S12].
- Transparency can reveal suspect issuance and log faults but does not itself prevent misissuance [S14].

Observed exactly in E34-A

- The enumerated universe contains 279,936 graphs, of which 63,288 satisfy the stated grounded three-root predicate.
- Key count, immediate-issuer diversity, terminal-component diversity, and anchor reachability without diversity are unsound for that predicate.
- Fail-closed bounded validation trades recall for soundness in the finite universe.
- Fail-open frontier promotion creates 109,944 false certifications at depth two.
- Truncated projections contain mixed truth classes until all possible delegate continuations are exposed.

Mixed or bounded evidence

- Existing PKI and software-update standards are strong analogies for trust-path and rotation design, but Enceavra's agent collaboration threat model includes semantic and organizational dependencies those standards do not directly evaluate.
- CoSi and threshold frameworks show scalable multiparty authorization, while their practical independence depends on governance and signer selection.
- Causal-inference research motivates common-root gating, but provides no cryptographic security result.

Open questions

- Can three genuinely distinct external roots be provisioned and maintained for an Enceavra pilot?
- Which dependency layers are necessary and sufficient for a given decision class?
- How often do valid real paths exceed operational verification budgets?
- Can root-gated collaboration improve answer quality at matched model compute?
- How much clean-case availability loss will users tolerate before bypassing the control?
- How should a trust basis recover after its existing threshold is compromised?

Enceavra hypotheses

- A complete grounded path plus AIC-33 root resolution will sharply reduce common-cause false authorization.
- A three-state result will preserve soundness better than a forced binary choice under missing evidence.

- Old/new threshold continuity will prevent silent in-band root replacement in ordinary rotation.
- Common-root gating will improve semantic reliability on shared-cause tasks without unacceptable clean-case regression when compared at fixed total compute.

16. Limitations

- 1 E34-A uses a functional graph with one parent per witness or delegate. Real PKI can have multiple candidate issuers, path constraints, bridge CAs, policy mappings, name constraints, and dynamic discovery.
- 2 The three anchors are assumed authentic and administratively distinct. The code cannot establish that assumption, and the real problem recurs at their governance boundary.
- 3 “Three distinct anchors” is a chosen policy predicate. Other actions may require different thresholds, weighted roots, jurisdictions, hardware properties, or adversary structures.
- 4 The universe weights every graph equally. Its precision and false-discovery rates are not population estimates.
- 5 Signature forgery, key theft, compromised validators, parser bugs, side channels, and denial-of-service branching are outside the exact model.
- 6 Revocation and temporal change are described in GTP-34 but not enumerated in E34-A.
- 7 Projection-equivalence counts are exact only for the recorded prefix representation and finite graph. Other observations can split those classes if they are independently authenticated.
- 8 B4 is exact by construction because it implements the declared truth predicate. It is a reference oracle inside the model, not evidence that a production implementation is secure.
- 9 The external sources establish protocol principles under their own scopes; they do not certify GTP-34 or Encevera.
- 10 The biological analogy is deliberately narrow and makes no claim of mechanistic equivalence.
- 11 No independent fixture author, external administrator, HSM, separate cloud control plane, or model runner participated.
- 12 No user-impact, production-operations, model-answer, or cost-effectiveness evidence was collected.

17. Readiness decision

The Round 33 readiness manifest is carried forward without credit for new production controls:

Status	Count
Present	29
Partial	4
Absent	1

The critical non-present controls remain:

- **A3:** genuine administrative separation;
- **I1:** independent observation;
- **R4:** independently administered witness;
- **R5:** receipt path outside the actor's control boundary;
- **C6:** model runner for an actual model-facing pilot.

Decision: BLOCKED. E34-A supplies no new external authority. It demonstrates why a chain of self-consistent attestations, cross-certificates, or witness keys cannot be used to mark these controls present unless paths terminate at independently provisioned and administratively resolved anchors.

No new Encevera answer benchmark was run. No model call was made. No production trust store, key, deployment, policy, or protected action was changed.

18. Next-round handoff

Round 35 should be an **external bootstrap and rotation intervention**, not another self-generated graph study.

Preferred intervention

- 1 Provision at least three trust anchors through separately administered channels, record their fingerprints and constraints out of band, and map their administrative roots under AIC-33.
- 2 Give an independent fixture author the frozen GTP-34 schema and ask for hidden valid paths, cycles, Sybil issuers, one-root and two-root concentration, revocations, rollbacks, and root rotations.
- 3 Require both old and new root thresholds for normal rotation; include a separately authenticated recovery fixture.
- 4 Run the Phase 1 structural arms at equal verification budgets and adjudicate labels only after outputs are committed.
- 5 Update A3, I1, R4, and R5 only from evidence actually supplied by those external systems.
- 6 Activate C6 and the Phase 2 model experiment only if all critical structural controls are present.

If external roots remain unavailable

Do not add another local “independent” signer. Ask an independent reviewer to attack the declared dependency basis or pause the sequence with the blocked state preserved. Repeating an internally authored simulation would improve implementation coverage but would not produce the external evidence the program now lacks.

19. Reproducibility artifacts

The accompanying package contains:

- this complete manuscript;
- `round34_trust_graphs.py`, the standard-library exact enumerator;
- `test_round34.py`, six deterministic tests;
- `summary.json`, complete aggregate results and digests;
- `worlds.jsonl.gz`, all 279,936 canonical world records;
- `policy_results.csv`, exact confusion matrices and derived metrics;
- `equivalence_classes.csv`, projection-equivalence counts;
- `counterexamples.json`, the first false-positive world for each unsound baseline;
- `gtp34_protocol.json`, the proposed decision procedure;
- `source_ledger.csv`, publication dates, inspected portions, support, and limits;

- README.md and SHA256SUMS.txt.

Reproduction commands:

```
``bash python repro/round34_trust_graphs.py --output results_final cd repro python -m unittest -v test_round34.py ``
```

Expected aggregate result digest:

ab8dcefb359d2862d25ecaaea920981c5880d0b90e9238387af6a8797a3b752b

20. Source inspection ledger

The machine-readable ledger is included as `source_ledger.csv`. The summary below records what was actually inspected and the evidentiary role of each source.

ID	Source and date	Material inspected	What it supports	Important limit
S01	Enceavra Round 9, archive	Full page and archive context	Verified experience and conditional procedure retention.	Does not address recursive trust paths.
S02	Enceavra Round 10, archive	Full page	Recoverable evidence, selective revision, matched-cost evaluation.	Does not establish trust anchors.
S03	Enceavra Round 32, 2026-10-23	Full canonical text and route	240/256 exact boundary classification; 16/16 shared-admin aliases accepted.	No external root was provisioned.
S04	Enceavra Round 33, 2026-10-24	Full manuscript, overview, index, publications, source tree	AIC-33, exact 3,375-topology result, blocked readiness, Round 34 handoff.	Stops at “externally anchored” without recursively testing issuer paths.
S05	RFC 5280, May 2008	Sections 6.1 and 6.2	Trust anchors are validation inputs selected locally.	Does not prove anchor independence.
S06	RFC 4158, September 2005	Graph/path-building and loop-detection sections	Paths can loop or dead-end; repeats must be detected.	Loop detection alone does not select a legitimate root.
S07	RFC 9480, November 2023	Trust-anchor provisioning update, Section 8.6	New anchors require authentication under existing trust or authorized policy.	Does not remove out-of-band assumptions.
S08	RFC 9925, February 2026	Abstract and trust-anchor/self-signature discussion	Self-signatures have limited trust-conveying value; anchors may be authenticated out of band.	Does not make unsigned objects generally trustworthy.
S09	RFC 9360, February 2023	Trust-establishment security considerations	A newly observed self-signed certificate cannot trigger trust-store modification by itself.	Not a complete Enceavra protocol.
S10	RFC 9334, January 2023	Roles, bootstrap, trust-anchor protection, full-path validation	Attestation remains grounded in anchor stores and explicit role trust.	Role separation is not organizational separation.
S11	TUF Specification 1.0.37, 2026-10-08	Root role and key migration	Preprovisioned roots, threshold roles, old/new rotation continuity, out-of-band recovery.	Thresholds do not automatically make administrators independent.
S12	Douceur, The Sybil Attack, 2002	Full paper	One entity can defeat redundancy through multiple identities; voucher chains require a grounding assumption.	Does not require one particular centralized implementation.
S13	Syta et al., CoSi, 2016	Abstract, threat model, construction, witness exceptions, evaluation summary	Scalable witness cosigning and auditability under a roster.	Does not prove roster independence.
S14	RFC 9162, December 2021	Abstract and introduction	Transparency supports detection and auditing, not automatic prevention.	Log inclusion does not establish certificate truth.
S15	Koerding et al., 2007-09-26	Full article abstract, model, results, discussion	Cue aggregation can depend on inferred common cause.	Analogy only; no cryptographic claim.
S16	E34-A, 2026-10-25	All worlds, aggregates, tests, digests	Exact results inside the declared finite model.	No production or model-quality inference.

21. Full sources

[S01] Enceavra Research. “Research Round 9.” Internal canonical archive, inspected October 25, 2026.

[S02] Enceavra Research. “Research Round 10: Evidence exchange, selective revision, and a testable path toward brain-inspired intelligence.” Internal canonical archive, inspected October 25, 2026.

[S03] Enceavra Research. “Research Round 32.” October 23, 2026. Internal canonical archive.

- [S04]** Encevera Research. “Research Round 33: Who Certifies Independence? Provenance Completeness and the Limits of Self-Attested Witness Quorums.” October 24, 2026. Internal canonical archive.
- [S05]** Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk. “Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile.” RFC 5280, May 2008. <https://www.rfc-editor.org/rfc/rfc5280>
- [S06]** Cooper, M., Dzambasow, Y., Hesse, P., Joseph, S., and R. Nicholas. “Internet X.509 Public Key Infrastructure: Certification Path Building.” RFC 4158, September 2005. <https://www.rfc-editor.org/rfc/rfc4158>
- [S07]** Brockhaus, H., von Oheimb, D., and J. Gray. “Certificate Management Protocol (CMP) Updates.” RFC 9480, November 2023. <https://www.rfc-editor.org/rfc/rfc9480>
- [S08]** Benjamin, D. “Unsigned X.509 Certificates.” RFC 9925, February 2026. <https://www.rfc-editor.org/rfc/rfc9925>
- [S09]** Schaad, J. “CBOR Object Signing and Encryption (COSE): Header Parameters for Carrying and Referencing X.509 Certificates.” RFC 9360, February 2023. <https://www.rfc-editor.org/rfc/rfc9360>
- [S10]** Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan. “Remote ATtestation procedureS (RATS) Architecture.” RFC 9334, January 2023. <https://www.rfc-editor.org/rfc/rfc9334>
- [S11]** Cappos, J., Kuppusamy, T. K., Lock, J., Moore, M., and L. Puehringer, editors. “The Update Framework Specification.” Version 1.0.37, last modified October 8, 2026. <https://theupdateframework.github.io/specification/latest/>
- [S12]** Douceur, J. R. “The Sybil Attack.” Peer-to-Peer Systems: First International Workshop, IPTPS 2002, pp. 251-260, 2002. <https://www.freehaven.net/anonbib/cache/sybil.pdf>
- [S13]** Syta, E., Tamas, I., Visser, D., Wolinsky, D. I., Jovanovic, P., Gasser, L., Gailly, N., Khoffi, I., and B. Ford. “Keeping Authorities ‘Honest or Bust’ with Decentralized Witness Cosigning.” 2016 IEEE Symposium on Security and Privacy, 2016. <https://dedis.cs.yale.edu/dissent/papers/witness.pdf>
- [S14]** Laurie, B., Messeri, E., and R. Stradling. “Certificate Transparency Version 2.0.” RFC 9162, December 2021. <https://www.rfc-editor.org/rfc/rfc9162>
- [S15]** Koerding, K. P., Beierholm, U., Ma, W. J., Quartz, S., Tenenbaum, J. B., and L. Shams. “Causal Inference in Multisensory Perception.” PLOS ONE 2(9): e943, September 26, 2007. <https://doi.org/10.1371/journal.pone.0000943>
- [S16]** Encevera Research. “E34-A Recursive Endorsement Graph Enumeration.” Generated October 25, 2026. Accompanying reproducibility bundle.

22. Final answer

Cross-certification can transmit trust, distribute authorization, and improve auditability. It cannot create a justified trust base merely by becoming deeper, denser, or circular. A relying party must eventually reach a root accepted through a process outside the path being evaluated, and an independence claim must count administratively distinct accepted roots rather than keys, names, reachable paths, or terminal graph components.

Within E34-A's declared finite universe, that boundary is exact. Fail-closed bounded validation loses coverage but preserves soundness. Fail-open frontier promotion restores coverage by manufacturing false roots and therefore fails. GTP-34 turns the result into an implementable, testable protocol while keeping its residual assumptions visible.

The next scientific step is not another local receipt. It is an externally administered bootstrap-and-rotation intervention with hidden graph fixtures, followed-only if the critical readiness controls are present-by a matched-compute model experiment measuring whether grounded independence improves answer reliability without harmful clean-case regressions.