

GOVERNED AUTONOMY

CSSL: Semantic change, accountable evidence, and operational continuity

A paper in the Governed Autonomy collection

Andrew Durden
Independent research
Publication edition 2026.10.7

Research on authority, evidence, and continuity in autonomous systems.

Content identifier: 30f2410373ee5d9211c2ab5e

Contents

CSSL: Semantic change, accountable evidence, and operational continuity 3

Abstract 3

1. Why a history of edits is insufficient 3

2. Identity, revision, and meaning 4

3. Information that actually supported the work 5

4. Coordinated change and accountable authority 5

5. Evidence supports particular claims 7

6. Effects, privacy, and replay 8

7. From changed source to observed operation 8

8. Four illustrative cases 9

9. Threat assumptions, limits, and research questions 10

Source guide 11

CSSL: Semantic change, accountable evidence, and operational continuity

Governed Autonomy · Publication 2026.10.7

Research study: conceptual exposition of the Continuous Semantic State Ledger specification

Source: Public exposition, successor 2026.10.7; historical research and evidence separately qualified

Abstract

The Continuous Semantic State Ledger (CSSL) studies how software can change under continuing human and machine participation without losing the relationship between meaning, authority, evidence, and consequence. Its subject extends beyond recording edits. A change can be internally consistent yet unauthorized, properly approved yet based on stale information, reproducibly built yet never deployed, or successfully dispatched without a known external result. Treating these conditions as one notion of success makes autonomous work difficult to govern.

CSSL proposes a common account of semantic identity, dependent information, coordinated change, evidence, information flow, and operational history. The central research contribution is their composition: a continuing system should be able to explain what changed, which assumptions supported it, who could authorize it, what was observed, and what remains uncertain. These questions remain distinct even when their answers are recorded together.

This paper develops that account through four illustrative cases: coordinated change, incomplete information, uncertain external effects, and revised evidence. The Sovereign-Closure revision retains open construction, convergence, and verification obligations. Its models and performance targets define requirements for evaluation; a complete implementation and demonstrated conformance remain to be established. The discussion identifies the practical questions such an evaluation would need to answer.

1. Why a history of edits is insufficient

Software work changes more than text. A renamed component may retain its purpose and dependencies. An unchanged component may acquire a different meaning because an interface, policy, environment, or dependency changed around it. A generated artifact can have the same bytes as an earlier artifact while belonging to a different authorized release. Conversely, a small edit can alter the conditions under which a large set of consumers remains valid.

The problem becomes more demanding when autonomous authors work continuously. They receive information through searches, tools, services, and previous work. They produce drafts, request execution, and respond to failures while other authors change the environment. A review that sees only the final textual difference cannot establish which

information was actually available, whether an apparent absence was searched comprehensively, or whether an approval still covers the resulting work.

CSSL therefore treats change as a relationship among an enduring subject, its particular revision, the information supporting the revision, and the authority and evidence associated with its consequences. A record of those relationships can support explanation and checking. Recording them cannot by itself make the information correct or the proposed change safe. This limitation matters throughout the study: a richer ledger improves the questions that can be asked, but the answers still depend on the quality and completeness of its producers.

The collection thesis places this problem within governed autonomy. The architecture study discusses the broader separation between reasoning and authority. Here the concern is the continuity of change across those boundaries. Public review scope.

2. Identity, revision, and meaning

An enduring entity and a particular representation of that entity answer different questions. Identity asks whether this is the continuing subject of earlier work. Revision asks which version is under discussion. Content identity asks whether particular bytes are equal. A name gives people or other components a way to refer to a subject, but a name alone does not establish its continuity.

These distinctions have practical consequences. Renaming a function need not create an unrelated function. Copying its text need not transfer its identity or the approvals attached to it. Deleting a subject and later creating another subject with the same name need not restore the earlier subject. Generated output also requires a lineage account: equal output does not explain which inputs, generator, or environment produced it.

Meaning introduces a further boundary. Different readers depend on different aspects of an artifact. One consumer may depend on a public interface; another may depend on behavior; a formatting tool may depend on exact syntax. A change that is irrelevant to one reader can invalidate another reader's assumptions. CSSL uses the idea of an explicit semantic interpretation to make this dependence discussable. Such an interpretation is itself a contract with limits. It cannot be assumed complete merely because it has a formal name.

The source consequently distinguishes stronger semantic understanding from conservative treatment when understanding is absent. A system must not announce that a change is irrelevant when its analysis does not cover the affected meaning. This is particularly important across language boundaries, generated artifacts, configuration, data, and operational policy. The available interpretation may differ substantially between these subjects.

Semantic continuity is therefore an obligation to preserve and explain relationships through change. It is not a claim that all software meaning can be mechanically captured. Public review scope.

3. Information that actually supported the work

A proposal is conditioned by what its author received. A search result, document, policy instruction, test definition, or environmental observation can influence the work even when it is not mentioned in the final explanation. CSSL asks for a record of governed information delivery so that checking can distinguish the information actually supplied from an author's later account of it.

This has a deliberately bounded meaning. A delivery record can support a claim that particular content was served under particular conditions. It cannot prove that a model understood the content, reasoned correctly from it, or was uninfluenced by its prior training. Nor does it establish that no information arrived through a channel outside the governed environment. The trust boundary must identify which channels are covered.

A coherent view of information is also different from a collection of individually authentic observations. Two authentic responses can describe states that never coexisted. A proposal may depend on both an interface and its consumer, or on a policy and the authority derived from it. If those facts are combined across an incompatible change, checking them separately does not establish that the proposal had a valid basis.

Search illustrates the issue sharply. Approximate retrieval can find useful material without establishing completeness. Returning no relevant result from such a search does not establish that no relevant result exists. An absence claim needs a search whose declared scope and completeness support that claim. Even then, the claim applies to that scope and time. A later addition or a change in the search definition can make the earlier conclusion stale.

This account changes the role of refresh and recovery. When supporting information changes, the proposed system would identify work that needs reconsideration while retaining useful history and unaffected activity. The intended benefit is more precise renewal of evidence. Whether it reduces rework at an acceptable cost remains an empirical question. Public review scope.

4. Coordinated change and accountable authority

Some changes only make sense together. An interface and its consumers may require a coordinated revision. A new policy and the work permitted under it may need a consistent interpretation. Showing only part of such a change can expose an intermediate state that satisfies none of the intended conditions.

CSSL treats coherent visibility as a separate problem from collecting approvals or storing drafts. Preparing a change does not mean that it has become authoritative. Evidence that participants are ready does not mean that ordinary readers may treat a partial result as the completed change.

4.1 An observable criterion for coordinated change

For a declared atomic change, the affected state is the complete set of participating revisions and their dependencies: for example, an interface, its implementation, the applicable policy, and the data representation on which both operate. The relevant readers are ordinary governed reads and the tools, builds, and decisions that rely on those reads. Diagnostic inspection of pending work is a separate observation mode; it cannot supply ordinary decision premises.

The intended invariant is **all-or-none visibility of the declared participant changes at an activation-aware, causally closed view**. A view including the change must include, or be able to obtain, every participating post-state and the dependencies supporting its activation. A view preceding it excludes those post-states. A replica that lacks a required revision may obtain the certified pending state or wait; it may not substitute an incompatible local revision. This is a visibility requirement. It does not itself prove serializability of arbitrary transactions, real deployment success, or soundness of every participating artifact.

Suppose a change declares that interface I1, implementation M1, policy P1, and data representation S1 replace I0, M0, P0, and S0 together. Assume these revisions depend on one another, so no cross-version compatibility has been established.

Reader's observation	Status under that declared change
I0, M0, P0, S0 in a view preceding activation	Permitted as the coherent old view. It cannot be relabeled as current authority after the governing policy changes.
I1, M1, P1, S1 with activation and every required dependency included or obtainable	Permitted as the coherent new view. Actual deployment still needs separate observation.
I1 with M0, or P1 with S0, as an ordinary view of this atomic change	Forbidden: part of the declared change has become visible without the rest.
I1, M1, P1, S1 shown after preparation or application but before activation	Forbidden as ordinary authoritative state; preparation and pending application are coordination evidence.
Pending M1 inspected in a diagnostic view while ordinary reads retain the old view	Permitted as explicitly pending evidence; it cannot establish completion or justify a dependent action.
New activation known, but S1 temporarily unavailable	Wait or recover the required state; do not return S0 alongside the new participants.

This example assumes all four revisions are participants. A separately declared compatible change may leave a component unchanged, and a declared staged change involving an independently controlled external effect has different permitted intermediate observations. Neither may be substituted silently for an atomic declaration. Concurrent work must still satisfy its own conflict and causal conditions; all-or-none visibility does not make conflicting edits compatible.

Interruption before a final decision does not authorize exposing prepared work. After a final decision to commit, the outstanding duty is to recover missing participant state and complete activation. Until that is possible, affected ordinary visibility remains coherent rather than exposing a completed subset. A timeout is not evidence that the decision never happened. Later revocation or policy change is assessed under its current conditions; a historical coherent view cannot renew expired authority.

Authority must also belong to the correct actor and moment. An author can describe an intention and propose a change. That author cannot establish future observations merely by including them in the proposal. An execution service can report what it observed during execution. A review decision can establish that an authorized reviewer approved a particular scope, subject to its conditions. These facts are related without being interchangeable.

The distinction between intent and authority remains central. A desired outcome can justify investigation and planning while still requiring additional authority for a consequential action. An estimate of the work's scope may prove too narrow. Expanding it then requires an appropriate scope decision; the original estimate is not a permanent exemption from review. Draft work and exploratory alternatives can remain useful without acquiring authority over a shared system.

The research question is whether an implementation can enforce these permitted and forbidden observations through concurrency, interruption, and policy change at an acceptable operational cost. Establishing that result requires the activation construction and its conformance evidence discussed in §9. The invariant states the intended evaluation criterion. Public review scope.

5. Evidence supports particular claims

Evidence has a subject, a method, a scope, and conditions of applicability. A test supports claims about the tested cases and environment. A formal argument supports a stated property under its assumptions. A compiler result establishes something about a particular translation. A runtime observation concerns an observed deployment and interval. A human approval expresses an authorized decision. These forms of evidence do not become substitutes for one another because they are collected in one system.

CSSL accordingly resists reducing assurance to a single universal score. Stronger evidence for one claim can remain irrelevant to another. A proof about an abstract model does not establish that a deployed binary implements the model. Reproducible output does not establish that the output is desirable. A monitoring interval without an alert does not establish the absence of a fault outside the monitor's coverage.

The same principle applies to provenance and independence. A signature can identify a signer and bind a statement without making the statement correct. Many agreeing results can share a common implementation, training source, or mistaken assumption. Their number alone does not establish independence. Nested or repeated attestations should not increase the underlying assurance merely by adding layers of presentation.

Currentness introduces another dimension. Evidence may remain a valid historical account while losing its ability to support a present decision. A dependency change, revoked

authority, withdrawn assumption, or newly discovered limitation may require affected decisions to be evaluated again. Erasing the old result would damage the history. Continuing to treat it as current would damage the decision.

This suggests a disciplined form of accountability: preserve what was known and accepted at the time, and separately establish what can support action now. The evidence statement applies this distinction to the publication itself. Public review scope.

6. Effects, privacy, and replay

Permission to perform an action and permission to use particular information are separate constraints. An authorized process can still be prohibited from sending a document to a particular destination or using it for an unrelated purpose. Information restrictions can apply to derived artifacts, search indexes, model inputs, logs, and metadata as well as obvious source files.

CSSL places these concerns alongside execution because data handling can be consequential even when no business transaction occurs. The fact that two tenants possess identical content does not authorize sharing their derived records. The fact that a model can process a document does not establish that the document may be sent to that model. Any release of information needs support from the relevant authority and purpose.

Execution raises a different uncertainty: what did the outside world actually do? A request may be accepted locally while its remote effect remains unknown. A response can disappear after the effect occurred. Repeating the request can be safe for some destinations and duplicative for others. No general promise of exactly one external effect follows merely from keeping a careful local history.

Recovery must therefore depend on the destination's actual guarantees and on available observations. Some actions can be reconciled, some can be safely retried under their conditions, and some require a decision about unresolved consequences. Compensation is another action with its own authority and outcome. It does not erase the original event or necessarily restore the earlier world.

Replay has a narrower purpose. Reconstructing a history or supplying retained observations to an isolated evaluation must not itself repeat live external actions. Moreover, a preserved commitment to historical content does not guarantee that the content remains available. Privacy restrictions, retention limits, and erasure can reduce what can be replayed while leaving some historical evidence intact. A deletion claim must state its covered stores and known destinations; it cannot establish control over every unknown copy or a model's learned parameters. Public review scope.

7. From changed source to observed operation

The source extends semantic continuity into builds, artifacts, releases, deployments, and incidents. This breadth follows from the underlying problem. A reviewable source change does not identify the artifact that reached an environment. A reproducible build does not establish that its release was authorized. A deployment command does not establish the resulting behavior.

The conceptual objective is to preserve the relationships between these stages while keeping their claims distinct. A build record should describe what was built from which declared inputs. An artifact record should identify the resulting subject. A release decision should apply to the intended source and artifacts under current conditions. Operational observations should identify the environment and interval they actually cover. Repair should connect to the incident and the conditions it is intended to restore.

Resource limits belong in this account. Expensive validation can delay work, and an incomplete search or interrupted computation can remain useful. But exhaustion of a budget cannot turn an incomplete result into complete evidence. Work needed for recovery and control also competes for resources; a design that spends everything on new activity can undermine its ability to govern that activity.

Offline work presents a related boundary. A participant may continue drafting and collecting bounded evidence while disconnected. That activity cannot silently acquire current authority over an environment it has not observed. Reconnection requires reconciliation with the current system and its policies. Similarly, rebuilding a derived view should recover the meaning of recorded events, rather than introduce new authority through the view service itself.

These are specification obligations. Their feasibility, responsiveness, and operational cost require implementation and measurement. Public review scope.

8. Four illustrative cases

Four hypothetical cases show what the proposed account of change would need to preserve. Their outcomes are illustrative.

A coordinated change succeeds. A team changes a reporting interface and the consumers that rely on it. The work begins with an identified view of the relevant interface, consumers, policies, and test requirements. Review concerns the proposed scope and its supporting evidence. The intended successful outcome is a coherent revision whose supporting records distinguish the author, review decision, resulting artifact, and later operational observations. A passing test remains evidence for the test's scope. It does not stand in for authorization or prove that the revised service has reached production.

An unsupported absence claim is denied. An assistant searches for callers of a component and receives no result from an approximate search. It proposes removal on the premise that no callers exist. The empty result does not support that premise. The objective can continue through a suitably scoped exact investigation or through a revised change that does not depend on the unsupported claim. The meaningful result is completion of authorized work with an adequate basis, rather than merely recording that an initial proposal was rejected.

An interrupted external action is reconciled. A release action sends an artifact to a destination, but the response is lost. The local process does not know whether publication occurred. Recovery examines the destination's observable state and the conditions of the original action. If reliable evidence establishes completion, the record can be settled accordingly. If the destination supports an appropriate safe retry, a retry may be possible

under current authority. If neither is established, uncertainty remains a real constraint. Replaying the historical record does not issue another publication request.

Later evidence changes a present decision. A dependency analysis supported a release, and a later discovery narrows that analysis's valid scope. The original release record remains part of history. A present deployment decision must consider whether its required assurance is still satisfied. Affected work can receive renewed analysis, additional evidence, or a revised scope. The new evidence does not make the earlier record disappear, and the existence of the earlier record does not settle the new decision.

Together these cases show why continuity requires more than append-only storage. It requires an account of which relationships continue to support which decisions as circumstances change.

9. Threat assumptions, limits, and research questions

The proposal assumes that authors and external services can be mistaken or adversarial. Signed statements require checking; independent observation matters; authority and information flow need enforceable boundaries. Its guarantees would also depend on the trustworthiness of the components that capture information, interpret meaning, enforce policy, and record consequences. Unobserved channels, compromised enforcement, unavailable evidence, and failures beyond the assumed trust model remain material limitations.

The Sovereign-Closure revision makes three construction requirements central to its foundational claims. Each has a distinct obligation and a bounded example against which a proposed construction can be evaluated.

Foundation	Construction requirement and bounded example	Claims that remain conditional
Binding a proposal to its premises and their discharge evidence	In the nonempty-discharge path, the proposal, premise account, and discharge certificate require one another's final content-derived identities. For a proposal requiring one discharged premise, an otherwise valid certificate cannot simply be attached after the proposal's final identity is chosen without revisiting what that identity binds. A qualifying construction must resolve this dependency in a form agreed by producer and verifier.	Canonical encoding and implementation conformance for that path; complete evidence binding for consequential proposals. An empty discharge list avoids this particular cycle without proving the rest of CSSL.
Binding activation to its supporting frontier	The activation identity depends on a frontier that includes that activation. For a two-participant change, having both post-states ready does not by itself establish the activation/frontier relationship. A qualifying construction must distinguish the commitment before activation from the resulting history after it.	Constructible activation records, activation-aware views, and implementation conformance for atomic cross-domain visibility. The all-or-none criterion in §4 remains the requirement to test.

Foundation	Construction requirement and bounded example	Claims that remain conditional
Complete rendering of accepted ordered placements	The retained traversal fragment renders concurrent blocks $A=[a]$ and $B=[b]$ as $[a,b]$, but can omit a valid later insertion $C=[c]$ whose origins are the adjacent a and b . Its checked scope supplies a detector for this omission. A general rendering argument must also cover supported histories beyond the fragment, including moves and tombstones.	Complete deterministic rendering and the general convergence and noninterleaving claims for supported event histories. A projection must reject missing, repeated, or unexpected placements rather than expose a result that omits an accepted placement.

The distinctions among stable identity, content, historical evidence, current authority, and ordinary visibility state the intended architecture. The source's final self-audit classifies the associated foundational statements as requirements without proof. Their stronger construction and conformance claims therefore depend on discharging the obligations above. Formal models and acceptance cases frame that evaluation; neither their presence nor a closure label establishes successful implementation. [Public review scope](#).

Several research questions follow. Can semantic interpretations be precise enough to reduce unnecessary invalidation while remaining conservative about unknown effects? How can the cost of complete dependency checking be bounded without weakening the claim? Which combinations of evidence provide useful assurance for a particular decision, and how should common causes be represented? How can coherent change remain available under realistic interruption and trust failures? What proportion of external uncertainty can be resolved by observation rather than repetition? How should privacy-driven unavailability change the claims a historical record can still support?

The proposed composition is valuable as a research agenda because these questions constrain one another. Faster change is useful only if the resulting evidence remains adequate. More retained information can aid reconstruction while increasing privacy obligations. Finer semantic distinctions can reduce unnecessary rework while making the interpretation itself more demanding to validate. CSSL's conceptual position is that these tradeoffs should be explicit in the system's account of change.

Source guide

This successor guide points to public arguments and their limits. The new private evidence linkage remains excluded. Previously public conceptual source guides, including any retained titles, document ordinals, sections and chapter maps, are not claimed to have been removed; see the [correction scope](#). The historical edition remains available unchanged.

Foundations. Review the relevant argument and assumptions in this paper, the [evidence guide](#), and the [research review route](#). Specification, reported results and production observations have distinct scope.

Identity and interpretation. Review the relevant argument and assumptions in this paper, the [evidence guide](#), and the [research review route](#). Specification, reported results and production observations have distinct scope.

Information and currentness. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

Coordinated work. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

Assurance. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

Effects and information governance. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

Operational continuity. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

No source-guide entry establishes a full implementation, a complete historical migration proof, or an independent audit.