

GOVERNED AUTONOMY

Architecture of Governed Autonomous Systems

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

Architecture of Governed Autonomous Systems 3

Abstract	3
1. A continuing objective has several kinds of state	3
2. Where the source works concentrate	4
3. Four kinds of relationship	6
4. Specified relationships and their limits	7
5. Interpretation, permission, and effect	9
6. State and evidence are about something	10
7. Interruption does not end responsibility	10
8. Improvement without inherited authority	11
9. Application understanding and continuity of work	12
10. Observation and stewardship are consequential interfaces	13
11. Three conceptual journeys	14
12. What composition must establish	15
13. Evidence and the next research claims	16
Source guide and scope of the synthesis	17

Architecture of Governed Autonomous Systems

Governed Autonomy · Publication 2026.10.7

Lead architecture synthesis: responsibilities, source relationships, and composition questions

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

Abstract

Persistent autonomous work joins several problems that are often evaluated separately: forming an objective, interpreting it, deciding what is permitted, changing the world, establishing what happened, and carrying responsibility through interruption and change. Each has its own subject and conditions for success. An institution needs the resulting decisions to remain connected when a model changes, a permission expires, or an external action leaves an uncertain result.

This paper presents a framework for reviewing those connections across the Governed Autonomy collection. At each consequential boundary, the framework identifies the subject, responsible owner, authority, evidence, and continuing obligation. It compares the works' different responsibilities, distinguishes expressly specified relationships from proposed combinations, and develops three conceptual journeys that make the obligations of composition concrete.

The synthesis draws on specifications, theoretical arguments, bounded reported reference results, and the retained 10 October 2026 author-directed, model-assisted examination of the implemented Freehold engine and Cuttlefish host. That historical source and dated AWS deployment evidence establish an implemented relationship between those two systems. The framework does not establish that every research family forms one integrated deployment. The final sections identify the additional implementation, evaluation, and adoption evidence needed for stronger claims. Proprietary construction details are outside the paper's scope.

1. A continuing objective has several kinds of state

An autonomous system is often described through its most visible activity: reasoning about a request and selecting a tool. The description becomes incomplete when work lasts longer than one interaction. The objective may persist while the available model changes.

Permission may expire while a plan is being prepared. A service may complete an action just before a connection is lost. A later worker may inherit a concise summary that omits the uncertainty. The visible conversation can continue while the responsible account of the work has already broken.

The collection's institutional thesis treats persistence as continuity of identity, obligations, authority, evidence, and consequence. These subjects are related but cannot be collapsed into one conversational state. A remembered instruction does not necessarily remain an

operative permission. A remembered result does not necessarily remain current evidence. Replacing a component does not discharge the institution's outstanding duties.

The architectural question is consequently broader than how to retain context. It asks which responsibilities must survive a transition and what would establish that they did. This can apply to a delegated research task, an operating product, or a proposed institution that forms its own internal agenda. Those modes have different origins of intent. They still need an account of consequential action that does not derive authority from a model's confidence or from the persistence of its narrative.

Five questions provide a common reading route:

CONCEPTUAL RESPONSIBILITY MAP

Purpose

What future condition is being pursued?
Whose objective or internally formed mission is it?

Authority

Who may act?
On what terms?
Who can change them?

Consequence

What may change?
Who is affected?
What remains owed?

Evidence

What was observed?
What is established?
What remains uncertain?

Continuity

What must survive?
How does work resume?
How can claims be revised?

Each responsibility has its own assumptions and evidence.

The map identifies responsibilities. An implementation may assign several to one component or distribute them across components. Either arrangement must preserve the distinctions among the decisions and their supporting assumptions.

2. Where the source works concentrate

The works share an interest in governed autonomy while addressing different objects. A comparison is useful when it explains the object of a claim before comparing the mechanisms proposed to support it.

Work	Main question developed in this collection
Institutional doctrine	What makes a persistent autonomous institution accountable across action, change, and succession?
MIRAGE	How can consequential computation remain inside a declared authority and observation boundary?
CSSL	How can semantic state changes retain their premises, authority, and causal history?
IC/POAR	How can intended product behavior remain identifiable through construction and operation?
AD8	How can responsibility remain connected across software change, external action, and learning?
Freehold	How can internally originated institutional activity coexist with bounded external stewardship?
Cuttlebone	How can reasoning, recommendation, and claims of completed effects retain their distinct authority?
SEPION	What can agreement among generated artifacts establish under declared evaluation conditions?
RCFS	How can competence and implementation change without inheriting new authority or losing obligations?
Cuttlefish Semantic OS	How can application understanding persist outside an agent's immediate reasoning while preserving current authority and verified results?
Cuttlefish Operating Environment	How can a user's work remain coherent across interfaces, resources, execution, and collaboration?
CHROMATIC	How can ambient maintenance, direct intent, and learned assistance remain useful under continuing human authority?
Experience Economy and ISE	When can accumulated experience support reusable knowledge or value?
The Reckoning	How can claims and judgments remain accountable to later evidence and calibration?
Human-LLM Capability Elicitation	How does an interaction help reveal a model's usable capability, and what limits the resulting claim?

This is a map of emphasis. It is not a decomposition into interchangeable modules. AD8 and RCFS both discuss learning, for example, while placing it in different accounts of responsibility. CSSL and the Semantic OS both concern meaning and state, but that overlap does not establish identical objects, authorities, or transition rules. A narrower work can supply an important distinction without being a mandatory subsystem of every broader proposal.

The topic papers are the appropriate starting points for their individual definitions and evidence limits. The glossary assists comparison while preserving scoped meanings. A familiar word should prompt a question about its subject, not an assumption of equivalence.

3. Four kinds of relationship

A diagram can make a collection look more integrated than its sources establish. Arrows can imply compatibility, precedence, or deployed communication that has never been specified. This synthesis therefore classifies a relationship before drawing conclusions from it.

READING RELATIONSHIPS PRECISELY

Explicit source contract

A named, scoped contract joins two works.
The contract still needs its stated proof or tests.

Conceptual correspondence

The works address related questions.
Shared vocabulary does not establish equivalence.

Proposed integration

A journey illustrates how roles might connect.
Additional contracts and evidence are required.

Unresolved adoption

Ownership or precedence has not been selected.
The public edition does not silently decide it.

Relationship class and implementation evidence are separate.

An **explicit source contract** is a relationship expressly defined by an identified source. The contract establishes what that source requires within its scope. It does not establish that its requirements have been implemented or verified. The doctrine-MIRAGE cross-corpus contract is an example: it distinguishes the institutional requirement, a mechanism within a declared effect scope, an intended realization, and bounded reference semantics. It records obligations connecting them rather than completed proofs of the entire chain.

A **conceptual correspondence** identifies a related question or responsibility. CSSL's concern with adequate and current decision premises corresponds to AD8's treatment of concurrent work. That comparison can expose common evaluation questions without

establishing that their objects can be exchanged or their permissions combined. Correspondence supports analysis; it does not confer compatibility.

A **proposed integration** describes a combination that might be useful. The journeys later in this paper are proposals of that kind. They show what responsibilities would need to remain connected. They are neither deployment instructions nor reports that the named works already operate together.

Unresolved adoption identifies an open decision about ownership, precedence, compatibility, or acceptance. If two proposals both describe a source of authoritative state, composition requires a decision about which subject each owns and how their relationship is governed. Until that decision is made, the comparison remains an account of alternatives.

The distinction is independent of assurance. A clearly specified relationship may remain entirely unimplemented. A reference implementation may demonstrate a useful local behavior while leaving integration unresolved. No relationship label promotes a source claim into a measured result.

4. Specified relationships and their limits

The following comparisons identify the source and scope of each relationship.

Relationship	Class	Source basis and practical limit
Institutional doctrine, MIRAGE, Native, and the P0 reference	Explicit source contract	The Cross-Corpus Contract, sections on refinement, state vocabularies, terms, and assurance, specifies scoped relationships and remaining obligations. It expressly denies that the bounded reference implements the whole institution.
Freehold and its proposed Cuttlefish host	Explicit source contract	The Freehold Implementation Contract, §0, assigns internal institutional activity and external operating authority distinct owners. The contract is proposed; its historical implementation comparisons lack a version-specific verification basis.
Freehold, Observatory, Atlas, Exchange, and stewardship	Explicit source contract	The Freehold supporting doctrine assigns these surfaces to the same institutional account while distinguishing observation, publication, relations, and human custody. A visual representation is not a new source of authority.
Freehold and the Experience Economy synthesis	Explicit source contract	The synthesis's opening dependency statement and §§1–2 place it under identified institutional doctrines and define a proposing role for synthesis. Its evidence statement calls the experience system an implementation target.

Relationship	Class	Source basis and practical limit
Freehold and The Reckoning	Explicit source contract	The Reckoning §§0.1–0.2 and 13–17 place its calibration account alongside identified institutional owners. Its receipt and revocation relationships still require live verification, and its training-weighting proposal is not a measured result.
Cuttlefish operating architecture and collaboration master	Explicit source contract	The collaboration master §3.1 and Appendix A identify adopted design domains from the operating architecture. The sources retain adoption and verification obligations; this is documented design lineage, not observed product integration.
CHROMATIC and its intended native owners	Proposed integration	CHROMATIC §44 proposes native integration; §§5.2 and 28–29 preserve other systems' ownership. Its dated alignment statements and planned checks do not establish current implementation.
CSSL, IC/POAR, and AD8	Conceptual correspondence	Their papers address semantic state, lasting product intent, current premises, and evidence of operation. The comparison does not adopt a universal shared state model or precedence rule.
AD8 and SEPION	Conceptual correspondence	Their generated-artifact discussions distinguish agreement, compliance, and the continued applicability of evidence. Shared concerns do not transfer a qualification result from one proposal to another.
Institutional doctrine and RCFS	Conceptual correspondence	Both distinguish continuing responsibility from the model or method currently used. RCFS §36.4 does not incorporate unrelated portfolio mechanisms into its constitution without a separate contract.
Cuttlebone and broader execution proposals	Conceptual correspondence	Cuttlebone concentrates on the consequence attributed to an artifact, its information origins, and its audience. Its conformance program still depends on actual execution and observation boundaries.
Human-LLM Capability Elicitation and the systems papers	Conceptual correspondence	The elicitation corpus supplies a methodology for attributing validated performance. Corpus 2 §2.2 and Corpus 18 §18.7 explicitly exclude a product or agent implementation specification.

Relationship	Class	Source basis and practical limit
A single combined implementation of the whole collection	Unresolved adoption	No common integration or authority arrangement has been established across the collection. Each proposed connection requires an explicit contract and an evaluation of the resulting responsibilities.

The dates and system versions remain separate. Semantic OS 3.1.0 is not a new version of the doctrine-MIRAGE package. A later date can identify a later source without establishing authority over another work. The public research map provides the public review route for each paper.

5. Interpretation, permission, and effect

Three questions recur across the collection: what is intended, what is permitted, and what actually happens. Each can have a well-formed answer while the others remain unsettled.

IC/POAR treats intended behavior as something that should remain identifiable through different implementations. This creates an interpretation obligation. A request to make a process easier does not automatically select a particular vendor, authorize a wider audience, or approve new uses of data. A system may infer necessary work, but it must distinguish necessity from preference and uncertain interpretation from accepted intent.

Cuttlebone examines a related boundary in the artifacts produced during work. A rehearsal may use genuine observations and remain a rehearsal. A recommendation may be externally visible and remain unexecuted. Repetition, persuasive wording, or reuse in another context cannot supply the missing authority. The information's origin, its audience, and the consequence attributed to it require separate treatment.

MIRAGE concentrates on the movement from a proposal toward an authorized effect within a declared boundary. Preparation is not permission, a decision is not completed execution, and observation and settlement remain additional responsibilities. This distinction prevents an authorization record from becoming a substitute for an actual outcome. It also prevents a detailed record of an outcome from retroactively creating permission for it.

The comparison suggests a useful composition question: at every transition, what new fact is being established, and which owner is entitled to establish it? Interpreting a request establishes a candidate meaning. Granting permission establishes an allowed scope. Executing changes a system or the world. Observing establishes evidence about that change. Settling discharges particular obligations under applicable rules. These may occur close together in an implementation, but a correct explanation should not merge their meanings.

The distinctions extend to apparently passive operations. Reading sensitive information can itself be consequential. An observation may reveal information to an unauthorized audience. A maintenance operation may affect the constraints governing later work. The relevant boundary follows consequences and authority, not a convenient label such as background activity.

6. State and evidence are about something

The collection uses words such as state, ledger, identity, receipt, and evidence in several contexts. Their subjects matter. A record of a software artifact is not a record of its deployment. A resource counter is not a complete account of an institution's solvency. A release history is not the same thing as adjudicated knowledge. Matching names cannot establish a shared semantic type.

CSSL's treatment of semantic change highlights two questions about a decision's basis. Was the basis adequate for the decision, and does it remain applicable? A worker can read every object it knows about while missing a newly eligible object or a changed rule defining the relevant group. Freshness of known data does not establish completeness of the premises. Conversely, an adequate basis can become stale while work continues.

This problem reaches beyond storage. Evidence about an artifact may depend on an evaluator, an input domain, a configuration, and a policy. A change to any relevant dependency can alter what the evidence supports. Historical records remain useful, but their continued use needs a reasoned account of applicability.

Provenance is one part of that account. W3C PROV-DM supplies a general vocabulary for entities, activities, and agents in provenance descriptions; factual correctness still requires appraisal. The IETF RATS architecture similarly distinguishes evidence, appraisal under policy, and the relying party's decision. These established distinctions provide comparison points for the collection's proposals. See the bibliography entries for PROV-DM and RATS.

The architecture must also distinguish a claim's authenticity from its meaning. A signature can help establish who made a statement under a key relationship. It does not establish that the signer observed the event, that the observation was accurate, or that the signer had power to authorize a different action. Multiple signers do not necessarily supply independent sources. Common ownership, instruments, training, or incentives can preserve a common failure beneath an apparently diverse record.

7. Interruption does not end responsibility

Recovery is a central composition test because it crosses the boundary between an internal attempt and an external objective. A process can restart while an earlier payment remains uncertain. A tool can report an error after a remote service completed the requested change. A model can be replaced while a duty to deliver or compensate remains open.

AD8 and RCFS both keep continuing responsibility distinct from the current method. The doctrine extends the issue to institutional identity, amendment, succession, and dissolution. These treatments do not imply one shared implementation, but they expose a common failure mode: a new local state can appear clean while the institution still owes something because of its earlier activity.

A recoverable failure calls for an authorized continuation. Depending on the actual conditions, this may involve obtaining current evidence, repairing an owned resource, choosing another qualified method, or reconciling an external outcome. It does not imply that every operation is safely repeatable. Unknown external effects and uncertain resource consumption remain relevant until there is an adequate basis for resolving them.

Compensation is another action, not an inverse operation on history. A refund can restore money while leaving delay, lost opportunity, notification, or a third party's response unchanged. Any claim of restoration needs to identify what was restored and what remains. A rollback of software likewise does not establish a rollback of data or of the outside world.

Continuity therefore requires an observation about the resumed objective: has the authorized outcome actually been completed, with its obligations accounted for? A record that recovery was attempted is useful evidence of an attempt. It is insufficient evidence of the outcome. A genuine need for external information or human authority can remain, but the system should not treat its own recoverable interruption as the natural end of achievable work.

8. Improvement without inherited authority

A system that learns from experience introduces a second form of change. It may retain new knowledge, choose a more efficient method, or train and deploy a different model. These changes can improve practical competence while creating new uncertainty about scope, provenance, privacy, and continuity.

RCFS makes the separation between competence and authority central. With the governing permission context held fixed, a better method does not acquire broader permission merely because it can now accomplish more. The reverse distinction also matters: a safety boundary can hold while a replacement method becomes less competent. Protection under substitution is not proof of performance under substitution.

AD8 distinguishes operational use, retained experience, learning, model evaluation, release, deployment, and selection. These activities have different affected parties and different authority requirements. Permission to use information for a current task cannot be presumed to authorize its incorporation into a shared model. Possession of model bytes cannot itself authorize their deployment.

SEPION examines another improvement claim: whether repeated generation supports a dependable behavioral result. Agreement and compliance remain different questions. Correlated generators can agree on an error; apparently different artifacts can satisfy the same bounded requirements. New instruments may distinguish behaviors that earlier instruments grouped together. Stronger evaluation therefore needs a precise account of what changed in the claim, rather than an unexplained comparison of scores.

The Experience Economy and ISE study distinguishes a history from a useful lesson and a useful lesson from demonstrated transfer. A procedure can work in its original setting while depending on conditions absent elsewhere. Retained knowledge can survive a model change while its interpretation becomes less reliable. Reuse therefore creates additional claims about applicability and permission rather than automatically carrying the original success forward.

CHROMATIC brings this issue into recurring human work. A repeated behavior can suggest a useful ability, but repetition is incomplete evidence of preference. A person can be following a temporary constraint or making an exception. Evaluating an ability can support competence within a domain while leaving permission and current user intent separate. The

shared question is whether improvement reduces work for the person without quietly expanding control over their environment.

The Reckoning adds the quality of probabilistic judgment. A system can accurately record its actions while systematically overstating their prospects. Evaluating its expectations requires an adequate population, prospective records, meaningful outcomes, and an account of changes in the process producing the forecasts. Calibration differs from task success, decision value, and permission. The source proposes that calibration can constrain existing scope; it cannot create authority.

Human-LLM Capability Elicitation studies the wider work episode: participants, information, representations, resources, authority, and validation. Its narrow elicitation claim holds the model and relevant support fixed while changing the interaction. Adding a source, tool, computational budget, or expert can also improve the outcome, but establishes a different kind of contribution. A useful system should record those contributions instead of attributing every gain to the model or the prompt.

The study is methodological and provisional, explicitly excluding a product or agent implementation specification. It helps ask whether a representation preserves the intended problem, whether additional candidates bring independent evidence, and whether a validator can actually change a claim's status. Elaborate procedure is not itself evidence of improved capability. A successful result may depend on human assistance or a favorable environment, while a remembered failure may later be corrected. Neither should become a permanent universal rule through repetition or compression.

9. Application understanding and continuity of work

The Semantic OS source asks whether knowledge of applications can persist outside the immediate reasoning context of a model. A task may require substantive judgment about a document while repeatedly rediscovering navigation, object identity, and the consequences of a particular edit. Preserving supported application understanding could reduce that repeated work. The claim concerns where understanding resides and how it remains valid, rather than eliminating the information needed for the actual decision.

Persistent understanding cannot be treated as permanent certainty. An application can change behavior without changing its familiar appearance. A known operation can have different consequences in another account or sharing arrangement. A compact description is useful only to the extent that its evidence and scope remain adequate. The source therefore makes currentness, exact answers, authority, and observed effects central to its ambition.

The operating environment sources ask a complementary question: can the user's work persist across applications, participants, models, views, and interruptions? The objective should remain identifiable while these means change. A conversation ending does not necessarily end the work, and an application view closing does not necessarily cancel an obligation. Conversely, continuing responsibility cannot extend permission indefinitely.

Collaboration adds boundaries that a single shared conversation can conceal. A person may contribute a permitted finding without giving colleagues access to the private environment in which it was developed. A customer may review a proposal without acquiring the right to

inspect its supporting internal workbook. Presentation, participation, and authority are different relationships to the same work.

The combined research question is how the system preserves the actual subject of work. A retrieved copy is not automatically the live document the person intended. A visible application is not proof that its state changed. A shared result is not a grant to operate the producing application. A summary of an earlier action does not establish whether an uncertain external effect occurred.

These correspondences do not establish that the Semantic OS and operating environment share one implemented contract. The later collaboration master explicitly identifies design dependencies on the operating architecture; the broader comparison with Semantic OS remains conceptual here. A future integration would need to preserve the source of application facts, each participant's authority, and the current evidence supporting every reported result.

The positive ambition matters alongside the boundaries. Refusing an unsupported action can preserve scope, but it does not demonstrate useful work in ordinary applications. A credible study would show that the intended result can actually be reached, including when applications were not designed for agent use, while measuring the costs of understanding, coordination, recovery, and human intervention.

10. Observation and stewardship are consequential interfaces

Freehold asks whether an institution can originate internal work while a human maintains its external operating conditions. This differs from the ordinary delegated-task mode in which the user supplies the objective. The distinction changes the meaning of the interface: an observation surface must not quietly become a route for assigning new internal missions.

The Observatory is intended to make activity and evidence intelligible. The World Atlas adds selected public claims about institutional history. Exchange concerns relations initiated by institutions, while stewardship concerns the human's bounded custody of conditions such as resources, access, and continued operation. These responsibilities belong in a substantive account of Freehold, rather than being reduced to decorations around a reasoning engine.

Each surface can introduce a new consequence. A public history can disclose sensitive information or overstate independence. A ranking can affect incentives. A funding decision can shape feasible activity without taking the form of a direct instruction. Pausing or ending operation can affect unfinished commitments. A claim of autonomy must therefore state the influence it excludes and the influence that remains possible.

This is also a useful test for user-facing architecture more generally. A control must reach its real owner. A success display must correspond to its claimed result. A retrieved status must be distinguished from a fresh observation when the difference matters. A polished interface is part of the user's evidence about the system; misleading status can undermine governance even when backend records are technically correct.

The collection does not infer legal personhood, sovereignty, or independent moral standing from its institutional terminology. Founding conditions, external obligations, and

stewardship still depend on applicable human authority and actual control. A distribution of keys or roles is not proof that one beneficial controller cannot dominate the system.

11. Three conceptual journeys

Three hypothetical journeys illustrate how responsibilities would need to connect. Their outcomes explain proposed requirements rather than report observed implementations.

11.1 A requested analysis reaches an authorized result

A user requests a comparison of approved records and an internal report. The intended outcome includes accurate scope, appropriate source use, and a report the user can retrieve. A candidate interpretation identifies the required work without expanding the audience or data rights. The system uses currently authorized records, preserves material uncertainty, and produces a result whose claims are tied to their supporting observations.

The useful outcome is the report. It does not need to cause a further external action to count as completed work. If a later suggestion proposes sending the report to another organization, that suggestion raises a new disclosure and authority question. Its usefulness does not enlarge the original request.

Several papers illuminate different parts of this journey: IC/POAR on lasting intended behavior, CSSL on decision premises, Cuttlebone on the distinction between analysis and consequential claims, and MIRAGE on declared effect boundaries. Describing those contributions together identifies a research agenda for their interfaces. It does not establish that one artifact format or implementation already connects them.

11.2 An interruption is resolved without duplicating an effect

A permitted transaction is submitted, but its response is lost. The objective remains open. A replacement worker receives the continuing responsibility, including the uncertainty about what happened. It does not infer non-occurrence from the absence of a local success message.

An authorized observation establishes that the destination completed the transaction. The system reconciles its account, discharges the appropriate obligation, and returns the supported result to the user. If evidence instead established that the transaction did not occur, a permitted continuation could be considered. If evidence remained unavailable, the precise external dependency would remain open.

This journey tests continuity across model, process, record, and external service. Its success depends on the destination's actual behavior and available evidence, as well as the internal architecture. AD8, RCFS, MIRAGE, and the doctrine each supply relevant requirements under different scopes. A completed composition study would have to show how those scopes meet at the actual effect boundary.

11.3 A more efficient method remains a qualified substitute

A repeated task becomes a candidate for a specialized procedure. Evaluation establishes a bounded domain in which the proposed replacement satisfies the required behavior. The service also examines protected exceptions, continuing obligations, data rights, and the conditions under which an alternative method remains necessary.

The replacement is used only where the qualification remains applicable. Changed evidence later reveals an exception outside the earlier comparison. The system limits the replacement's scope, preserves ongoing work, and returns affected tasks to a suitable method while further evaluation proceeds. The successful continuation is the user's completed authorized work with an adequate current basis.

RCFS helps frame substitution and qualification; SEPION helps distinguish supported behavioral comparison from universal equivalence; AD8 raises learning and release authority. None makes the other questions disappear. Demonstrating a net efficiency gain would also require measuring evaluation, maintenance, and recovery costs, rather than only the fast path.

12. What composition must establish

Composition can create a failure even when every component appears to satisfy its local description. One component may interpret an identifier as an artifact version while another treats it as a continuing institutional identity. One may report that a request was accepted while another displays that the operation completed. A permission may lose its purpose restriction when translated. An old observation may cross an interface without the conditions under which it remains useful.

The appropriate review follows the actual claim and effect across the boundary. It asks whether the same subject remains identifiable, whether the responsible owner remains authoritative, whether restrictions survive translation, and whether evidence supports the statement made by the consumer. It also asks how rejection, partial application, interruption, and recovery are represented. A correspondence that covers only a successful ordinary path leaves important behavior unaccounted for.

Established work on protection offers relevant principles. Saltzer and Schroeder's account of least privilege and complete mediation helps frame why a checked interface is insufficient when an alternative consequential route remains available. Lamport's account of event ordering helps distinguish causal ordering from assumptions about physical time. Hohfeld's legal analysis provides a reading route for keeping different legal positions distinct rather than treating every power or claim as the same kind of right. These are comparative foundations, not evidence of this collection's correctness or a claim that a technical system can decide legal relations by itself. *Bibliography.*

A source-level refinement argument must also reach the implementation it describes. An abstract rule can forbid an unauthorized action while a concrete administrator, shared credential, network route, or recovery mechanism still permits it. The observed deployment must satisfy the assumptions of the argument. The public exposition withholds implementation details, so it cannot itself supply the evidence needed to certify that closure.

13. Evidence and the next research claims

13.1 The implemented Freehold-Cuttlefish relationship

The retained 10 October 2026 author-directed, model-assisted examination identified an actual Freehold engine and Cuttlefish host operating in AWS. The new AD8 corpus and reference analysis adds no product or infrastructure examination. The implementation evidence records the source revisions declared by examined images, dated workload observations, and the limits of each inspection. It supports implemented responsibilities for institutional history, founding, charter and resource checks, external authority, receipts, writeback, and cognition lifecycle. It does not supply conformance proofs for every proposed architecture in this synthesis.

The continuity exhibit is an additional, public reference result. Its lost acknowledgement, process exit, successor reconciliation, replay, revocation, and evidence checks are executed against a controlled witness. It supplies a reproducible bounded example of continuity; it is not a production acceptance test or a model-replacement experiment. Prior-art comparison places these responsibilities alongside established authorization, provenance, durable execution, and formal-methods work.

13.2 Claim-specific research obligations

The evidence guide distinguishes six kinds of support. A specification establishes what is required or proposed. A reference run establishes behavior on its inputs. Machine checking establishes a property of a stated formal model or domain. Refinement connects identified descriptions under assumptions. Empirical evaluation measures behavior on an acquired workload. Independent audit adds an identified external assessment with its own scope and conflicts. These forms of support do not form an automatic ladder.

The topic papers state the propositions, premises, scope, and attribution of disclosed reference checks and conditional arguments. Their comparison criteria distinguish a permitted observation, a forbidden combination, and a relationship not established by the available evidence. Authored traces remain illustrations, and the collection supplies no acquired workload from which to estimate deployment or escalation rates. Evaluating those rates requires a defined population and actual observations.

Stronger architecture claims require a bounded but complete operating subject. A useful study would name the objective class, allowed effects, external assumptions, baseline, acquisition method, and observable criteria for success. It would include legitimate work, denied expansions of scope, stale evidence, changed authority, interruption, and recovery. It would retain incomplete and censored outcomes rather than silently removing them from the denominator.

A useful continuity study would replace a model or component while duties remain outstanding and establish what actually survives. A useful authority study would inspect every relevant decision and effect route within its declared boundary, including indirect and privileged paths. A useful learning study would separate capability gains from changes in data use and permission. A useful human study would test whether people understand consequential choices better, including whether the interface induces accidental broadening of authority.

These studies would turn the conceptual comparisons into stronger claims about implemented systems. The present contribution is the set of explicit, contestable obligations they would test: preserving the subject of work, its authority, its evidence, and its continuing responsibilities through composition and change.

Source guide and scope of the synthesis

The institutional basis is the definitive third edition of *The Doctrine of Governed Autonomous Institutions at Scale* and its two supporting doctrines, as identified in the thesis. The doctrine-MIRAGE comparison relies on the *Cross-Corpus Contract*, particularly its sections on specification-to-semantics refinement, indexed state vocabularies, scoped terms, assurance, and missing measurements. The MIRAGE paper explains the associated trust and evidence limits.

The Freehold relationship relies on the opening evidence statement and §0 of its *Implementation Contract*, together with the *Observatory, World Atlas, and Stewardship Doctrine*. The Freehold source guide provides the edition and section route. The Cuttlebone source notes identify its narrower consequence-boundary commitments.

The Semantic OS source guide identifies the canonical 3.1.0 architecture and its review companions. The environment source guide records the operating architecture's adoption gates and the collaboration master's explicit dependencies. The CHROMATIC source guide preserves its proposed integration status. The Experience Economy source guide identifies the synthesis's doctrinal dependency and separates intended transfer measurement from observed results.

The Reckoning source guide distinguishes prospective forecast records, statistical evaluation, and proposed institutional consequences. The capability-elicitation source guide identifies the baseline research corpus and its explicit limits on implementation and empirical claims.

The remaining comparisons use the identified sections in each topic paper. In particular, the IC/POAR source guide, AD8 source guide, SEPION source guide, and RCFS source guide identify the basis for conceptual, construction, and conformance claims. The sources and method describe the collection's scope.

Adopting a proposed relationship requires a versioned decision about the subjects and responsible owners, together with evidence for the resulting claims. The independently versioned works retain their obligations until such a change is made through their applicable authority arrangements.