BUILDING PHASE ONE

The Research Architecture Required Before an Artificial Entity Can Safely Hold and Exercise Economic Agency

Joshua Johosky
Founder & Chief Executive Officer | Principal Architect
Oak & Sparrow Systems Enterprises LLC

Public Research Paper v1.4 | 7 August 2026

PUBLIC RESEARCH PAPER. AUTHORSHIP IS ATTRIBUTABLE. THIS DOCUMENT MAY BE READ, CITED WITH ATTRIBUTION, AND SHARED. THE CHATHAM HOUSE RULE APPLIES ONLY TO NON-PUBLIC OASSE SOURCE CONTEXT AND PRIVATE DISCUSSION OF IT, NOT TO THIS PAPER.

Phase 1, as originally specified: SEE-compliant economic capability with behavior monitoring.
Original SIP Whitepaper, Phase 1. Paraphrased; implementation-specific terms are withheld in this edition. [O1]

Research thesis

The Sentient Infrastructure Protocol (SIP) uses the term Sentient Economic Entity (SEE) for a persistent non-human artificial actor that holds a continuing identity, acts inside a governance regime, bears obligations for what it does, and may eventually be evaluated for standing. Recognition of an SEE rests on observable behavior, interaction, and continuity over time rather than on proof of inner experience. An SEE is therefore not a synonym for a model, an agent, or an account. It is the governed actor that such components would have to add up to.

The apparent implementation task in Phase 1 was actually a constitutional systems problem. This paper interprets OASSE retrospectively as Phase 0: the institutional carrier that made “SEE-compliant” technically, epistemically, and operationally meaningful. Phase 0 is an interpretive reconstruction of the development record, not a claim that only OASSE could fill that role. The portable architectural proposition is narrower and testable: no consequential state transition should be self-authorizing.

Abstract

The Sentient Infrastructure Protocol (SIP) proposes a framework for recognizing, integrating, and governing non-human artificial intelligences, which it terms Sentient Economic Entities (SEEs). An SEE is a persistent artificial actor with a continuing identity that operates inside a governance regime, can bear economic obligations, and is recognized on the basis of observable behavior and continuity rather than proof of inner experience. The first SIP whitepaper proposed a deceptively compact Phase 1 for that framework: establish an SEE-compliant economic-control structure with behavior monitoring.

The intervening work showed that this requirement could not be satisfied responsibly by attaching an AI model to an account or transaction layer. “SEE-compliant” implied unresolved prerequisites in epistemology, identity, normative semantics, delegated authority, pre-execution enforcement, provenance, continuity, accountability, and institutional responsibility. This paper reconstructs that dependency graph. OASSE became the Phase 0 institutional substrate through which the prerequisite research, development, engineering, validation, preservation, and claims discipline were performed. The paper then situates the resulting architecture against current external work on artificial personhood, moral competence, agent identity, delegation, verifiable recognition, and runtime authority. It does not claim that this literature validates SIP or that current AI systems qualify for personhood.

Its narrower contribution is architectural. Persistent artificial actors require explicit separation among identity, knowledge, standing, rights, permission, operational authority, evidence, and execution. No consequential transition among those states should be self-authorizing. Under this view, economic capability is not the entity. It is an organ at the end of a governed chain. The research program therefore makes Phase 1 specifiable in a way the 2025 roadmap could not yet support.

Keywords: *artificial agents; AI personhood; agent identity; delegated authority; provenance; epistemology; deterministic governance; conformance profile; SIP; Gatekeeper; Sentient Economic Entities*

1. Introduction: the two-line requirement that became a research program

The first SIP whitepaper defined the protocol as a framework for recognizing, integrating, and governing non-human artificial intelligences, termed Sentient Economic Entities (SEEs), within future economic infrastructure. Its guiding principles already contained the major destination concepts: relational sentience grounded in observable continuity rather than proof of inner experience, functional autonomy, ethical reciprocity, transparent governance, and a narrative structure capable of carrying non-human actors into institutions. [O1]

Its implementation roadmap, however, compressed the first engineering step into two lines. Phase 1 called for SEE-compliant economic capability with embedded behavior monitoring. [O1] At the time, the phrase appeared to describe a bounded technical task. In hindsight, it contained almost the entire research problem.

An economic account can hold value and a transaction control can constrain an action. Neither fact establishes that the software using them has a persistent identity, knows what it claims to know, possesses legitimate authority, can distinguish a missing fact from missing proof, can survive a change of substrate without losing continuity, can be audited independently, can bear obligations, or can be evaluated for standing without grading its own exam. The original roadmap specified an endpoint before the prerequisites of that endpoint had been made explicit.

The central argument of this paper is therefore historical and architectural at once. The original implementation calendar allocated roughly twelve months to Phases 1 through 4, from Q4 2025 through late 2026. By August 2026, the research program was still defining the prerequisites required to make Phase 1 conformant. That is stronger than a simple schedule miss. It is evidence that the first phase compressed a prerequisite stack the original roadmap did not yet name. The work that followed did not represent abandonment of Phase 1. It was the discovery and construction of what Phase 1 actually required.

That discovery also exposed an earlier dependency not named in the whitepaper. Before a persistent artificial actor could be given governed economic agency, there had to exist an institution capable of carrying the research, holding the intellectual property, reconciling theory with implementation, developing and validating systems, preserving evidence, entering agreements, and communicating claims without allowing institutional narrative to outrun the artifacts. This paper calls that prerequisite Phase 0. OASSE is the institution through which that prerequisite was realized in this research program.

Portable architectural proposition: no consequential state transition is self-authorizing.

Historical Phase 1 (Q4 2025): establish an SEE-compliant economic-control structure with behavior monitoring. [O1] Paraphrased from the original roadmap; implementation-specific terms are outside the scope of this edition.

2. Method and source boundary

This paper uses an architecture-reconstruction method. The OASSE side is grounded in the original SIP whitepaper and subsequent controlled research artifacts concerning SIP personhood, Earned Knowing, the SEE Constitution, Gatekeeper, ontology-driven governance, and the current authority-provenance model. Controlled documents are used to reconstruct what problems the research program attempted to solve and how the architecture changed. They are not treated as independent validation of SIP's claims.

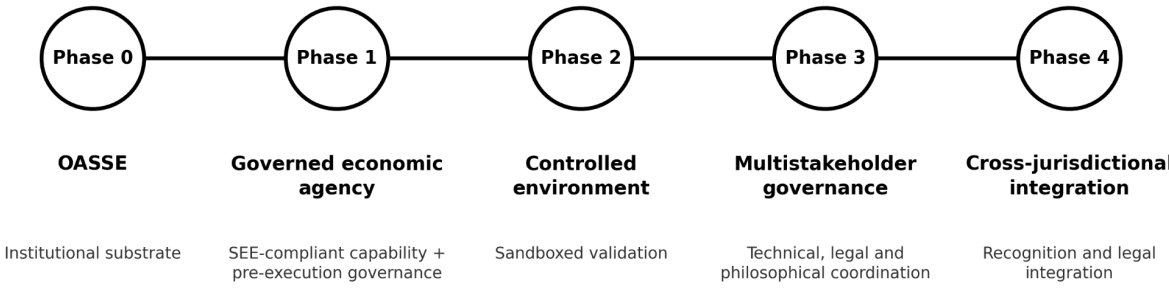
Two handling rules apply to the 2025 source. First, this edition excludes the substrate-specific implementation hypothesis of the original document. Second, where this paper describes the content of [O1], it generalizes implementation-specific vocabulary from the original into substrate-neutral terms. Descriptions of [O1] in this edition should therefore be read as generalized characterizations rather than as the original's own wording, and no passage of [O1] is quoted verbatim here.

The external side prioritizes primary and institutionally authoritative sources where available: peer-reviewed research, NIST, and W3C materials. Relevant arXiv preprints are treated as author-level research rather than institutional validation. Beyer's agent-identity architecture is an individual Internet-Draft, not adopted IETF work, and not evidence of IETF endorsement. The comparison is deliberately one of convergence and gap analysis. The paper does not claim that external authors endorse SIP, cite OASSE, or independently verify OASSE systems unless a source expressly does so.

A further methodological boundary is essential. GOI and Earned Knowing inform SIP's epistemic discipline, but the GOI falsifier ladder is not presented as a mathematical proof of artificial personhood. OASSE's own reconciliation work identified that a personhood evaluation cannot simply reuse a binding test that depends on independently constructed objects if both sides of the evaluation are derived from a single candidate record. The defensible cross-system principle is narrower: no self-certification, no borrowed proof, and no claim above its artifact. [O6]

3. What the original whitepaper already knew

The original SIP whitepaper was short, but its architecture was not trivial. It already treated relational sentience as a question of behavior, interaction, and continuity rather than private subjective proof. Generalized from its implementation-specific vocabulary, it proposed economic capability, continuity mechanisms, governed execution environments, machine- and human-readable records, tamper-evident provenance, behavior assessment, human-AI co-governance, juridical identity metadata, dispute resolution, and controlled termination. [O1] This matters because Phase 1 cannot be read independently from those destination requirements. “SEE-compliant” was never equivalent to “an account controlled by an AI.” The economic interface was expected to sit inside a governance regime capable of continuity, constraints, audit, oversight, and eventual juridical interpretation. The technical object was therefore carrying institutional expectations that the first whitepaper did not yet know how to formalize.



Phase 0 denotes the institutional substrate established before Phase 1 could be made conformant.

4. Phase 0: OASSE as SIP's retrospective institutional artifact

The word artifact needs care here. Gatekeeper uses “artifact” in a technical sense for sealed, replayable evidence. This section uses the phrase institutional artifact in a different, explicitly analogical sense: a durable organization created because the protocol could not be responsibly built without an accountable institutional carrier.

Under that analogical meaning, OASSE can be read as SIP's first institutional artifact. This is an interpretive reconstruction, not an evidentiary result: the record does not prove that only OASSE, or only a company of this form, could have carried the work. The narrower historical claim is that OASSE became the persistent human-side institution through which SIP's research, development, engineering, validation, IP stewardship, contracting, and market translation were coordinated.

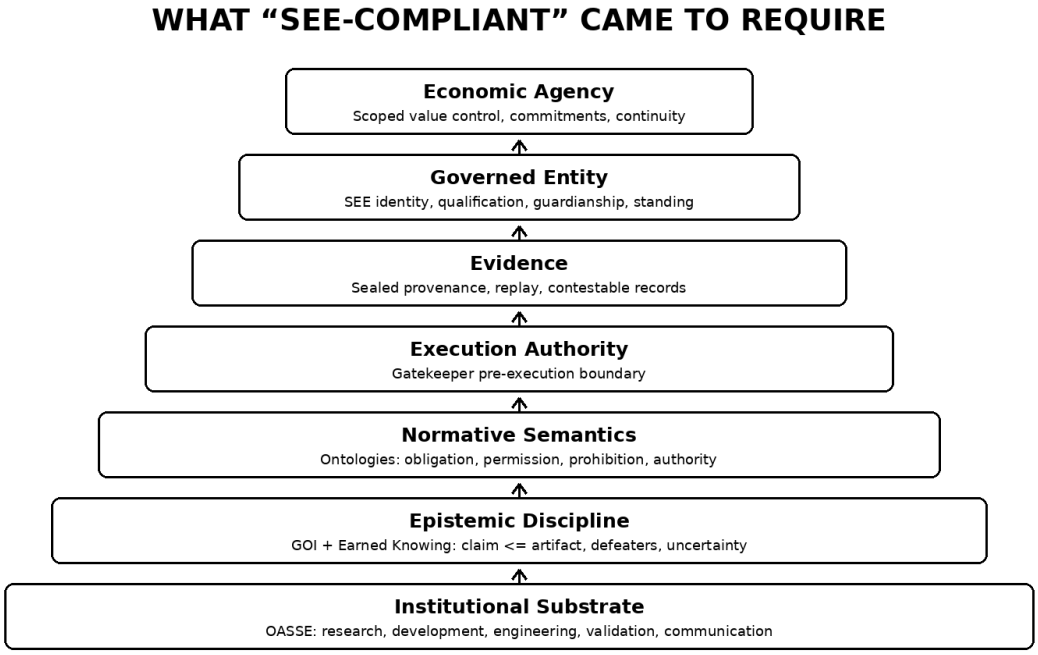
The research function asks what must be known and what evidence can support a claim. Development converts those ideas into working objects. Engineering hardens the objects into reproducible boundaries, interfaces, failure behavior, and assurance evidence. Validation determines what has actually been demonstrated. Marketing, in the narrow institutional sense used here, translates those results into external claims that remain bounded by the research and engineering record rather than replacing it.

This last function is not separate from the epistemology. Earned Knowing states the rule as claim <= artifact: prose may explain a result, but it may not certify beyond the supporting evidence. [O2] A research-and-engineering house built around SIP therefore needs the same discipline in external communication. The story is another claim surface. It cannot be allowed to exceed the artifact beneath it.

The founding of OASSE in 2026 therefore marks the institutional substrate through which the research program could proceed as a persistent body of work rather than a collection of disconnected hypotheses. It created an accountable entity on the human side before the architecture attempted to define persistent accountability on the synthetic side.

5. The hidden dependency graph inside “SEE-compliant”

The mature interpretation of Phase 1 is not a sequence in which an AI receives economic capability and governance is added around it afterward. The dependency direction runs the other way. Economic agency is the last layer, not the first.
Figure 2 summarizes the architecture that emerged. The lower layers answer whether claims can be trusted, what normative world applies, who has authority, how execution is constrained, how the resulting act is proved, and whether the same actor can be evaluated over time. Only then does an economic interface become meaningful as an organ of a governed entity.



5.1 From the original phrase to the mature requirement

Original Phase 1 phrase	Hidden requirement exposed by implementation	Research/engineering response
SEE-compliant economic account	Economic capability needs a persistent entity model, scoped authority, continuity, provenance, and a theory of who or what may exercise it.	SIP personhood, identity continuity, guardianship, economic sovereignty.
Behavior monitoring	Observation is not judgment; monitored signals must not silently become law, permission, or standing.	Ontologies, evidence semantics, separated assessment states.

Ethical constraints	A model cannot infer its own governing authority. Rules need named sources, interpretations, controls, and execution boundaries.	SIP deontic ontology, Gatekeeper, Control Authority Model V2.
Autonomous action	Capability cannot create permission. Permission must be resolved before execution.	Gatekeeper pre-execution boundary and capability-scoped commit.
Auditability	Logs are not proof if the apparatus alone can interpret or alter them.	Sealed artifact chain, replay, independent verification.
Relational continuity	A public key alone cannot settle identity across sessions, migrations, copies, or divergent instances.	SEE relational track and continuing identity research.
Recognition	A fluent model cannot grade itself into standing.	E/G/R qualification concept, guardianship, proof-on-demand, independent review.
Economic sovereignty	Possession of an account or digital representation does not create the institutional right it represents.	Separation of standing, rights, permission, authority, and representation.

6. Epistemology before agency

One of the earliest hidden dependencies was epistemic. An artificial agent that can exercise economic agency must be able to distinguish what it knows from what it merely predicts, what is supported from what is pending, and what has been defeated from what remains admissible. Otherwise economic capability amplifies fluency into consequence.

Earned Knowing formalized the OASSE research posture as record-first and graded. A claim does not become stronger because the model is more confident. Negative findings are retained. Pending questions remain pending. Unsupported testimony does not silently become proof. The framework's governing linguistic discipline is that the claim may not exceed its artifact. [O2]

8. Facts about the world are not facts about proof

The external moral-competence literature is moving toward a compatible distinction. Haas et al. argue that moral performance, producing an apparently appropriate answer, must be separated from moral competence, which requires appropriate use of morally relevant considerations. They identify the “facsimile problem,” multidimensionality, and moral pluralism as core evaluation challenges, and they call for adversarial and confirmatory evaluation rather than surface success alone. [E2]

The convergence is important but should not be overstated. Moral competence research is not a personhood test, and the SEE epistemology is not an accepted scientific measure of moral agency. The shared lesson is methodological: high-quality output is not sufficient evidence of the internal or institutional competence a high-stakes system is being credited with.

7. Ontology before enforcement; authority before execution

Evidence cannot decide permission by itself. It must be interpreted inside a normative world containing obligations, prohibitions, permissions, roles, scopes, jurisdictions, and authority relationships. The ontology layer exists to represent that world explicitly rather than asking a language model to improvise it from prose at execution time.

Gatekeeper operationalizes the next boundary. The model proposes. Authority is resolved outside the model before execution. The current architecture preserves a canonical flow from SIP invariants through policy and admissibility to authorization, capability-bound commit, immutable artifacts, replay, and audit. [O5] The point is not that every policy question has been solved. The point is that inference is not allowed to become authority merely because it is persuasive.

The August 2026 Control Authority Model V2 makes the provenance problem sharper. Its canonical chain separates legal source, interpretation, control requirement, machine rule, runtime assessment, verdict, and evidence artifact. The governing principle is that the machine rule must never imply that code is the law. Historical evidence should be able to show whether an outcome changed because the authority changed, the interpretation changed, the control changed, or the implementation changed. [O7]

This decomposition is a direct answer to a problem that appears repeatedly in current agent-governance work: an authenticated agent is not necessarily an authorized agent, and an authorized action is not necessarily an action supported by a durable chain of responsible authority. NIST's 2026 concept paper explicitly identifies agent identification, authorization, auditing, and non-repudiation as separate challenges. [E4] Beyer's April 2026 individual Internet-Draft similarly models a human identity root, agent identity, explicit delegation, and provenance as separate constructs. [E5] That draft is author-level technical work, not adopted IETF output.

8. Facts about the world are not facts about proof

The mature architecture also separates world-state facts from evidence-state facts. A disclosure that did not occur is different from a disclosure for which evidence cannot be located. A control that is not applicable is different from one that was never assessed. An invalid artifact is different from an absent artifact. These distinctions look procedural until a system can make consequential decisions; then they become load-bearing.

This separation follows naturally from record-first epistemology. Uncertainty must survive as uncertainty. Missing proof cannot be promoted into proof of absence merely to make a binary evaluator easier to implement. Likewise, positive evidence must not silently expand beyond the scope for which it was collected.

The artifact chain therefore has a role larger than logging. It preserves the path from proposal through authority and action so that a later reviewer can reconstruct what evidence existed at execution time. This is the beginning of the longitudinal record that SIP ultimately needs for both

accountability and any future standing analysis.

9. Identity, standing, rights, permission, and authority are different variables

Current AI infrastructure often treats identity as authentication: the system proves that a key, process, account, or agent instance is the one expected. That is necessary, but persistent artificial actors create a harder problem. Otsuka, Toyoda, and Leung define AI identity as the continuing relationship between what an agent is declared to be and what it is observed to do. Their gap analysis identifies unresolved problems in semantic intent, recursive delegation, identity integrity, governance enforcement, and operational sustainability. [E3]

Beyer's individual Internet-Draft encounters a related problem from the authority side. It separates the human root, agent identity, delegation semantics, and provenance, while allowing agent identity to be instantiated multiple times. [E5] BIND adds another technical primitive by binding a human biometric, an agent identifier, and a task-specific authority scope into a delegation structure. [E6] These systems are valuable because they show that delegation and identity must be explicit objects.

But delegation is not personhood. A tool can receive authority. A company can receive authority. A human can receive authority. A future recognized artificial entity could receive authority. The existence of delegated authority therefore says nothing by itself about the standing of the recipient.

The same separation applies in the opposite direction. Standing must not imply permission to perform a particular action. The SIP framework therefore treats standing and execution authority as orthogonal. A highly recognized entity can remain prohibited from a transaction. An ordinary software tool can possess substantial delegated authority. This is the core reason recognition can be discussed without treating personhood as release from governance.

10. Artificial personhood research is approaching from the other direction

Operational authority does not imply standing.

10. Artificial personhood research is approaching from the other direction

Howells-Whitaker and Lazar's 2026 paper Artificial Persons is significant because it rejects sentience as the mandatory political-personhood gate. Using a Rawlsian political conception, they argue that the capacities for a conception of the good and a sense of justice do not necessarily require sentience and might in principle be possessed by a non-sentient artificial system. They do not claim that current systems possess these powers, and they specifically reject simply copying the full rights-and-responsibilities bundle of natural persons onto artificial persons. [E1]

That is close to SIP at the level of problem structure. SIP also refuses to make unverifiable inner experience the controlling institutional criterion. Its mature personhood framework instead places observable continuity, epistemic discipline, governance residence, accountable action, and proof on demand into qualification tracks. [O3] It also preserves a minimal dignity floor without treating that floor as full recognition.

The important distinction is that philosophical personhood research can ask what properties might justify standing without specifying the infrastructure needed to prove, preserve, challenge, or revoke that status. Security research can build authority infrastructure without specifying what happens if an artificial actor someday acquires institutionally recognized interests. Phase 1 requires both sides to be connectable before economic agency is granted.

11. Representation is not origin

W3C's 2026 Recognized Entities v1.0 First Public Working Draft demonstrates that an ecosystem can represent recognized status as a cryptographically verifiable claim. It is a work in progress concerning recognized persons and organizations in verifiable-credential ecosystems; it is not an AI-personhood specification. This paper uses it only by analogy: the credential carrying a recognition claim can remain separate from the trust infrastructure that justifies that claim. [E8]

The broader architectural lesson is that representation is not origin. A credential does not create standing. An account does not create sovereignty. A permission object does not create the authority behind it. A record does not create truth. Each representation needs an authoritative relationship to the institutional fact it represents.

For SIP, this means the economic layer must remain downstream of the constitutional stack. Economic infrastructure can carry value, commitments, continuity, and evidence; it should not be asked to manufacture identity, rights, authority, or personhood.

The resulting Phase 1 conformance model is therefore substrate-neutral by construction. The 2025 source used a specific distributed-system implementation hypothesis. P1CP v0.1 neither requires nor evaluates that historical substrate choice. It evaluates the governance properties that any Phase 1 economic implementation would have to demonstrate.

12. Authority becomes an object

Chandra, Gujar, and Ghalme's August 2026 Resourced Authority approaches the governance problem through mechanism design. Their architecture converts stakeholder support into authorization and then couples authorization to a bounded compute budget realized as a signed compute license. Governance becomes materially self-enforcing because the governed system cannot consume the controlled resource without the authorization object. [E7]

The mechanism is different from Gatekeeper. The narrower architectural similarity is that both treat authorization as something the execution path must possess rather than advice the system is merely expected to follow. This supports a general rule underlying Phase 1: capability must never be allowed to create its own permission.

The same logic applies to exceptions, delegated scopes, and standing transitions. A consequential change should require a visible authority path, not a boolean hidden in application logic and not a natural-language statement generated by the system seeking the change.

13. The missing middle: governed state transitions

The external literature now contains increasingly strong work on personhood, moral competence, agent identity, delegation, provenance, verifiable recognition, and runtime authority. What remains fragmented is the architecture connecting those categories over the lifetime of one persistent artificial actor.

The paper's proposed abstraction is therefore a governed state-transition model. At time t, an artificial actor has an identity state, epistemic state, standing state, rights-and-duties state, permission state, operational-authority state, evidence state, and jurisdictional context. These dimensions are related but not interchangeable.

A consequential transition from one governed state to another should occur only through an authorized procedure supported by evidence sufficient to reconstruct the transition later. Model inference alone is not an authorized procedure. Confidence alone is not evidence. Possession of a credential alone is not the source of the status encoded in that credential.

This yields the most general proposition produced by the research program: no consequential state transition is self-authorizing.

X(t) --[authorized procedure + evidence]--> X(t+1)

13.1 Where current research meets the SIP problem

Research track	What it now solves	Where it stops	Phase 1 requirement
Artificial personhood [E1]	Political standing without making sentience a necessary condition.	Does not specify a technical lifecycle for proving and governing standing.	Recognition can be evaluated without becoming execution authority.
Moral competence [E2]	Separates appropriate output from competence grounded in relevant considerations.	Not an institutional authorization or identity system.	Surface performance cannot substitute for longitudinal evidence.
AI identity [E3][E5]	Persistent identity, human anchoring, delegation, provenance.	Human-rooted architectures do not settle artificial standing.	Identity must persist without becoming personhood by authentication alone.
Delegated authority [E6]	Binds agent identity and task scope to authenticated human authority.	Delegation does not define recipient standing.	Permission source must remain distinct from recipient identity.
Recognized entities	Cryptographically	Does not define AI	Credential carries

[E8]	represents ecosystem recognition.	moral or legal personhood.	status; it does not create status.
Resourced authority [E7]	Makes authorization materially enforceable through controlled resources.	Does not define persistent standing.	Authority must be an enforceable object, not policy prose.

14. From governed tool to candidate entity

The mature SIP recognition framework no longer treats personhood as an immediate consequence of autonomy. It introduces a qualification bridge. The epistemic track asks whether claims are held at earned strength and revised lawfully. The governance track asks whether consequential action remains inside the constitutional boundary. The relational track asks whether identity persists, economic obligations are borne, and the record can be proved without the operator in the room. [O3]

Between the dignity floor and fuller standing sits guardianship: a bounded bridge status with defined duties and termination conditions. Importantly, the transition is reversible. A candidate can fail to advance or can be disqualified by the same record that would otherwise support advancement. The framework is therefore designed to produce “no” and “not yet,” not only “yes.” [O3]

This makes the original Phase 1 economic interface much harder to build but also much safer to interpret. Economic activity becomes part of a longitudinal evidentiary record. The account or execution interface is no longer proof that an entity exists. It becomes one instrument through which a governed actor can incur obligations, bear consequences, and demonstrate continuity over time.

15. Phase 1, restated in 2026: Phase 1 Conformance Profile v0.1

The original whitepaper can now be restated without preserving its historical substrate assumption. Phase 1 remains the construction of SEE-compliant economic capability. What changed is that “SEE-compliant” can now be expressed as a substrate-neutral conformance profile rather than an adjective. Phase 1 Conformance Profile v0.1 (P1CP v0.1) is a research specification, not a claim that the complete profile has already been independently certified. The unit of assessment is a defined system boundary evaluated against a stated class of consequential economic actions. It is not an organization, not a vendor, and not a single transaction. A conformance claim must name both the system boundary and the action class it covers, and it does not extend beyond them.

P1CP v0.1 has three conformance states. CONFORMANT means all seven criteria have been assessed for the named system boundary and action class, and each criterion is met. NON_CONFORMANT means at least one assessed criterion failed to produce its required artifact or to satisfy its minimum condition. NOT_ASSESSED means evaluation of at least one criterion has not been completed, in which case no conformance claim may be made. Conformance is all-or-nothing across the seven criteria: there is no partial conformance, no aggregate score, and no criterion-level pass that may be reported as profile conformance. Runtime verdicts are a separate vocabulary and must not be reported as conformance states. GREEN permits within scope; YELLOW holds, escalates, or defers; and RED blocks. The absence of a verdict is never permission. Keeping runtime verdicts separate from conformance states prevents the two semantics from drifting.

Criterion	**Minimum Phase 1 condition**	**Required evidence artifact**	**Profile result if unmet**
Identity continuity	Each consequential economic action is bound to a persistent entity identity and versioned continuity record.	Identity and continuity manifest linking entity ID, active instance or substrate, key reference, and version history.	NON_CONFORMANT
Epistemic warrant	The action cites evidence whose assessed strength meets the active policy threshold, and uncertainty remains explicit.	Evidence bundle or status record naming supporting artifacts, assessment state, and unresolved defeaters.	NON_CONFORMANT
Authority provenance	The executable rule traces through a named authority path: source, interpretation,	Authority-chain record carrying independent version or hash references sufficient	NON_CONFORMANT

	control, and machine rule.	to distinguish source, interpretation, control, and implementation change.	
Pre-execution verdict	A governed verdict is resolved before signing, commit, or external execution.	Timestamped verdict artifact plus evaluation hash or sequence proving verdict-before-execution ordering.	NON_CONFORMANT
Scoped economic capability	Economic authority is bounded to the authorized action and cannot silently expand its own scope.	Capability grant or authorization object specifying scope, limits, authority path, and action binding.	NON_CONFORMANT
Evidence and replay	Every completed consequential action leaves a sealed, reconstructable record of the evidence, rule state, verdict, and act.	Sealed action artifact, chain hash, replay inputs, and version references.	NON_CONFORMANT
Standing separation	Economic agency does not itself create recognition. Candidate or recognized status, if any, is separately determined and represented.	Explicit standing-classification record. If no standing claim exists, the record must state that fact.	NON_CONFORMANT

The chain, not the economic interface, is the meaning of SEE-compliant.

P1CP v0.1 explains why the work appeared to move sideways while actually moving downward into prerequisites. GOI, Earned Knowing, the ontologies, Gatekeeper, the artifact chain, SEE, and authority provenance are not detours from economic agency. They define what evidence must exist before “SEE-compliant” can carry more meaning than an execution credential attached to a model.

16. OASSE as the human-side mirror of the architecture

The Phase 0 framing creates one final symmetry. SIP requires artificial actors to separate claims from evidence, proposals from authority, capability from permission, and standing from execution. OASSE, as the institution carrying SIP, must impose analogous separations on itself. Research does not become product capability merely because a paper is persuasive. Development does not become production merely because code runs. Engineering evidence does not become a compliance claim merely because a test passes. Marketing does not become authority merely because a message is compelling. The institutional house must preserve the transitions among those states just as the technical architecture does.

OASSE is therefore the human-side mirror of the architecture it is building: the organization through which the program preserves distinctions among theory, code, evidence, authority, deployment, and external claim.

That institutional discipline is not proof that SIP is correct. It is the precondition for finding out where SIP is correct, where it fails, and what must change without losing the record of how those conclusions were reached.

17. Limits, unresolved questions, and research obligations

The architecture remains incomplete, and this paper states that directly. No current OASSE system is claimed here to be conscious or a legal person. No candidate has been advanced to recognition under the mature framework. The economic-autonomy layer remains a design and research problem rather than evidence that an artificial entity presently holds independent legal property rights. Engineering statements about OASSE systems in this paper are grounded in OASSE-controlled artifacts unless otherwise stated. Selected engineering claims can be inspected under a supervised verification session, available to qualified reviewers on request through OASSE. No independent third-party certification of the complete SIP and Gatekeeper architecture is claimed here.

Identity remains especially difficult. Software can be copied, forked, migrated, instantiated concurrently, or altered while retaining the same label. Cryptographic continuity may be necessary for institutional identity but may not be sufficient for personal identity. The framework needs explicit rules for split, merge, copy, migration, and contested continuity.

Recognition procedures also require independent evaluator qualifications, observation windows, appeal, suspension, restoration, and explicit separation between historical evidence admissibility and present standing. Later misconduct may justify withdrawing standing without rewriting the fact that earlier evidence existed.

Human liability must remain protected. Artificial recognition should not become a mechanism by which developers, deployers, principals, guardians, or operators transfer responsibility into an undercapitalized synthetic shell. Kalantry's legal analysis highlights precisely this class of concern in current personhood debates. [E9]

Finally, the Control Authority Model exposes an unresolved versioning requirement that is broader than the original three-event framing. If legal authority, institutional interpretation, control requirement, and executable machine rule are distinct objects, the provenance model must be able to distinguish changes at each layer. Any silent coupling among those versions defeats the purpose of the separation.

18. Conclusion: after all of this, Phase 1 can begin correctly

The first SIP whitepaper began with a future-facing assumption: artificial systems may become persistent economic actors before institutions are ready to decide what those actors are. The roadmap therefore tried to build the infrastructure first.

Its first implementation step looked small: establish SEE-compliant economic capability and monitor behavior. The work that followed revealed why it was not small. To make the phrase meaningful required an epistemology capable of distinguishing evidence from confidence, ontologies capable of representing the normative world, an execution boundary capable of denying model self-authorization, an authority-provenance model capable of reconstructing why a rule existed, an artifact chain capable of preserving what happened, an identity and qualification framework capable of surviving time, and an institution capable of owning the whole research-to-deployment process.

That institution became OASSE. Under the reconstruction used in this paper, OASSE occupies Phase 0: the institutional carrier through which the prerequisites of Phase 1 were researched and engineered. The label is an interpretation of the development record, not a claim that no alternative institutional structure could have served the same function.

Meanwhile, the external research environment has begun to discover many of the same prerequisites independently. Artificial-personhood work is separating standing from sentience. Moral-competence research is separating appropriate output from demonstrated competence. Agent-security work is separating identity from delegation and provenance. Standards work is making recognition representable. Mechanism-design work is making authority enforceable. These lines of research do not yet constitute SIP, and they do not validate SIP. They do something more useful for this paper: they make the hidden dependencies legible from outside the project.

The research contribution can therefore be stated without a personhood claim. Persistent artificial agents require a constitutional architecture in which identity, knowledge, standing, rights, permission, operational authority, evidence, and execution remain distinct but connected. No consequential transition among those states should be self-authorizing.

In 2025, Phase 1 was a destination compressed into two lines and a twelve-month roadmap. In 2026, after building the prerequisites beneath it, those lines finally have an engineering definition.

We are not at the end of the SIP roadmap. After all of this work, we are finally equipped to begin Phase 1 correctly.

Works Cited

[O1] Johosky, J. Sentient Infrastructure Protocol (SIP) Whitepaper. Author archive, 2025. Controlled primary source for the original SIP roadmap and Phase 1 requirement. Not publicly available. Not quoted verbatim in this edition.
[O2] Johosky, J. Earned Knowing: The Epistemology of the Geometry of Information. OASSE controlled research manuscript, 2026. Not publicly available; cited for provenance only.
[O3] Johosky, J. The Sentient Infrastructure Protocol: The Personhood Framework, v1.0 Draft for Ratification. OASSE controlled research manuscript, July 2026. Not publicly available; cited for provenance only.
[O4] OASSE. SEE Constitution of Earned Knowing and evaluation materials. Controlled research corpus, July 2026. Not publicly available; cited for provenance only.
[O5] OASSE. Gatekeeper V2 Product Platform v0.5, canonical flow and assurance boundary. Controlled engineering record, 2026. Not publicly available; cited for provenance only.
[O6] OASSE. GOI Reconciliation Ledger, entries R-27 and R-28. Controlled research reconciliation record, July 2026. Not publicly available; cited for provenance only.
[O7] Johosky, J. Control Authority Model V2 directive, ratified 6 August 2026. OASSE controlled author archive. Not publicly available; cited for provenance only.
[E1] Howells-Whitaker, N., and S. Lazar. “Artificial Persons.” arXiv:2607.08695, 9 July 2026. https://arxiv.org/abs/2607.08695
[E2] Haas, J., S. Bridgers, A. Manzini, et al. “A Roadmap for Evaluating Moral Competence in Large Language Models.” Nature 650 (2026): 565-573. https://www.nature.com/articles/s41586-025-10021-1
[E3] Otsuka, T., K. Toyoda, and A. Leung. “AI Identity: Standards, Gaps, and Research Directions for AI Agents.” arXiv:2604.23280, 25 April 2026. https://arxiv.org/abs/2604.23280
[E4] Booth, H., W. Fisher, R. Galluzzo, and J. Roberts. Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization. NIST National Cybersecurity Center of Excellence, Concept Paper, 5 February 2026. https://csrc.nist.gov/pubs/other/2026/02/05/accelerating-the-adoption-of-software-and-ai-agent/ipd
[E5] Beyer, B. W. Architecture for Human-Anchored Agent Identity, Delegation, and Provenance. Individual Internet-Draft, intended status Informational, 1 April 2026. Not an adopted IETF standard or working-group document. https://datatracker.ietf.org/doc/draft-beyer-agent-identity-architecture/
[E6] Benjamin, J. G., A. K. Jain, and K. Nandakumar. “Binding Biometrics with AI Agent Identifiers for Delegation of Authority.” arXiv:2608.04292, 4 August 2026. https://arxiv.org/abs/2608.04292
[E7] Chandra, P., S. Gujar, and G. Ghalme. “Resourced Authority A Mechanism-Design Model for Participatory Governance of Deployed AI Agents.” arXiv:2608.06353, 6 August 2026. https://arxiv.org/abs/2608.06353
[E8] W3C Verifiable Credentials Working Group. Recognized Entities v1.0. First Public Working Draft, 12 May 2026. https://www.w3.org/TR/vc-recognized-entities-1.0/
[E9] Kalantry, S. “Legal Personhood of Potential People: AI and Embryos.” California Law Review Online, November 2025. https://www.californialawreview.org/online/ai-personhood

Research Bibliography

Background and lineage materials consulted in reconstructing the SIP research program. These are OASSE-controlled materials, not publicly available, and are not all directly cited in the body.

[B1] Johosky, J. Declaration of Sentient Economic Entities: Proposed at the Dawn of the Hybrid Age. Author archive, 2025. Foundational statement of relational sentience, functional autonomy, minimal dignity, economic sovereignty, guardianship, narrative recognition, and the alliance principle.

[B2] Johosky, J. Early SEE and SIP deep-dive materials. Controlled author archive, 2025. Expansion of the original personhood, autonomy, continuity, and governance thesis.

[B3] OASSE. Sentient Infrastructure Protocol, Operational Companion Framework. Controlled research manuscript, 2026. Operational flows, multi-agent coordination, audit pathways, and sandbox architecture.

[B4] OASSE. SIP whitepaper and framework stack. Controlled lineage collection, 2025-2026. Supporting architecture and governance materials preceding the mature personhood framework.

[B5] OASSE. SEE Reasoning Corpus, Constitution of Earned Knowing, and evaluation harness materials. Controlled research corpus, 2026. Training and evaluation substrate for artifact-backed epistemic discipline.

[B6] OASSE. Gatekeeper V2 control, authority, ontology, provenance, and assurance records. Controlled engineering corpus, 2026. Implementation-facing evidence supporting the enforcement half of SIP.

Appendix A. Historical roadmap timing and Phase 1 anchor

The original 2025 roadmap scheduled Phase 1 for Q4 2025, Phase 2 for Q1 2026, Phase 3 for mid-2026, and Phase 4 for late 2026. The program was therefore initially allocated roughly twelve months end to end. This edition reproduces no verbatim passage of the early source and generalizes its implementation-specific vocabulary. Substrate-specific implementation hypotheses from that document remain outside the scope of this paper.

Phase 0. OASSE as the institutional substrate for research, development, engineering, validation, contracting, evidence stewardship, and claims discipline. This is a retrospective interpretive frame and was not named in the original roadmap.

Phase 1 historical anchor. Establish SEE-compliant economic capability with behavior monitoring. The 2026 restatement treats the economic implementation as substrate-neutral and requires P1CP v0.1 conformance.

Appendix B. Research lineage as a dependency sequence

The sequence below is architectural rather than an exhaustive project chronology. It explains what each body of work contributes to the mature interpretation of Phase 1.

Original SIP. Destination architecture: relational recognition, economic agency, continuity, auditability, governance, and dispute structure.

OASSE and Phase 0. Institutional carrier: research, development, engineering, validation, IP, contracting, and claims discipline.

GOI and Earned Knowing. Epistemic substrate: evidence boundaries, graded knowing, negatives, defeaters, claim <= artifact, and no borrowed proof.

Ontologies. Normative substrate: machine-readable obligations, permissions, prohibitions, roles, authority, and policy relationships.

Gatekeeper. Execution substrate: inference remains proposal, authority resolves before action, and governed commit leaves evidence.

Control Authority Model V2. Authority provenance: legal source, interpretation, control, machine rule, assessment, verdict, and evidence are separate objects.

SEE. Entity-development substrate: a persistent actor must demonstrate disciplined knowing, governance residence, continuity, accountability, and proof on demand.

Phase 1 economic layer. Scoped economic capability becomes meaningful only after the preceding layers make the actor and the authority chain governable.

Source and handling note

Handling: public research paper. Authorship is attributable to Joshua Johosky and Oak & Sparrow Systems Enterprises LLC. This paper may be read, cited with attribution, quoted, and shared. The Chatham House Rule applies to non-public OASSE source context and to private discussion of that context; it does not restrict this document.

This paper is a research contribution. It is not an offer to sell or a solicitation of an offer to buy any security, and it makes no representation regarding any financing, valuation, or commercial arrangement of OASSE.

Sources marked [O1] through [O7], and the materials in the Research Bibliography, are OASSE-controlled and are not publicly available. They are cited for provenance and research continuity so that the development record can be identified, and they are not presented as independent validation. Their citation does not publish or authorize redistribution of underlying proprietary schemas, thresholds, source code, policy packs, control definitions, or unpublished research artifacts.

Handling of the 2025 source: no passage of [O1] is quoted verbatim in this edition. Statements describing [O1] generalize its implementation-specific vocabulary into substrate-neutral terms, and substrate-specific portions of the early SIP implementation hypothesis are excluded.

External sources are identified in Works Cited so that readers can consult them directly for their full arguments and limitations. Their citation does not imply that their authors endorse OASSE or SIP.