OSAI: A Neutral Naming and Reference Substrate for Software and AI Agents

Standards-Facing Memo (Public)

Version:	v1.0
Status:	Issued
Effective Date:	March 5, 2026
Document Identifier:	OSAI-FW-MEM-01
Publication Classification:	Public (Standards Submission Artifact)
Governing Record:	OSAI Document Control & Revision History (OSAI-FW-L2-DC-01)
Author:	Austin M. Hall
Scope:	Conceptual, non-operational; model-agnostic naming/reference substrate for AI agent identity discussions.
Reference Artifacts:	OSAI-FW-ABS-01 (Abstract & Preface); OSAI-FW-RO-02 (Reader's Orientation); OSAI-FW-L2-DC-01 (Document Control & Revision History)

1) Purpose

This memo outlines a narrow, upstream concept: a neutral naming and reference substrate for software and AI agents. The intent is to support interoperability, auditability, and accountability across downstream implementations (authentication, authorization, policy, enforcement) without competing with those downstream functions.

The intent of this document is to describe conceptual reference primitives and a separation principle; it does not define requirements, standards, or implementation.

2) Problem Statement

As autonomous and semi-autonomous agents move into operational environments, organizations face a recurring failure mode: actions occur without a stable, attributable identity reference that persists across systems, vendors, and logs.

These conditions tend to surface across sectors once agents operate autonomously. They are structural: auditability and authorization cannot scale when the object being governed is not nameable in a consistent way. In practice, this creates:

- Identity ambiguity: agents inherit human credentials or appear as opaque non-human workloads

- Audit incoherence: logs cannot reliably bind actions to a consistent agent identity reference
- Lifecycle confusion: ephemeral agents, delegated authority, and rotated credentials fragment traceability
- Interoperability friction: different stacks name and represent “agents” differently, preventing consistent cross-system reasoning

3) OSAI Concept (Upstream Only)

OSAI is positioned as a reference layer: a way to name an agent (or agent instance) such that downstream systems can attach authentication, authorization, policy, and audit evidence without forcing a single vendor or standards profile.

OSAI is not:

- an Identity and Access Management (IAM) product
- an authentication mechanism (authN)
- an authorization system (authZ)
- an enforcement layer, gateway, or control plane
- a governance body or regulatory framework

OSAI is intended to be identity-model agnostic: it is designed to map into multiple identity systems and credential schemes without replacing them.

4) Minimum Properties for a Useful Agent Reference Layer

A neutral naming substrate is only useful if it is stable enough to be cited and mapped. In practice, a reference identifier benefits from:

- Uniqueness: a globally unambiguous reference, within defined issuer semantics
- Issuer / provenance: who asserted the identifier and under what authority
- Lifecycle: persistent vs ephemeral identifiers; rotation; supersession; revocation references
- Linking: the ability to associate audit evidence and policy decisions across systems
- Non-operational posture**:** the identifier does not imply permission, trust, or compliance on its own

This memo does not propose an implementation standard. It proposes a separation principle: naming / reference is best kept distinct from authentication and authorization to preserve interoperability across downstream identity and control systems.

5) Mapping Posture (How OSAI Coexists with Downstream Standards)

Downstream systems may use OAuth/OIDC profiles, workload identity standards, enterprise identity governance platforms, or sector-specific authorization mechanisms. OSAI is designed to map into these frameworks:

- OSAI identifiers can be represented as claims, attributes, or metadata inside existing identity and authorization flows
- OSAI does not require a specific credential type
- OSAI does not dictate a policy engine or permission model
- OSAI provides a consistent reference anchor that downstream systems can use to bind:
  - permissions and scopes
  - delegation chains
  - audit logs and evidence artifacts
  - ownership attribution and accountability assignment

6) Standards-Facing Considerations (Non-Normative)

To reduce fragmentation and avoid premature lock-in, standards efforts may consider:

- Preserving separation: naming/reference ≠ authentication ≠ authorization ≠ enforcement
- Representing provenance semantics: "who asserted this identifier" should be expressible
- Supporting lifecycle semantics: ephemeral agents and rotated credentials should remain auditable
- Remaining model-agnostic: avoid embedding vendor-specific semantics as default "agent identity models"
- Treating auditability as first-class: identifiers are most useful when usable as evidence anchors across logs and systems

7) References (Issued Public Artifacts)

OSAI-FW-L2-DC-01:
https://www.InfrastructureOSAI.com/assets/OSAI-FW-L2-DC-01.pdf
Document Control & Revision History (Issued)

OSAI-FW-ABS-01:
https://www.InfrastructureOSAI.com/assets/OSAI-FW-ABS-01.pdf
Abstract (Issued)

OSAI-FW-RO-02:
https://www.InfrastructureOSAI.com/assets/OSAI-FW-RO-02.pdf
Reader's Orientation (Issued)