OSAI Framework - Layer 2

Identity and Naming as Infrastructure for Artificial Intelligence

Conceptual Reference White Paper

Author	Austin M. Hall

Version	v1.0
Original Date	February 9, 2026
Public Issue Date	June 8, 2026
Document Identifier	OSAI-FW-L2-01
Status	Issued
Publication Classification	Public Conceptual Reference White Paper
Governing Record	OSAI-FW-L2-DC-01
Website	InfrastructureOSAI.com

Document Control

Document Identifier	OSAI-FW-L2-01
Title	OSAI Framework - Layer 2: Identity and Naming as Infrastructure for Artificial Intelligence
Version	v1.0
Original Date	February 9, 2026
Public Issue Date	June 8, 2026
Author	Austin M. Hall
Status	Issued
Publication Classification	Public Conceptual Reference White Paper
Governing Record	OSAI-FW-L2-DC-01
Related Artifacts	OSAI-FW-ABS-01; OSAI-FW-RO-02; OSAI-FW-L2-DC-01; OSAI-FW-MEM-01; OSAI-FW-TCN-01; OSAI-FW-GRM-01; OSAI-FW-SCH-01; OSAI-FW-RSL-01; OSAI-FW-DMO-01; OSAI-FW-DIF-01; OSAI-FW-REL-01; OSAI-FW-NVP-01
Supersession	None. Initial public issue of the Layer 2 White Paper.
Copyright	

Table of Contents

Publication Note - Relationship to Later OSAI Materials

The Layer 2 White Paper is the conceptual root of the OSAI Framework - Layer 2 Series. The white paper defines identity and naming primitives for artificial intelligence systems as a conceptual, upstream, non-operational, and non-normative reference framework.

Later OSAI materials, including the Standards-Facing Memo, Technical Concept Note, Canonical Identifier Grammar, Three-Record Schema Template Pack, Resolver Illustration, Narrow Demonstration, Differentiation Note, Release Note, and Namespace and Vocabulary Posture Note, provide companion context for how the conceptual distinctions described in the white paper may be represented in later controlled materials.

Those later materials do not modify, replace, expand, or retroactively alter the white paper's scope, non-scope, interpretive boundaries, or non-operational posture. They should be read as companion materials, not as retroactive changes to the conceptual framework.

OSAI remains positioned as an upstream canonical reference identity and naming framework. OSAI does not authenticate, authorize, enforce, certify, govern, monitor, control, or determine legal responsibility for artificial intelligence systems or software agents.

Contemporary Context Note

Recent activity around software and AI agents has increased attention on agent discovery, authentication, authorization, non-human identity governance, runtime control, auditability, and machine-mediated commerce. Those developments reinforce the need for stable reference identity, but they remain downstream or adjacent to the scope of this white paper.

OSAI does not replace operational identity, authorization, security, protocol, compliance, certification, registry, or runtime-control systems. The OSAI Framework defines a neutral naming and reference posture for describing AI systems and software agents so that authority, attestation, sectoral context, and jurisdictional context may be associated with a stable referent without collapsing those downstream functions into OSAI itself.

Abstract

Artificial intelligence systems are increasingly deployed as operational components within regulated and high-liability environments, including healthcare, finance, energy, logistics, and public administration. As these systems transition from experimental tools to embedded decision-support and coordination mechanisms, existing institutional frameworks encounter persistent difficulty addressing foundational questions of attribution, jurisdiction, authority, and accountability.

This work proceeds from the premise that effective governance, auditability, and enforceable constraint presuppose the ability to reliably identify and situate artificial intelligence systems within institutional contexts. For purposes of this framework, identity refers to system-level attribution and classification rather than legal personhood or independent standing. In the absence of standardized, machine-resolvable identity, downstream governance functions are implemented through fragmented controls that do not scale and are difficult to evaluate under regulatory or legal scrutiny.

The OSAI framework is presented as a naming and identity substrate intended to address this structural deficiency. Rather than prescribing regulatory outcomes or governance models, it provides a consistent mechanism for identifying, classifying, and referencing artificial intelligence systems by sector, jurisdiction, and operational scope. Policy, regulatory, and legal determinations remain the responsibility of appropriate authorities; this framework is limited to defining identity primitives that may support such downstream determinations.

This document adopts an infrastructure-oriented perspective. It asserts that responsible deployment of artificial intelligence depends not solely on model capability or performance, but on the existence of stable identity primitives that enable institutional oversight, accountability, and constraint. Absent such primitives, organizations are compelled to rely on manual, contractual, or ad hoc mechanisms that degrade under scale, automation, and cross-jurisdictional operation.

Status & Intent

This document is issued as a conceptual reference framework. It does not propose a technical specification, governance model, or standard, nor does it seek formal adoption, endorsement, or implementation. The purpose of this document is to clarify identity and naming primitives for artificial intelligence systems at a time when such concepts are increasingly discussed but inconsistently defined. The framework is intended to support analysis, dialogue, and downstream development by providing a neutral, non-normative point of reference.

Scope, Non-Scope, and Interpretive Boundaries

Scope

This document defines identity and naming primitives for artificial intelligence systems as a foundational reference layer. The objective is to support clear reasoning, attribution, and cross-institutional dialogue regarding artificial intelligence systems operating across organizational, sectoral, and jurisdictional boundaries.

Within scope are:

- Artificial intelligence systems operating autonomously or semi-autonomously
- Identity as a descriptive and referential construct, independent of runtime state
- Naming as a prerequisite for attribution, continuity, and accountability
- Jurisdictional anchoring of AI system identity
- Sectoral classification as a stabilizing context for governance
- Sovereignty boundaries separating system identity, operator identity, and institutional authority

This document is intended to function as an upstream reference layer, enabling downstream actors, such as regulators, insurers, standards bodies, and system operators, to reason about AI systems without presupposing specific technical implementations or enforcement mechanisms.

Non-Scope

This document explicitly does not define, prescribe, or recommend:

- Authentication mechanisms
- Authorization or access-control models
- Identity and Access Management (IAM) architectures
- Security tooling or cryptographic systems
- Runtime enforcement or monitoring mechanisms
- Compliance regimes, certification programs, or audit requirements
- Legal adjudication, liability assignment, or dispute resolution processes
- Regulatory authority, policy mandates, or governance structures

Any interpretation of this document as proposing or endorsing operational, technical, or regulatory solutions is outside its intended scope.

Interpretive Disclaimer

The identity and naming constructs described in this document are conceptual reference primitives. They are provided for the purpose of analysis, reasoning, and structural clarity. They do not constitute normative specifications, technical standards, policy directives, or implementation requirements.

This document does not assert authority over artificial intelligence systems, nor does it seek to regulate, enforce, or control their behavior. It does not offer legal analysis, conflict-of-law guidance, or jurisdictional resolution mechanisms.

References to jurisdiction, sovereignty, or sectoral context are descriptive, not prescriptive, and are intended solely to clarify how identity may be reasoned about in complex, multi-domain environments.

Forward Compatibility Statement

Where this document references potential downstream constructs, such as identity credentials, operational boundaries, or illustrative downstream constructs, such references are illustrative only. They do not represent proposals, requirements, or claims of ownership by the OSAI Framework.

Their inclusion is intended to demonstrate how identity primitives may support future institutional reasoning once identity exists, not to define how such mechanisms should be designed, implemented, or enforced.

2. Problem Statement: The Identity Gap in Artificial Intelligence Systems

2.1 The Pre-Identity Assumption in Digital Systems

Most contemporary digital systems are designed around an implicit assumption: that meaningful action originates from a human principal operating within a bounded organizational and jurisdictional context. Identity frameworks, access controls, audit mechanisms, and governance processes are therefore structured to associate actions with identifiable human actors, roles, or institutions.

Artificial intelligence systems increasingly violate this assumption.

AI systems may initiate actions without direct human invocation, operate continuously across time, interact with multiple systems simultaneously, and function across organizational and geographic boundaries. In such environments, identity models that presume a human origin or a single institutional boundary become insufficient for clear attribution and reasoning.

As described in the OSAI Preface, identity is treated here as a precondition for governance, rather than an outcome of control mechanisms.

The result is not a failure of existing systems, but a mismatch between emergent system behavior and foundational identity assumptions.

2.2 The Distinction Between Capability and Identity

Current discourse around artificial intelligence frequently focuses on capability: what systems can do, how autonomous they are, and how their performance compares to human or institutional benchmarks. While capability is relevant to safety, efficiency, and governance, it does not resolve the more basic question of what a system is within a broader socio-technical environment.

Identity is distinct from capability.

An AI system's identity does not describe how it performs tasks, but rather how it is referenced, contextualized, and attributed within and across systems. Without a stable identity construct, it becomes difficult to reason about responsibility, jurisdictional applicability, or sectoral relevance, regardless of the system's technical sophistication.

This distinction is critical: governance mechanisms that attempt to manage AI systems solely through capability-based controls risk overlooking the structural ambiguity created by absent or inconsistent identity definitions.

2.3 Attribution Without Identity

Attribution is commonly invoked as a requirement for accountability, auditability, and trust. However, attribution presupposes the existence of an identifiable subject to which actions can be ascribed.

In the absence of clear identity primitives, attribution becomes indirect or inferred, often relying on proxies such as:

- the deploying organization,
- the hosting infrastructure,
- the developer of a model,
- or the human operator most closely associated with a system's output.

These proxies may be useful in specific contexts, but they are not substitutes for identity. They can obscure distinctions between origination, operation, delegation, and responsibility, particularly when AI systems interact across institutional or sectoral boundaries.

The identity gap therefore manifests not as a lack of accountability mechanisms, but as uncertainty about what entity those mechanisms should apply to.

2.4 Jurisdictional and Sectoral Ambiguity

Artificial intelligence systems are increasingly deployed in contexts where jurisdictional and sectoral boundaries overlap or conflict. A single system may be developed in one jurisdiction, deployed in another, operate across multiple sectors, and affect stakeholders subject to different regulatory regimes.

Without an identity framework capable of anchoring AI systems to jurisdictional and sectoral reference points, such complexity is addressed reactively through case-by-case interpretation or post-hoc analysis rather than structurally.

This ambiguity complicates:

- regulatory reasoning,
- insurance and risk assessment,

- cross-border coordination,
- and institutional accountability.

The absence of identity does not prevent action, but it makes consistent interpretation difficult.

2.5 The Need for an Upstream Identity Layer

The challenges described above are often approached through downstream mechanisms such as security controls, access management, or governance processes. While necessary, such mechanisms presuppose that the entities they govern are already identifiable, nameable, and contextualized.

Consistent with the positioning established in the OSAI Abstract, this document treats identity as an upstream reference layer, logically prior to security, governance, and enforcement mechanisms.

Such a layer does not enforce behavior or assign liability. Instead, it provides the referential clarity necessary for other systems to function coherently across institutional, sectoral, and jurisdictional boundaries.

2.6 Framing the Remainder of This Document

The remainder of this document addresses the identity gap by defining foundational identity and naming primitives for artificial intelligence systems. These primitives are intended to clarify how AI systems may be referenced and contextualized prior to the application of operational, security, or governance mechanisms.

They are not presented as requirements or standards, but as conceptual constructs designed to support clearer reasoning, consistent with the non-operational posture established in Layer-1 materials.

Figure 1: Conceptual Positioning of AI Identity and Naming

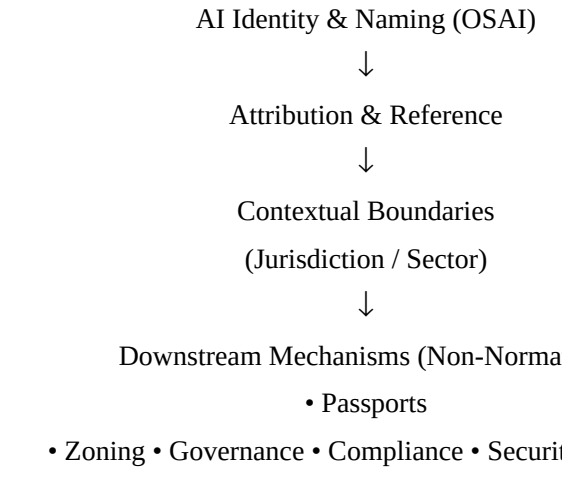

3. Foundational Definitions

This section establishes the foundational terms used throughout this document. These definitions are intended to support consistent reasoning and reference across institutional, sectoral, and jurisdictional contexts. They are descriptive in nature and should not be interpreted as technical specifications, regulatory classifications, or operational requirements. These definitions are informed by the identity gap described in Section 2 and are consistent with the non-operational, upstream posture established in the OSAI Abstract and Preface.

Where terms have common usage in other domains, they are defined here only as they apply within the OSAI Framework.

3.1 Artificial Intelligence System

For the purposes of this document, an Artificial Intelligence System is defined as:

A non-human computational system that performs tasks involving inference, decision support, or action generation, and that may operate autonomously or semi-autonomously within or across digital environments.

This definition is intentionally broad. It does not distinguish between models, agents, services, or composite systems, nor does it require a particular level of autonomy or technical architecture. The purpose of this definition is not classification by capability, but inclusion within a shared referential category for identity reasoning.

3.2 Identity (OSAI Context)

Within the OSAI Framework, identity refers to:

A stable, descriptive reference construct used to distinguish one artificial intelligence system from another across time, context, and environment.

Identity, in this context:

- Is independent of runtime state
- Is independent of access credentials
- Is independent of enforcement mechanisms
- Persists across deployments, interactions, and operational boundaries

Identity does not describe how a system behaves, performs, or is controlled. It describes how a system is recognized and referenced within a broader socio-technical environment.

3.3 Naming

Naming refers to the structured assignment of a unique, persistent identifier to an artificial intelligence system.

Within this framework, naming serves three primary functions:

- Distinction between systems
- Continuity across system lifecycles
- Support for attribution and reference

Naming is a prerequisite for identity but is not equivalent to identity itself. A name provides a reference point; identity provides contextual meaning. The specific format, syntax, or governance of names is outside the scope of this document.

3.4 Attribution

Attribution refers to the ability to associate actions, outputs, or effects with a specific artificial intelligence system through its identity.

Attribution, as used here:

- Is descriptive rather than punitive
- Does not imply fault, liability, or intent
- Does not assign responsibility

Attribution enables institutions to reason about which system acted, without determining who is accountable or what consequences should follow. Those determinations are explicitly downstream of identity.

3.5 Authority of Origination

Authority of origination refers to the human or institutional source under whose auspices an artificial intelligence system was initially created or authorized to exist.

This concept is distinct from:

- Operational control
- Deployment authority
- Runtime supervision

Authority of origination provides historical and institutional context for identity without asserting ongoing control or responsibility. It exists to clarify lineage, not governance.

3.6 Jurisdictional Anchor

A jurisdictional anchor is a reference point that associates an artificial intelligence system’s identity with one or more legal or geographic jurisdictions.

Jurisdictional anchoring:

- Is descriptive, not determinative
- Does not resolve conflicts of law
- Does not imply exclusive authority

An artificial intelligence system may have multiple jurisdictional anchors over its lifecycle. The purpose of anchoring is to support reasoning and interpretation, not to assert legal primacy.

3.7 Sectoral Classification

Sectoral classification refers to the contextual association of an artificial intelligence system with one or more sectors of activity (e.g., healthcare, finance, energy, transportation).

Within this framework:

- Sector refers to institutional and regulatory context, not use case
- Sectoral classification may change over time
- Multiple sectoral classifications may apply simultaneously

Sectoral classification is intended to reduce ambiguity in governance reasoning by providing contextual grounding for identity, not to impose regulatory categorization.

3.8 Sovereignty Boundary

A sovereignty boundary delineates the conceptual separation between:

- The identity of an artificial intelligence system
- The identities of human operators or institutions
- The authority of jurisdictions or governing bodies

Sovereignty boundaries clarify where identity ends and governance begins. They do not assert control, ownership, or enforcement authority. Their function is to prevent conflation of system identity with institutional power or legal responsibility.

3.9 Relationship Among Definitions

The terms defined in this section are interdependent but non-hierarchical. Identity provides the foundational reference. Naming supports identity. Attribution relies on both. Jurisdictional anchoring, sectoral classification, and sovereignty boundaries provide contextual structure. This relationship reflects the upstream identity framing described in Section 2 and does not imply sequencing, enforcement, or governance authority.

None of these constructs imply operational control, enforcement, or compliance. They exist to enable consistent reasoning in environments where artificial intelligence systems operate across traditional institutional and jurisdictional boundaries.

4. Identity and Naming Primitives

This section defines the identity and naming primitives referenced throughout this document. These primitives provide a stable conceptual foundation for reasoning about artificial intelligence systems in environments where traditional identity assumptions no longer hold, as described in Section 2.

Consistent with the non-operational posture established in the OSAI Abstract and Preface, and with the definitions set forth in Section 3, the primitives described here are descriptive reference constructs. Unless otherwise noted, they are non-operational, non-normative, and do not prescribe technical implementation, governance processes, or enforcement mechanisms. References to what a primitive “enables” or “supports” refer solely to conceptual reasoning and reference clarity.

4.1 Purpose of Identity and Naming Primitives

Identity and naming primitives establish referential clarity: the ability to reason consistently about what an artificial intelligence system is prior to determining how it may be managed.

As discussed in Section 2, many downstream challenges attributed to artificial intelligence arise not from insufficient control mechanisms, but from ambiguity about the entities those mechanisms are intended to govern. Identity and naming primitives address this ambiguity by defining the minimum structural elements required for attribution, contextualization, and cross-domain reasoning.

4.2 Primitive: System Identity

System identity is the persistent reference construct that distinguishes one artificial intelligence system from another, as defined in Section 3.2.

System identity:

- Exists independently of runtime state or deployment environment
- Persists across system interactions and lifecycles
- Does not encode permissions, behavior, or capability

By providing continuity across changing operational contexts, system identity establishes a stable subject for reasoning and attribution, addressing the identity gap described in Sections 2.1 and 2.3.

4.3 Primitive: Naming

Naming operationalizes system identity by providing a structured, unique identifier, as defined in Section 3.3.

Naming:

- Distinguishes systems from one another
- Enables continuity across time and context
- Serves as a prerequisite for attribution

Naming provides a reference point upon which contextual reasoning may rely. The format, syntax, governance, and lifecycle management of names remain explicitly outside the scope of this document, consistent with the upstream framing described in Section 2.5.

4.4 Primitive: Authority of Origination

Authority of origination, defined in Section 3.5, associates an artificial intelligence system with the human or institutional source under whose auspices it was initially authorized to exist.

This primitive provides historical and institutional context for identity. It clarifies lineage without implying ongoing control, deployment authority, or operational responsibility, supporting attribution reasoning without collapsing identity into governance.

4.5 Primitive: Jurisdictional Anchor

A jurisdictional anchor, as defined in Section 3.6, associates an artificial intelligence system’s identity with one or more legal or geographic reference points.

Jurisdictional anchoring supports jurisdictional reasoning in multi-domain environments, as discussed in Section 2.4. It is descriptive rather than determinative and does not resolve conflicts of law or assert legal authority. Anchors may change over time as systems evolve.

4.6 Primitive: Sectoral Classification

Sectoral classification, defined in Section 3.7, associates an artificial intelligence system with one or more institutional or regulatory domains relevant to its operation.

Sectoral classification:

- Provides contextual grounding for governance reasoning
- Is independent of specific use cases
- May be plural and dynamic over a system’s lifecycle

This primitive addresses sectoral ambiguity identified in Section 2.4 by enabling contextual interpretation without prescribing regulatory categorization.

4.7 Primitive: Sovereignty Boundary

A sovereignty boundary, defined in Section 3.8, delineates the conceptual separation between:

- System identity
- Institutional authority
- Jurisdictional governance

Sovereignty boundaries prevent conflation of system identity with legal power or institutional control. They clarify where identity ends and governance begins without asserting ownership, enforcement authority, or regulatory mandate. This distinction reinforces the upstream positioning described in Section 2.5.

4.8 Interdependence of Primitives

The primitives described in this section are interrelated but non-hierarchical. System identity provides the foundational reference. Naming supports identity. Authority of origination, jurisdictional anchoring, sectoral classification, and sovereignty boundaries provide contextual structure.

As noted in Section 3.9, these relationships do not imply sequencing, workflow, or enforcement authority. They exist to support coherent reasoning across institutional, sectoral, and jurisdictional boundaries prior to the application of downstream mechanisms.

4.9 Interpretive Boundary

The identity and naming primitives defined in this section are conceptual constructs. They are not requirements, standards, or implementation guidance. Their purpose is to clarify how artificial intelligence systems may be referenced and reasoned about in complex environments, consistent with the non-operational posture established in Layer-1 materials.

5. Jurisdictional Scope and Multi-Domain Operation

This section examines how artificial intelligence systems operate across jurisdictional and institutional boundaries and why identity and naming primitives are necessary to reason about such environments coherently. The discussion is descriptive and does not offer legal analysis, conflict-of-law guidance, or jurisdictional resolution mechanisms.

Consistent with the non-operational posture established in the OSAI Abstract and Preface, and with the identity framing set forth in Sections 2 through 4, jurisdiction is treated here as a contextual reference, not as an assertion of authority or control.

5.1 The Limits of Single-Jurisdiction Assumptions

Many existing governance and compliance frameworks assume that a system:

- is developed within a single jurisdiction,
- is deployed within a bounded legal environment,
- and operates under a coherent regulatory regime.

Artificial intelligence systems increasingly violate these assumptions.

An AI system may be created in one jurisdiction, trained using data sourced from multiple regions, deployed by an institution operating globally, and produce effects subject to diverse regulatory expectations. In such cases, attempts to assign jurisdiction reactively based on infrastructure location, operator domicile, or contractual arrangements often obscure rather than clarify responsibility.

Jurisdictional relevance should not be inferred solely from infrastructure location, data residency, or hosting arrangements.

This challenge reflects the identity gap described in Section 2.4: without a stable referential identity, jurisdictional reasoning becomes fragmented and inconsistent.

5.2 Jurisdiction as Context, Not Control

Within the OSAI Framework, jurisdiction is treated as a descriptive context associated with system identity, rather than as a controlling or determinative attribute.

As described in Section 4.5, a jurisdictional anchor provides a reference point that supports reasoning about applicability, relevance, and oversight. It does not:

- assert legal authority,
- resolve jurisdictional conflicts,
- or establish compliance obligations.

Jurisdictional context may include state, supranational, sectoral, or institutional authorities, depending on the environment in which a system operates.

This distinction is critical. Treating jurisdiction as context allows multiple jurisdictional considerations to coexist without forcing premature resolution or hierarchical ordering.

5.3 Multi-Jurisdictional Identity Over the System Lifecycle

Artificial intelligence systems may acquire, lose, or modify jurisdictional anchors over time. Such changes may occur due to:

- redeployment across regions,
- transfer of operational control,
- modification of system scope,
- or integration into new institutional environments.

Identity and naming primitives, as defined in Sections 3 and 4, enable continuity across these transitions. By preserving a stable system identity independent of deployment context, institutions can reason about jurisdictional relevance without conflating identity with location or control.

This lifecycle perspective reinforces the upstream positioning described in Section 2.5, ensuring that identity remains referential even as operational contexts evolve.

5.4 Jurisdictional Overlap and Interpretive Ambiguity

In multi-domain environments, it is common for multiple jurisdictions to assert relevance over a single AI system. Such overlap may arise from:

- data protection regimes,
- sector-specific regulations,
- cross-border service provision,
- or institutional mandates.

The OSAI Framework does not attempt to resolve such overlaps. Instead, it provides identity constructs that allow these claims to be recognized, referenced, and analyzed without collapsing into singular or simplified interpretations.

This approach supports institutional reasoning by making jurisdictional complexity explicit rather than implicit.

5.5 Relationship Between Jurisdiction and Sovereignty Boundaries

As described in Section 4.7, sovereignty boundaries delineate the conceptual separation between system identity and governance authority. In jurisdictional contexts, these boundaries serve to prevent the conflation of:

- an AI system’s identity,
- the legal authority of institutions,
- and the regulatory reach of jurisdictions.

By maintaining this separation, identity primitives enable clearer reasoning about where jurisdictional considerations apply without asserting control or ownership over the system itself.

5.6 Implications for Downstream Reasoning (Non-Prescriptive)

While this document does not prescribe governance mechanisms, the presence of clear jurisdictional references at the identity layer supports downstream reasoning in areas such as:

- regulatory interpretation,
- insurance and risk assessment,
- cross-border coordination,
- and institutional accountability.

These implications are descriptive, not directive. They illustrate how identity clarity can support consistent reasoning without dictating outcomes or processes.

5.7 Interpretive Boundary

This section does not provide legal advice, jurisdictional determinations, or conflict-of-law analysis. References to jurisdiction are intended solely to clarify how identity and naming primitives may support reasoning in multi-domain environments.

Any interpretation of this section as proposing jurisdictional authority, regulatory compliance frameworks, or enforcement mechanisms is outside its intended scope.

6. Sectoral Taxonomy as a Stabilizing Context

This section examines the role of sectoral taxonomy in providing contextual stability for artificial intelligence system identity. Sectoral taxonomy is treated as a descriptive reference construct that supports reasoning about governance, risk, and applicability without prescribing regulatory classification or enforcement.

In the context of artificial intelligence systems, sectoral context often determines the interpretive, ethical, and regulatory relevance of system behavior more than technical capability alone. Consistent with the non-operational posture established in the OSAI Abstract and Preface, and with the identity framing set forth in Sections 2 through 5, sectoral taxonomy is presented here as a contextual stabilizer, not as a mechanism of control.

6.1 The Limits of Capability-Centric Classification

Contemporary discussions of artificial intelligence frequently classify systems by capability, model type, or degree of autonomy. While such classifications may be useful for technical evaluation, they do not provide sufficient context for institutional or regulatory reasoning.

Capability-centric classification:

- changes as systems evolve,
- varies across deployment contexts,
- and often obscures institutional relevance.

As described in Section 2.2, identity is distinct from capability. Sectoral taxonomy addresses this distinction by anchoring identity in institutional context, rather than technical performance.

6.2 Sector as Institutional Context

Within the OSAI Framework, a sector refers to an institutional and regulatory domain in which an artificial intelligence system operates or produces effects. Examples may include healthcare, finance, energy, transportation, or other domains defined by shared governance, oversight, or societal impact.

Sectoral context:

- is independent of specific use cases,
- reflects institutional relevance rather than technical function,
- and may shape how downstream actors interpret risk, oversight, or applicability.

Sectoral context should not be conflated with individual use cases, applications, or workflows, which may vary widely within the same sector.

As with jurisdiction, sector is treated as context rather than control.

6.3 Sectoral Classification and Identity Stability

Artificial intelligence systems may be repurposed, redeployed, or integrated into environments spanning multiple sectors. In such cases, identity anchored solely to capability or deployment context becomes unstable.

Sectoral classification, as defined in Section 3.7, supports identity stability by:

- providing a consistent reference frame for interpretation,
- allowing comparison across systems operating in similar institutional domains,
- and reducing ambiguity when systems cross organizational or technical boundaries.

This stabilizing effect aligns with the upstream identity layer described in Section 2.5.

6.4 Multi-Sector Operation and Overlap

It is increasingly common for artificial intelligence systems to operate across multiple sectors simultaneously or sequentially. Such overlap may arise from:

- shared infrastructure,
- integrated service models,
- or cross-sector institutional mandates.

The OSAI Framework does not require exclusive sectoral classification. Multiple sectoral contexts may apply concurrently, and sectoral relevance may change over time.

Sectoral taxonomy, as used here, enables these overlaps to be explicitly represented, rather than implicitly assumed or ignored.

6.5 Relationship Between Sectoral Classification and Jurisdiction

Sectoral and jurisdictional contexts are related but distinct. Jurisdiction anchors identity to legal or geographic reference points, as discussed in Section 5. Sectoral taxonomy anchors identity to institutional and regulatory domains.

Maintaining this distinction prevents conflation between:

- where an AI system operates (jurisdiction),
- and the institutional context in which its operation is evaluated (sector).

Together, these contexts support more coherent reasoning without asserting authority or prescribing outcomes.

6.6 Sectoral Taxonomy and Governance Reasoning (Non-Prescriptive)

Clear sectoral context at the identity layer allows institutions to:

- distinguish between systems operating in materially different domains,
- recognize domain-specific considerations without embedding them into identity,
- and assess relevance without relying on ad hoc interpretation.

These observations are descriptive. This document does not prescribe how sectoral classifications should be defined, governed, or enforced.

6.7 Interpretive Boundary

Sectoral taxonomy, as defined in this section, is not a regulatory classification scheme, certification framework, or compliance mechanism. References to sectors are intended solely to clarify how identity and naming primitives may support reasoning in institutional contexts.

Any interpretation of this section as proposing sector-specific requirements, obligations, or governance structures is outside its intended scope.

7. Relationship to Existing Systems (Non-Integration Statement)

This section clarifies the relationship between the OSAI Framework and existing technical, organizational, and regulatory systems. The purpose of this section is to prevent misinterpretation of OSAI as an implementation framework, integration layer, or operational dependency.

Consistent with the non-operational posture established in the OSAI Abstract and Preface, and with the identity framing set forth in Sections 2 through 6, OSAI is positioned strictly upstream of existing systems. In this context, “upstream” refers to conceptual reasoning and reference clarity that precedes technical design, governance frameworks, and enforcement mechanisms.

7.1 Non-Integration with Technical Systems

The OSAI Framework does not integrate with, depend upon, or prescribe interaction with any technical system, including but not limited to:

- identity and access management (IAM) platforms,
- security tooling or control planes,
- authentication or authorization systems,
- model orchestration or deployment platforms,
- data governance or compliance tools.

OSAI does not define APIs, schemas, interfaces, or protocols. It does not require implementation within software systems and does not assume the existence of any particular technical architecture.

The absence of integration is intentional and fundamental to the framework’s scope, not a temporary or transitional state.

7.2 Distinction from Identity and Access Management (IAM)

OSAI should not be interpreted as an extension of, replacement for, or enhancement to identity and access management systems.

IAM systems address:

- access control,
- authentication,
- authorization,
- and operational enforcement.

By contrast, the OSAI Framework addresses:

- referential identity,
- attribution context,
- jurisdictional and sectoral grounding,
- and sovereignty boundaries.

These concerns operate at different conceptual layers. OSAI exists prior to, and independently from, access control or enforcement decisions.

7.3 Distinction from Security, Risk, and Compliance Frameworks

The OSAI Framework is not a security standard, risk management framework, or compliance mechanism. It does not:

- assess threats or vulnerabilities,
- define control requirements,
- establish risk thresholds,
- or prescribe mitigation strategies.

While identity clarity may support downstream reasoning in security, risk, or compliance contexts, such applications are outside the scope of this document and are not implied by its contents.

7.4 Distinction from Regulatory or Standards Bodies

OSAI is not issued by, affiliated with, or endorsed by any regulatory authority, standards organization, or governmental body. It does not assert normative authority or seek to define requirements, obligations, or certifications.

This document is intended as a neutral reference framework that may inform future discussion, analysis, or standardization efforts without asserting precedence or authority over them.

7.5 Relationship to Future Extensions (Non-Binding)

Concepts such as jurisdictional credentials, digital zoning constructs, or operational boundary mechanisms may be discussed in future appendices or related documents. Such discussions, if present, should be understood as illustrative extensions rather than requirements or mandates.

Nothing in this section or elsewhere in this document obligates adoption, implementation, or alignment with any future construct.

7.6 Interpretive Boundary

This section exists to establish clear boundaries between the OSAI Framework and existing systems. Any interpretation of OSAI as an integration layer, technical dependency, governance mechanism, or enforcement framework is inconsistent with its intended scope.

8. Scope, Non-Scope, and Explicit Limitations

This section formally defines the scope and limitations of the OSAI Framework. It consolidates boundaries established implicitly throughout prior sections and makes explicit what this document does and does not address.

The purpose of this section is not to narrow discussion, but to ensure interpretive stability across institutional, regulatory, and academic contexts.

8.1 Scope of the OSAI Framework

The OSAI Framework is limited to the conceptual definition of identity and naming primitives for artificial intelligence systems.

Within scope are:

- referential identity constructs for artificial intelligence systems;
- naming as a prerequisite for identity and attribution;
- attribution as a descriptive association between system identity and system action;
- jurisdictional anchoring as contextual reference;
- sectoral taxonomy as institutional context;
- sovereignty boundaries as conceptual separation between identity and authority.

These elements exist solely to support consistent reasoning about artificial intelligence systems across domains and over time.

For clarity, this framework does not attempt to distinguish artificial intelligence systems from adjacent forms of software or automation, except insofar as such systems require identity reasoning beyond traditional software constructs.

8.2 Explicit Non-Scope

The OSAI Framework explicitly excludes the following:

- technical implementation guidance or system design;
- software architecture, APIs, schemas, protocols, or interfaces;
- identity and access management (IAM) mechanisms;
- authentication, authorization, or access control systems;
- security, risk management, or compliance frameworks;
- governance models, regulatory requirements, or enforcement mechanisms;
- liability allocation, accountability determination, or legal responsibility;
- certification, accreditation, or audit processes;
- operational policies, procedures, or workflows.

The exclusion of these areas is intentional and foundational to the framework’s purpose.

8.3 No Normative or Prescriptive Authority

This document does not assert normative authority. It does not define standards, requirements, or obligations, and it does not prescribe actions for institutions, regulators, or system developers.

Any use of the terms “should,” “may,” or “can” within this document is descriptive and interpretive only, and should not be construed as directive or prescriptive. Nothing in this document should be interpreted as establishing best practices, recommendations, or preferred approaches.

8.4 No Operational or Enforcement Intent

The OSAI Framework is not designed to be operationalized, enforced, or implemented directly. It does not define compliance pathways, enforcement models, or operational controls.

Nothing in this document implies:

- mandatory adoption;
- regulatory endorsement;
- institutional obligation;
- or system-level implementation.

Any downstream use of this framework for operational or governance purposes occurs outside the scope of this document and is not implied by its contents.

8.5 Assumptions and Preconditions

This framework assumes:

- the continued deployment of artificial intelligence systems across multiple jurisdictions and sectors;
- the presence of institutional actors who must reason about such systems without shared identity constructs;
- the absence of a universally accepted identity and naming reference layer for artificial intelligence at the time of writing.

This document does not assume consensus, adoption, or alignment by any particular stakeholder group.

8.6 Known Limitations

The OSAI Framework does not attempt to:

- resolve jurisdictional conflicts;
- harmonize regulatory regimes;
- define sectoral boundaries;
- address technical performance or capability evaluation;
- predict future governance models or institutional outcomes.

These limitations are deliberate and reflect the framework’s upstream positioning.

8.7 Interpretive Boundary

This document should be read as a reference framework only. Any interpretation that treats OSAI as a standard, implementation guide, governance model, or enforcement mechanism is inconsistent with its stated scope.

The absence of operational detail is not a deficiency, but a defining characteristic of the framework.

8.8 Status of the Document

This document is issued as a conceptual reference intended for scholarly, institutional, and policy discussion. It does not represent policy, regulation, or institutional position.

Future documents may expand upon or reference this framework. Such documents, if produced, should be interpreted independently and without retroactive effect on this document’s scope or intent.

Appendix A - Relationship to Companion Bridge Artifacts

This appendix identifies the relationship between the Layer 2 White Paper and later OSAI companion artifacts. The listed artifacts provide controlled context for the white paper without modifying, replacing, expanding, or retroactively altering its conceptual framework.

OSAI-FW-MEM-01

Standards-Facing Memo. Introduces the upstream separation principle and public standards-facing posture.

OSAI-FW-TCN-01

Technical Concept Note. Bridges the conceptual framework into canonical reference identity language supported by companion technical artifacts.

OSAI-FW-GRM-01

Canonical Identifier Grammar. Defines the narrow OSAI ID grammar used for stable referenceability.

OSAI-FW-SCH-01

Three-Record Schema Template Pack. Defines Canonical Identity Record, Authority Binding Record, and Attestation Record structures.

OSAI-FW-RSL-01

Minimal Resolver Illustration. Shows lookup-only retrieval posture for linked reference artifacts.

OSAI-FW-DMO-01

Narrow Demonstration. Demonstrates stable identity with changing authority in a single controlled example.

OSAI-FW-DIF-01

Differentiation Note. Distinguishes OSAI canonical reference identity from downstream authentication, authorization, protocol, certification, runtime-control, and procurement-adjacent systems.

OSAI-FW-REL-01

Companion Artifact Index and Release Note. Identifies the controlled release package and publication-conversion conditions.

OSAI-FW-NVP-01

Namespace and Vocabulary Posture Note. Clarifies namespace posture, issuer-scoped identity, collision handling, controlled vocabulary, and registry non-operation boundaries.

Interpretive Note

The companion artifacts are not incorporated into the body of the white paper. They should be read as later contextual and bridge materials. The white paper remains the conceptual reference root for the Layer 2 series.

Appendix B - Standards-Facing / Public Comment Context

Related OSAI materials were submitted through public comment channels, including the CAISI/NIST RFI process. This statement is included only as public-record context for the document family.

Nothing in this appendix should be interpreted as endorsement, approval, adoption, validation, review, affiliation, ranking, sponsorship, or authorization by NIST, CAISI, or any governmental body.

This appendix does not alter the white paper's conceptual scope, non-scope, interpretive boundaries, or non-operational posture.

References / Related Materials

OSAI-FW-ABS-01
OSAI Framework: Abstract and Preface
https://www.InfrastructureOSAI.com/assets/OSAI-FW-ABS-01.pdf

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

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

OSAI-FW-MEM-01
OSAI: A Neutral Naming and Reference Substrate for Software and AI Agents
https://www.InfrastructureOSAI.com/assets/OSAI-FW-MEM-01.pdf

OSAI-FW-TCN-01
OSAI: Technical Concept Note for Canonical Reference Identity
https://www.InfrastructureOSAI.com/assets/OSAI-FW-TCN-01.pdf

OSAI-FW-GRM-01
OSAI: Canonical Identifier Grammar for Software and AI Agents
https://www.InfrastructureOSAI.com/assets/OSAI-FW-GRM-01.pdf

OSAI-FW-SCH-01
OSAI: Three-Record Schema Template Pack for Software and AI Agents
https://www.InfrastructureOSAI.com/assets/OSAI-FW-SCH-01.pdf

OSAI-FW-RSL-01
OSAI: Minimal Resolver Illustration for Software and AI Agents
https://www.InfrastructureOSAI.com/assets/OSAI-FW-RSL-01.pdf

OSAI-FW-DMO-01
OSAI: Narrow Demonstration of Stable Identity with Changing Authority
https://www.InfrastructureOSAI.com/assets/OSAI-FW-DMO-01.pdf

OSAI-FW-DIF-01
OSAI: Differentiation Note for Canonical Reference Identity and Downstream Agent Systems
https://www.InfrastructureOSAI.com/assets/OSAI-FW-DIF-01.pdf

OSAI-FW-REL-01
OSAI: Companion Artifact Index and Release Note
https://www.InfrastructureOSAI.com/assets/OSAI-FW-REL-01.pdf

OSAI-FW-NVP-01
OSAI: Namespace and Vocabulary Posture Note
https://www.InfrastructureOSAI.com/assets/OSAI-FW-NVP-01.pdf

Revision History

Version	Date	Status	Description
v1.0	June 8, 2026	Issued	Initial public issue of the OSAI Framework Layer 2 conceptual reference white paper, incorporating document-control front matter, companion-material references, contemporary context note, standards-facing public-comment context, references, and revision history.