OSAI Framework

Abstract and Preface — Institutional Reference Edition

Document Status: Reference · First Published: February 4, 2026 · Author: Austin Hall

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 supplying the identity primitives upon which such determinations depend.

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.

Preface — Identity as a Precondition for Artificial Intelligence

Artificial intelligence has entered operational environments more rapidly than the institutional, legal, and technical frameworks governing those environments were designed to accommodate.

Across sectors with material consequence, artificial intelligence systems increasingly influence decisions that affect safety, capital allocation, access to services, and public trust. Despite this expanded role, a foundational deficiency remains unresolved: artificial intelligence systems generally lack standardized identity, jurisdictional attribution, and formal attachment to authority or operational scope.

This deficiency is structural rather than conceptual. Organizations deploying artificial intelligence are routinely required to address questions their systems were not designed to resolve, including system classification, jurisdictional applicability, scope of authorization, and accountability assignment.

In the absence of stable identity primitives, institutions rely on documentation, contractual representations, and informal governance practices to manage risk and responsibility. While such mechanisms may suffice at limited scale, they become increasingly fragile as systems operate across organizational boundaries, jurisdictions, and automated workflows.

The OSAI framework proceeds from a foundational assumption: governance, accountability, and enforceable constraint cannot be meaningfully implemented without first establishing reliable system identity. Identity, in this context, denotes the ability to consistently reference, classify, and situate artificial intelligence systems within institutional structures.

OSAI is not an artificial intelligence system, regulatory authority, or compliance regime. It is a naming and identity substrate intended to support downstream governance, policy development, and enforcement without embedding normative decisions within the identity layer itself.

As artificial intelligence systems evolve toward greater autonomy, interoperability, and institutional integration, identity transitions from an implicit assumption to an explicit prerequisite. The purpose of this work is to articulate an identity framework capable of supporting that transition while preserving institutional discretion and authority.