OSAI: Canonical Identifier Grammar for Software and AI Agents

Technical Companion Artifact

Version:	**v1.0**
Status:	Issued
Effective Date:	May 7, 2026
Document Identifier:	OSAI-FW-GRM-01
Publication Classification:	Public (Companion Technical Artifact)
Governing Record:	OSAI Document Control & Revision History (OSAI-FW-L2-DC-01)
Author:	Austin M. Hall
Scope:	Conceptual, non-operational canonical identifier grammar for an upstream canonical reference identity substrate for software and AI agents.
Companion Artifacts:	OSAI-FW-MEM-01; OSAI-FW-SCH-01; OSAI-FW-RSL-01 OSAI-FW-DMO-01; OSAI-FW-RO-02; OSAI-FW-L2-DC-01

1) Purpose

This document defines a narrow upstream artifact: a canonical identifier grammar for software and AI agents expressed within the OSAI naming/reference posture. Its purpose is to make a subject consistently referenceable across downstream systems without collapsing naming, authority, attestation, authentication, authorization, policy, or enforcement into a single construct.

This document does not define an implementation standard, credential profile, policy engine, or registry architecture. It defines only the canonical form and interpretive boundaries of an identifier grammar intended to support stable referenceability across systems and contexts.

2) Scope and Non-Scope

This document applies only to the canonical identifier form used to represent a software or AI-related subject within the OSAI reference posture. It defines syntax, normalization, segment meaning, permanence expectations, and exclusions.

This document does not define:

- authentication mechanisms
- authorization semantics
- permission models
- policy or enforcement behavior
- trust scoring or compliance status
- transport protocols
- registry deployment requirements
- vendor-specific identity models

Consistent with the memo's upstream separation principle, the identifier defined here is intended to coexist with downstream identity, control, and audit systems without replacing them.

3) Canonical Syntax

osai:<jurisdiction>:<sector>:<class>:<token>

This grammar defines a narrow canonical reference form for software and AI-related subjects. The grammar is intentionally compact. Its purpose is not to encode all available metadata, but to provide a stable anchor that downstream systems can cite, map, and bind to separate records.

In canonical form:

- osai is the fixed namespace prefix
- jurisdiction expresses a bounded jurisdictional context
- sector expresses a bounded operational or industry context
- class expresses the subject category

- token expresses a stable subject token assigned within issuer context

The canonical identifier is a naming/reference artifact only. It does not, by itself, assert authority, trust, permission, compliance, attestation, or runtime state.

3.1) Interpretive Rule

The canonical identifier should be read as a structured reference string, not as an executable control object. Its purpose is to make a subject stably citeable across systems, logs, and contexts while leaving downstream control functions separate.

3.2) Fixed Grammar Posture

For v1.0, the canonical grammar is fixed at five segments. Additional metadata, provenance, authority, attestation, or lifecycle detail should be represented in linked records rather than embedded into the identifier itself.

4) Normalization Rules

Canonical OSAI identifiers should conform to the following normalization rules:

1. Canonical form is lowercase only.
2. Canonical form uses ASCII characters only.
3. Segment delimiter is the colon (:).
4. Hyphens may be used within a segment where needed.
5. Spaces are not permitted.
6. Underscores are not permitted in canonical form.
7. Empty segments are not permitted.
8. Each canonical identifier must contain exactly five ordered segments.
9. The canonical identifier must not embed authority state, environment state, attestation state, credential material, or policy state.
10. The canonical identifier must remain interpretable without vendor-specific parsing assumptions.

4.1) Reserved Semantic Exclusions

The following should not be encoded directly into the canonical identifier:

- permission scopes
- role grants
- environment labels such as sandbox or production
- trust levels
- approval status
- attestation method
- credential identifiers
- endpoint addresses
- timestamps

Those semantics belong in linked records, not in the identifier grammar.

4.2) Token Normalization Posture

The token segment should be durable, readable, and structurally conservative. In v1.0, it should:

- begin with a lowercase letter or digit
- contain only lowercase letters, digits, and hyphens
- avoid semantic overload
- avoid embedding dates, approval states, or environment names unless those are truly part of the stable subject identity

A token like grid-balance-001 is preferable to a token like prod-approved-grid-balance-20260413.

5) Segment Definitions

The canonical OSAI identifier is composed of five ordered segments.

Segment	Required	Meaning	Notes
osai	Yes	Fixed namespace prefix	Always the literal osai
jurisdiction	Yes	Jurisdictional context anchor	Geographic, sovereign, or administrative context only
sector	Yes	Operational or industry context anchor	Controlled value; not a policy label
class	Yes	Subject category	Controlled value such as agent, service, workflow
token	Yes	Stable subject token	Issuer-assigned durable reference token

5.1) osai

The first segment is always the fixed namespace prefix osai. It indicates that the identifier is being expressed within the OSAI canonical reference posture.

5.2) jurisdiction

The jurisdiction segment expresses a bounded jurisdictional, sovereign, or administrative context anchor. It is intended to answer a narrow contextual question: within what jurisdictional frame should this subject be interpreted for reference purposes.

This segment does not indicate who currently authorizes the subject. It indicates contextual placement, not delegated authority.

Starter examples for v1.0:

- global
- us
- us-la

5.3) sector

The sector segment expresses a bounded operational or industry context anchor. It exists to preserve contextual interpretability across systems without collapsing sector semantics into policy semantics.

This segment should remain narrow and controlled in v1.0.

Starter examples for v1.0:

- general
- energy
- healthcare
- finance
- logistics
- public-sector

Use of general is acceptable where a narrower sector value is not yet required.

5.4) class

The class segment expresses the subject category represented by the identifier.

Starter values for v1.0:

- agent
- service
- workflow

- model
- dataset
- system

These are reference categories only. They are not trust or privilege categories.

5.5) token

The token segment is a stable issuer-assigned subject token intended to distinguish the referenced subject within the preceding context. It should be durable enough to support lifecycle traceability without embedding transient operating conditions.

Examples:

- research-001
- grid-balance-001
- intake-routing-002

6) Controlled Vocabulary Posture

This grammar permits controlled vocabularies for selected segments, especially jurisdiction, sector, and class. In v1.0, those vocabularies should remain intentionally narrow.

The purpose of this posture is twofold:

1. It stabilizes the grammar before vocabulary expansion begins.
2. It reduces premature complexity while preserving cross-system legibility.

6.1) Vocabulary Governance Principle

Controlled vocabularies may expand in later companion artifacts, but vocabulary growth should not alter the canonical grammar structure itself. Grammar stability takes precedence over vocabulary breadth.

6.2) v1.0 Starter Vocabulary Rule

For v1.0, only starter values needed for the first worked examples and later bridge artifacts should be included. Do not broaden vocabularies merely because future use cases can be imagined.

That restraint increases credibility.

6.3) Non-Scope of This Document

This document defines the grammar and the starter posture for controlled values. It does not define the full authoritative vocabulary set for future public use across all OSAI sectors or jurisdictions.

7) Permanence and Lifecycle Rules

An OSAI canonical identifier is intended to function as a stable reference anchor. For that reason, the following lifecycle rules apply.

7.1) Stability Rule

The same subject should retain the same canonical identifier across environments where the underlying subject remains the same.

A sandbox deployment and a production deployment do not, by themselves, require distinct canonical identifiers if the underlying referenced subject is the same.

7.2) Authority Separation Rule

A change in delegated authority does not, by itself, create a new canonical identifier.

Authority belongs in a separate binding layer. If authority changes, the authority binding record changes. The canonical identity reference does not.

7.3) Attestation Separation Rule

A new attestation does not, by itself, create a new canonical identifier.

Attestation belongs in a separate evidentiary layer. If a new attestation is issued, the attestation record changes. The canonical identity reference does not.

7.4) Lifecycle Expression Rule

Rotation, supersession, retirement, correction, or revocation references should be represented through linked lifecycle semantics rather than identifier churn.

This preserves auditability and longitudinal reference coherence.

7.5) Non-Sufficiency Rule

A canonical identifier is not sufficient by itself to establish:

- permission
- trust
- validity
- compliance
- safety
- authorization

7.6) Replacement Threshold Rule

A new canonical identifier should only be created when the underlying referenced subject is materially different, not merely because a downstream system has changed, a role has changed, or a new evidentiary artifact has been issued.

8) Exclusions and Anti-Drift Notes

An OSAI canonical identifier is not:

- a credential
- a secret
- a permission set
- a role definition
- a trust score
- a runtime endpoint
- a policy object
- an attestation artifact
- a compliance certificate
- an enforcement decision

These exclusions are included to prevent downstream functions from being collapsed into the canonical identifier itself. The identifier is intended to support binding and reference, not to replace the downstream systems that make operational decisions.

9) Illustrative Examples

Example 1

osai:global:general:agent:research-001
Illustrative interpretation: a globally scoped general-purpose agent reference expressed in canonical OSAI form.

Example 2

osai:us-la:energy:agent:grid-balance-001
Illustrative interpretation: an energy-context agent reference anchored to a Louisiana jurisdictional context.

Example 3

osai:us:healthcare:workflow:intake-routing-002
Illustrative interpretation: a U.S.-anchored healthcare workflow reference expressed as a stable canonical subject.

These examples do not express permissions, trust status, runtime environment, or attestation claims. They express only canonical identity reference in structured form.

10) Relationship to Companion Artifacts

This document defines only the identifier grammar.

Subsequent companion artifacts may define:

- field-level record structures for canonical identity, authority binding, and attestation
- minimal resolver behavior for reference lookup
- a narrow demonstration showing stable identity with changing authority contexts

Those later artifacts should remain consistent with the separations preserved here. The grammar is the base layer, not the whole system.

References

OSAI-FW-SCH-01
OSAI: Three-Record Schema Template Pack for Software and AI Agents
Issued Companion Technical Artifact, v1.0, issued May 7, 2026.
https://www.InfrastructureOSAI.com/assets/OSAI-FW-SCH-01.pdf

OSAI-FW-RSL-01
OSAI: Minimal Resolver Illustration for Software and AI Agents
Issued Companion Technical Artifact, v1.0, issued May 7, 2026.
https://www.InfrastructureOSAI.com/assets/OSAI-FW-RSL-01.pdf

OSAI-FW-DMO-01
OSAI: Narrow Demonstration of Stable Identity with Changing Authority
Issued Companion Technical Artifact, v1.0, issued May 7, 2026.
https://www.InfrastructureOSAI.com/assets/OSAI-FW-DMO-01.pdf

OSAI-FW-MEM-01
OSAI: A Neutral Naming and Reference Substrate for Software and AI Agents
Standards-Facing Memo (Public), v1.0, issued March 5, 2026.
https://www.InfrastructureOSAI.com/assets/OSAI-FW-MEM-01.pdf

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

OSAI-FW-L2-DC-01
Document Control & Revision History
Issued publication control record, v1.3, issued May 7, 2026.
https://www.InfrastructureOSAI.com/assets/OSAI-FW-L2-DC-01.pdf