ZLAR Glossary

Curated Terms for the NCCoE Concept Paper Response

Supporting reference for public comment on software and AI agent identity and authorization

This glossary contains terms directly relevant to the NCCoE concept paper's six focus areas. Where a NIST or standards-body term exists for a ZLAR concept, the NIST term leads and the ZLAR term follows in brackets. Terms are organized by category: architecture, enforcement, evidence, identity, standards, and regulatory.

1. Architecture

Data Plane Enforcement Point [Execution Boundary]

NIST (SP 800-207) / ZLAR

The exact point where an AI agent's reasoning becomes a real-world action. The location where the Policy Enforcement Point operates. SP 800-207 distinguishes the control plane (PDP/PE) from the data plane (where the PEP mediates access between subjects and resources).

ZLAR applies the PEP concept to the specific location where an agent's action selection (non-deterministic) meets policy enforcement (deterministic). SP 800-207 was written for network/data access; execution boundary names this stratum for AI agents.

The Governance Primitive

ZLAR

Authorization before irreversible action, with tamper-proof evidence — expressed as a three-property design pattern. Unifies SP 800-207's PEP/PDP model, SP 800-53 AU-10's non-repudiation requirement, and EU AI Act Article 14's human oversight mandate into a single architectural property.

Three properties: (1) Cryptographic evidence — you can prove what happened. (2) Structural enforcement — policy rules cannot be overridden by the agent. (3) Formal verification path via Cedar — you can prove the enforcement rules are correct before deployment (PoC validated).

Access Enforcement [Enforcement Layer]

NIST (SP 800-53 AC-3) / ZLAR

The layer of the architecture that intercepts agent actions before execution and decides whether they proceed. Maps to SP 800-53 AC-3 (Access Enforcement).

ZLAR's enforcement layer contains no machine learning, no natural language processing, and no probabilistic reasoning — it evaluates signed policy deterministically. Two surfaces: bash gate (hook-based) and MCP gate (JSON-RPC proxy), both sharing the same signed policy and audit trail format.

Continuous Monitoring [Observation Layer]

NIST (SP 800-137 / SP 800-53 SI-4) / ZLAR

The layer that reads the evidence trail after the fact, finds patterns, and surfaces them to the human. Maps to SP 800-137 (Information Security Continuous Monitoring) and SI-4 (System Monitoring).

The observation layer is structurally forbidden from enforcement. It produces facts for the human to interpret, not alerts for the system to act on.

Identity-to-Execution Gap

EXT (industry / NCCoE)

The structural gap between agent authentication (knowing who the agent is) and execution governance (controlling what the agent does and proving what happened). Only 28% of organizations can trace agent actions to a human sponsor (CSA/Strata Identity, 2026).

The six NCCoE standards collectively address identity but none provides deterministic per-action enforcement at the execution boundary. ZLAR fills this gap as the PEP/PDP (SP 800-207) applied to agent actions.

---

2. Enforcement

Policy Enforcement Point [Gate]

NIST (SP 800-207) / ZLAR

The non-bypassable component that intercepts every tool call an AI agent makes, evaluates it against a signed policy, and returns allow/deny/ask. SP 800-207 Section 3.1 defines the PEP as “responsible for enabling, monitoring, and eventually terminating connections between a subject and an enterprise resource.”

ZLAR’s PEP contains no machine learning, no natural language processing, and no probabilistic reasoning. It evaluates signed policy deterministically and is not susceptible to prompt-based manipulation.

Policy Decision Point (PDP) / Policy Engine (PE)

NIST (SP 800-207 Section 3.1)

The component responsible for the ultimate decision to grant access. In ZLAR’s current production deployment, the JSON policy evaluator is the PDP — it evaluates signed rules top-to-bottom with deterministic, default-deny semantics.

Cedar is the proven migration path (PoC with 14/14 tests, formal verification via SMT solver), adding formal policy analysis capabilities the current JSON evaluator does not have. Both share the same property: deterministic evaluation that cannot be persuaded.

Signed Authorization Policy [Policy]

NIST (SP 800-207, SP 800-53 SI-7(6)) / ZLAR

The signed, versioned JSON document that defines what the PEP allows, denies, and asks about. Satisfies SI-7(6) (cryptographic protection of information integrity). Ed25519-signed (FIPS 186-5) — agents cannot modify the rules that govern them.

If the signature is invalid, the gate refuses to load the policy and denies everything (fail-closed per SP 800-207 default-deny posture).

Default-Deny Posture / Fail-Secure [Fail-Closed]

NIST (SP 800-207 Section 3, SP 800-53 SC-24) / ZLAR

Any failure in the gate results in denial, never in allowing an action to proceed. Implements the default-deny architectural assumption of SP 800-207’s ZTA model and SC-24 (Fail in Known State).

Five governance-specific failure modes (missing policy, invalid signature, gate crash, human unreachable, unwritable audit trail) — each defaults to deny.

Implicit Deny [Default Action]

NIST (SP 800-207 Section 3, SP 800-53 AC-3) / ZLAR

The policy-level fallback when no rule matches a tool call. Anything not explicitly addressed by a rule is blocked.

Risk Assessment [Risk Profile]

NIST (SP 800-53 RA-3, SP 800-30) / ZLAR

A three-dimensional numeric assessment per policy rule: irreversibility, consequence, and blast_radius (each 0–100). Maps to RA-3 applied at the per-action level.

NIST distinction: SP 800-30 covers organizational risk across the C/I/A triad. ZLAR adds irreversibility as a first-class dimension — the same action can be low-consequence but high-irreversibility.

Least Privilege / Delegation of Authority [Permission Attenuation]

NIST (SP 800-53 AC-6) / ZLAR

The monotonic reduction of capabilities at each step in a delegation chain. Implements AC-6 (Least Privilege) across multi-agent delegation hops.

As of March 2026, zero of eight major multi-agent frameworks implement permission attenuation in delegation chains.

---

3. Evidence & Audit

Cryptographic Audit Trail [Evidence Layer]

NIST (SP 800-53 AU-9(3), AU-10) / ZLAR

Proves what happened — cryptographically, tamper-evidently, with non-repudiable human attribution. Satisfies AU-9(3) (“cryptographic mechanisms to protect the integrity of audit information”) and AU-10 (non-repudiation via digital signatures).

Per-entry FIPS 186-5 EdDSA signatures satisfy AU-10. FIPS 180-4 SHA-256 hash chaining satisfies AU-9(3). As of March 2026, no ratified standard specifies a format for cryptographically verifiable agent action audit trails.

Audit Records with Cryptographic Protection [Audit Trail]

NIST (SP 800-53 AU-2, AU-3, AU-9(3), AU-12) / ZLAR

Satisfies AU-2 (Event Logging), AU-3 (Content of Audit Records — all six required fields present), AU-9(3) (cryptographic integrity protection), and AU-12 (automatic generation at execution boundary).

Tamper-Evident Hash Chain [Hash Chain]

NIST (SP 800-53 AU-9(3), FIPS 180-4) / ZLAR

Each audit entry’s prev_hash field contains the FIPS 180-4 SHA-256 hash of the previous entry. Any modification, insertion, deletion, or reordering invalidates all subsequent entries.

SP 800-92 identifies the need for log integrity but specifies no mechanism. ZLAR provides a specific, deployed implementation.

Non-Repudiation / Decision Attribution [Authorizer]

NIST (SP 800-53 AU-10) / ZLAR

Every audit entry records who or what made the decision. Satisfies AU-10 (Non-repudiation) via per-entry Ed25519 digital signatures. Required at SP 800-53B High baseline.

Distinguishes policy-auto, human-approved (by specific user ID), timeout, rate-limit, and error authorizers in each entry.

Security Attestation [Proof Layer]

NIST (SP 800-53A) / ZLAR

A self-contained, cryptographically sealed attestation bundle for consumption by regulators, insurers, courts, and auditors. An external party can independently verify without access to ZLAR's running infrastructure.

---

4. Cryptography

EdDSA per FIPS 186-5 [Ed25519]

NIST (FIPS 186-5 Section 7, SP 800-186)

NIST-approved digital signature algorithm using the Edwards25519 curve. Used by ZLAR for policy signing and per-entry audit trail integrity. CMVP-validated modules exist (Bouncy Castle BC-FJA 2.1.0, Certificate #4943, FIPS 140-3).

Cite as: "FIPS 186-5-approved EdDSA per Section 7, using the Edwards25519 curve recommended in SP 800-186."

SHA-256 per FIPS 180-4

NIST (FIPS 180-4)

Cryptographic hash function used for the audit trail's hash chain. SHA-256 is quantum-safe — unlike Ed25519 signatures, the hash chain is not vulnerable to quantum attacks.

ML-DSA per FIPS 204 [CRYSTALS-Dilithium]

NIST (FIPS 204)

Post-quantum digital signature algorithm. ZLAR implements ML-DSA-44 (NIST Level 2) alongside Ed25519 with hybrid composite signing mode. Three modes: ed25519 (default), ml-dsa-44, and hybrid (both must verify).

Algorithm choice is configuration, not code — satisfying cryptographic agility per IR 8547.

Cryptographic Agility [PQC Phase 1 → Phase 2]

NIST (IR 8547, SP 800-131A Rev. 2) / ZLAR

The ability to interchange cryptographic algorithms via configuration without code changes. NIST IR 8547 schedules Ed25519 for deprecation after 2030 and prohibition after 2035.

ZLAR's lib/crypto.sh provides algorithm-agnostic signing, verification, and key management across Ed25519, ML-DSA-44, and hybrid modes.

---

5. Identity & Authorization Standards

Zero Trust Architecture [Zero Trust for AI Agents]

NIST (SP 800-207) / ZLAR

Extension of SP 800-207 zero trust principles to agent governance: never trust the agent, always verify against signed policy before every action.

ZLAR implements SP 800-207 at the AI agent execution boundary: the JSON policy evaluator functions as the PDP, the gate as the non-bypassable PEP, and the Ed25519-signed audit trail satisfies Tenet 7 (comprehensive telemetry) and SP 800-137 monitoring requirements.

OAuth 2.0

EXT (IETF) — one of six NCCoE standards

ZLAR operates downstream of OAuth: it assumes OAuth-issued tokens exist and consumes them as identity context, then enforces at the per-action level within the token's authorized scope. OAuth issues tokens at session boundaries; ZLAR enforces at every action within that session. The two are complementary.

SCIM (System for Cross-domain Identity Management)

EXT (IETF) — one of six NCCoE standards

SCIM provisions identities; ZLAR governs actions taken by those identities. ZLAR's registry derives agent inventory from governance contact rather than explicit SCIM provisioning.

NGAC (Next Generation Access Control)

NIST (SP 800-178, NIST IR 7987) — one of six NCCoE standards

A NIST-developed access control standard referenced in the NCCoE concept paper. Cedar implements NGAC's functional objectives with formal verification via Lean theorem prover and faster evaluation than OPA/Rego (Cedar team, OOPSLA 2024).

AuthZEN

EXT (OpenID Foundation)

The Authorization API 1.0 standard defining a Subject/Action/Resource/Context evaluation interface for PEP-to-PDP communication. Implementer's Draft approved by the OpenID Foundation.

ZLAR's gate evaluation maps to AuthZEN's Access Evaluation API. ZLAR could expose an AuthZEN-compatible endpoint, enabling interoperability with any AuthZEN-compliant PEP.

AIMS (Agent Identity Management System)

EXT (IETF)

The conceptual model in draft-klrc-aiagent-auth-00 (March 2026), composing WIMSE, SPIFFE, and OAuth 2.0 into a coherent agent identity framework. AIMS explicitly declares policy format and execution-level enforcement out of scope (Section 12). ZLAR provides the execution-boundary PEP and cryptographic audit trail that AIMS leaves to implementers.

WIMSE (Workload Identity in Multi-System Environments)

EXT (IETF)

Standards for workload identity using URI-based identifiers compatible with SPIFFE IDs. WIMSE answers “what workload is this?” ZLAR answers “what is this workload permitted to do right now?” Together they implement a complete zero trust architecture for AI agents.

Non-Human Identity (NHI)

EXT (industry / NCCoE)

An identity assigned to a non-human entity. 82-to-1 ratio of NHIs to human identities in enterprises (Rubrik Zero Labs). SP 800-63-4 explicitly excludes NHIs from scope. ZLAR addresses NHI governance at the execution boundary.

Cedar

EXT (AWS / CNCF)

A policy language with default-deny semantics and formal verification via SMT solver. CNCF Sandbox project. Used by AWS Bedrock AgentCore.

ZLAR’s current production PDP uses JSON rules; Cedar is the proven migration path (PoC with 14/14 tests) adding formal policy analysis. Cedar’s own benchmarks report 42–80x faster evaluation than OPA/Rego (Cedar team, OOPSLA 2024).

MCP (Model Context Protocol)

EXT (Anthropic / AAIF) — one of six NCCoE standards

An open protocol for connecting AI models to external tools. ZLAR’s MCP gate is a vendor-agnostic TCP proxy that intercepts tools/call JSON-RPC requests. MCP defers authorization entirely to OAuth and provides no per-tool enforcement. ZLAR fills this gap externally.

6. Regulatory & Compliance

SP 800-53 Rev 5 AU Controls

NIST

ZLAR’s audit trail satisfies: AU-2 (Event Logging), AU-3 (Content of Audit Records — all six required fields), AU-9(3) (Cryptographic Protection via EdDSA signatures), AU-10 (Non-repudiation — required at High baseline), AU-11 (Audit Record Retention), AU-12 (Automatic generation at execution boundary).

SR 11-7 / OCC 2011-12 (Model Risk Management)

EXT (Federal Reserve / OCC)

Interagency guidance on model risk management. ZLAR maps to SR 11-7’s three pillars: (1) Model development → policy authoring and signing ceremony; (2) Model validation → evidence trail and witness layer; (3) Governance, policies, and controls → signed policy, fail-closed defaults, human authority chain.

SR 11-7 requires “effective challenge” — ZLAR’s deny path implements effective challenge at the execution boundary with cryptographic recording.

SEC Rule 17a-4 (Audit-Trail Alternative)

EXT (SEC)

The October 2022 amendments explicitly recognize cryptographic methods (hash chains, digital signatures, Merkle trees) as valid compliance mechanisms for broker-dealer recordkeeping. ZLAR's hash-chained, Ed25519-signed JSONL architecture directly satisfies 17a-4(f)(2)(ii)(A) requirements.

FINRA

EXT (US financial regulator)

FINRA's 2026 Annual Regulatory Oversight Report (December 2025) and Regulatory Notice 24-09 (June 2024) establish supervisory expectations for AI agent deployments. FINRA Rule 3110 (Supervision) requires a supervisory system "reasonably designed to achieve compliance" — ZLAR's gate is structurally a supervisory control.

Note: these are supervisory expectations, not rules with the force of law.

NCCoE (National Cybersecurity Center of Excellence)

NIST

NIST's industry collaboration program. The concept paper "Accelerating the Adoption of Software and AI Agent Identity and Authorization" (February 2026) identifies six focus areas: Identification (I1–I4), Authorization (A1–A5), Access Delegation (D1–D4), Logging and Transparency (N1–N4), Tracking Data Flows (T1–T3), Prompt Injection Prevention (P1–P2).

References six standards: MCP, OAuth 2.0, OIDC, SPIFFE/SPIRE, SCIM, NGAC. Comment deadline: April 2, 2026.

OWASP Top 10 for Agentic Applications

EXT (OWASP)

Published December 2025. OWASP recommends "comprehensive signed audit logs of agent actions" as a "non-negotiable security control" — directly validating ZLAR's per-entry Ed25519-signed audit trail.

Prepared for NCCoE concept paper response (comment deadline April 2, 2026). Full glossary and implementation available at https://github.com/ZLAR-AI/ZLAR