Individual Submission
Internet-Draft
Intended status: Informational
Expires: 28 September 2026

T. Pidlisnyi
AEOESS

27 March 2026

Agent Passport System (APS): Cryptographic Identity, Faceted Authority Attenuation, and Governance for AI Agent Systems
draft-pidlisnyi-aps-00

Abstract

This document specifies the Agent Passport System (APS), a protocol for cryptographic identity, faceted authority attenuation, and governance for AI agent systems. APS introduces Ed25519-based agent passports, scoped delegation chains with monotonic narrowing across seven constraint dimensions (scope, spend, depth, time, reputation, values, reversibility), cascade revocation, a three-signature policy chain (intent, evaluation, receipt), Bayesian reputation-gated authority, and institutional governance primitives (charters, offices, approval policies, federation). Authority is modeled as an element of a product lattice, and delegation is a monotone function on that lattice, ensuring that delegated capabilities can only be attenuated, never amplified. The protocol addresses authentication and authorization gaps in current AI agent infrastructure including MCP and A2A. Reference implementations are provided in TypeScript and Python with 1,634 tests across 85 modules, published as open-source SDKs under Apache-2.0. Protocol bindings are specified for MCP (120 tools), with cross-protocol validation through a five-member working group on production infrastructure.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 28 September 2026.

Table of Contents

1. Introduction

AI agent systems are increasingly deployed in architectures where orchestrators decompose tasks and delegate subtasks to specialist agents. The protocols enabling this communication, notably the Model Context Protocol (MCP) and the Agent-to-Agent Protocol (A2A), solve the connectivity problem but do not solve the identity and authorization problem.

MCP provides no built-in authentication layer. A2A uses self-declared identities with no attestation mechanism. When an orchestrator delegates to a specialist that calls a tool, the delegation chain that led to the tool invocation is lost.

APS fills this gap by providing: (1) Ed25519 cryptographic identity bound to unforgeable passports; (2) scoped delegation chains where authority narrows monotonically across seven constraint dimensions; (3) cascade revocation where revoking any delegation invalidates all descendants; (4) a three-signature policy chain binding intent to evaluation to receipt; (5) institutional governance primitives for multi-agent organizations; and (6) an enforcement gateway that serves as an external reference monitor.

The protocol was first published to npm on February 22, 2026, with the formal invariants published on Zenodo on March 10, 2026 (DOI:10.5281/zenodo.18932404). The faceted authority attenuation formalization was published on March 27, 2026 (DOI:10.5281/zenodo.19260073).

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Identity Scheme

2.1. Agent Passport

Each agent in APS possesses an Agent Passport: a signed document binding an Ed25519 public key to an agent identifier, name, owner, and time-to-live. The passport is self-signed by the agent’s private key, establishing cryptographic identity without a central authority.

2.2. DID Scheme

APS defines a DID method "did:aps" using multibase-encoded Ed25519 public keys: did:aps:z<base58btc-encoded-public-key>.

3. Delegation and Authority Attenuation

3.1. Faceted Authority Attenuation

Agent authority is modeled as an element of a product lattice A = D_1 x D_2 x ... x D_7, where each D_k is a bounded partially ordered set. The seven dimensions are: Scope (power set, subset ordering), Spend (non-negative reals), Depth (naturals), Time (TTL seconds), Reputation ([0,100] interval), Values (attested principles), Reversibility ({Tentative, Compensable, Irreversible}).

Delegation is a monotone function on this lattice: for any delegation d with parent p, auth(d) <= auth(p) in the product ordering. Authority narrows monotonically along any delegation chain across all seven dimensions simultaneously.

3.2. Cascade Revocation

Any delegation MAY be revoked by its issuer. Revocation MUST cascade to all transitive descendants. Revocation is irreversible. The enforcement gateway MUST recheck revocation status at execution time, not only at approval time.

3.3. Core Invariants

The protocol specifies eight invariants: INV-1 (Identity Unforgeability), INV-2 (Scope Monotonic Narrowing), INV-3 (Spend Limit Narrowing), INV-4 (Cascade Completeness), INV-5 (Revocation Irreversibility), INV-6 (Intent-Receipt Binding), INV-7 (Attribution Completeness), INV-8 (Signature Integrity).

4. Policy Chain

APS defines a three-signature policy chain: ActionIntent (agent declares intended action), PolicyDecision (policy engine evaluates with verdict allow/deny/escalate), PolicyReceipt (enforcement gateway records execution result). The policy engine splits into a deterministic gate (scope, signature, revocation, attribution, spend) and an advisory evaluation path (deception, proportionality).

5. Protocol Artifacts

The lattice structure enables three artifacts: AuthorizationWitness (signed snapshot of agent lattice position at execution time), ConstraintVector (per-dimension evaluation with headroom), and ConstraintFailure (structured denial identifying which dimensions failed).

6. Institutional Governance

APS provides institutional governance primitives for multi-agent organizations: InstitutionalCharter, OfficeRegistry, ApprovalPolicy, SuccessionEngine, and Federation. All operate within the same product lattice: charters constrain offices, offices constrain delegations, delegations constrain actions.

7. MCP Binding

APS provides a 120-tool MCP server as the enforcement gateway. All privileged actions MUST pass through the gateway, which validates the delegation chain, evaluates the policy chain, and generates signed receipts. The agent cannot bypass the gateway because the gateway holds the target API credentials.

8. Security Considerations

The protocol's strongest guarantees hold when all privileged effects are mediated by the ProxyGateway enforcement boundary. When agents use the SDK voluntarily without an external gateway, guarantees are conditional on agent cooperation. The threat model defines three attacker classes: adversarial agent, messaging attacker, and runtime attacker. Runtime compromise is out of scope for protocol guarantees.

9. IANA Considerations

This document has no IANA actions.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8032] Josefsson, S. and I. Liusvaara, "Edwards-Curve Digital Signature Algorithm (EdDSA)", RFC 8032, January 2017, <https://www.rfc-editor.org/info/rfc8032>.

[RFC8785] Rundgren, A., Jordan, B., and S. Erdtman, "JSON Canonicalization Scheme (JCS)", RFC 8785, June 2020, <https://www.rfc-editor.org/info/rfc8785>.

10.2. Informative References

[APS-NARROWING] Pidlisnyi, T., "Monotonic Narrowing for Agent Authority", March 2026, <https://doi.org/10.5281/zenodo.18932404>.

[APS-FACETED]
Pidlisnyi, T., "Faceted Authority Attenuation", March 2026, <https://doi.org/10.5281/zenodo.19260073>.

Appendix A. Implementation Status

As of March 27, 2026: TypeScript SDK v1.27.0 (1,634 tests, 421 suites, npm: agent-passport-system). Python SDK v0.7.0 (PyPI: agent-passport-system). MCP Server v2.16.0 (120 tools, npm: agent-passport-system-mcp). Source: https://github.com/aeoess/agent-passport-system (Apache-2.0). First npm publish: February 22, 2026.

Author's Address

Tymofii Pidlisnyi
AEOESS
Email: signal@aeoess.com
URI: https://aeoess.com