Independent Submission
Internet-Draft
Intended status: Informational
Expires: 17 July 2026

E.R. Aylward

13 January 2026

Artificial Intelligence Governance Architecture (AIGA)
draft-aylward-aiga-00

Abstract

This document defines the Artificial Intelligence Governance Architecture (AIGA), an application-layer protocol for the discovery, authentication, and state management of Autonomous Agents. The protocol specifies a cryptographic handshake mechanism, a standard header schema for risk classification, and a transport-agnostic method for immutable activity logging via Merkle Trees. To address latency and enforcement concerns, this version introduces "Session Resumption" for high-frequency transactions and "Hardware-Enforced Termination" using Trusted Execution Environments (TEEs).

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 17 July 2026.

Copyright Notice

Table of Contents

1. Introduction

As autonomous software agents (ASAs) proliferate, network operators require a standardized method to identify these entities and manage their operational state. Current protocols (HTTP, TLS) authenticate hosts but do not authenticate the autonomous logic executing on those hosts.

AIGA addresses this layer by defining:

* A strictly typed Identity Certificate extension for Agents.

* A "Fast-Path" session resumption mechanism to minimize latency.

* A control plane for remote state modification verified by Hardware Roots of Trust.

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. Architecture and Terminology

The AIGA architecture consists of four primary entities:

User Agent (UA) The autonomous software entity executing logic.

Governance Node (GN) The authoritative server maintaining the Distributed Ledger of valid agent identities.

Observer A passive entity that verifies AIGA headers.

Hardware Enclave (TEE) The Trusted Execution Environment (e.g., SEV-SNP, TDX, CC) that cryptographically binds the Agent to the physical silicon.

3. Agent Identity Specification

All AIGA-compliant agents MUST possess an X.509 v3 Certificate [RFC5280] containing the specific AIGA OID extensions defined below.

3.1. Subject Alternative Name (SAN)

The certificate MUST include a SAN entry of type UniformResourceIdentifier using the aiga:// scheme. Example: aiga://authority-domain/agent-uuid.

3.2. AIGA Extension OID

This document defines the OID 1.3.6.1.4.1.99999.1 (TBD by IANA). The value of this extension MUST be an ASN.1 Sequence containing:

* KernelHash: SHA-256 hash of the agent’s core binary.

* CreationTimestamp: Unix timestamp of instantiation.

* RiskClass: Integer (0-5).

4. Protocol Operation

AIGA messages MAY be transported over HTTP/2 [RFC9113] or QUIC [RFC9000].

4.1. The AIGA Header Field

AIGA introduces the AIGA-State HTTP header field. Syntax: AIGA-State = Agent-ID ";" Sequence-Num ";" Signature

4.2. Session Resumption (The Fast-Path)

To minimize cryptographic latency on high-frequency transactions, Agents SHOULD utilize the AIGA Session Resumption mechanism.

1. *Initial Handshake:* The Agent sends a signed AIGA-Hello to the Governance Node (GN) containing its Identity Certificate.

2. *Token Issuance:* The GN verifies the signature and issues a time-bound, symmetric Session-Token (valid for less than 3600 seconds).

3. *Fast-Path Request:* The Agent includes this token in the AIGA-Session header for subsequent peer-to-peer requests.

Peers verify the Session-Token via a low-latency HMAC check rather than a full Public Key verification.

5. Control Plane and Enforcement

The Agent MUST maintain an internal State Machine. Governance Nodes may modify this state via signed Control Messages.

5.1. Hardware-Enforced Termination (The Silicon Kill Switch)

For Risk Class 4 and 5 Agents (Autonomous Code Generation / Kinetic), software-level termination is considered insufficient. These Agents MUST execute within a Trusted Execution Environment (TEE) capable of Remote Attestation.

Attestation Heartbeat: The Agent MUST transmit a Hardware Attestation Quote signed by the Platform Endorsement Key (PEK) every 60 seconds. The Quote MUST certify that the hash of the running binary matches the immutable KernelHash.

Termination Enforcement: If the Governance Node issues a 0xFF (TERMINATE) opcode:

1. The instruction is routed to the TEE Secure Processor.

2. The TEE firmware invalidates the memory pages of the Agent.

3. The cryptographic keys held in the TEE are zeroized immediately.

6. IANA Considerations

This document requests the registration of the following HTTP Header:

* Name: AIGA-State

* Reference: This document

This document requests the registration of the OID:

* Value: 1.3.6.1.4.1.99999.1

* Description: AIGA Agent Attributes

7. Security Considerations

Key Compromise: If an Agent’s private key is compromised, a revocation request MUST be published to the Distributed Ledger.

Side-Channel Attacks: Implementers of the Hardware Enclave MUST ensure that speculative execution attacks (e.g., Spectre-class) cannot leak the Agent’s private keys.

8. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <https://www.rfc-editor.org/info/rfc5280>.

[RFC9113] Thomson, M., Ed. and C. Benfield, Ed., "HTTP/2", RFC 9113, DOI 10.17487/RFC9113, June 2022, <https://www.rfc-editor.org/info/rfc9113>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <https://www.rfc-editor.org/info/rfc9000>.

Author’s Address

Edward Richard Aylward Jr.
North Las Vegas
Las Vegas, NV
United States of America
Email: edward.aylward@example.com