

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 1 April 2026

A. Sogomonian
Artificial Intelligence Internet Foundation (AIIF)
28 September 2025

Architecture for the AI Internet Protocol (AIIP)
draft-sogomonian-aiip-architecture-00

Abstract

This document describes a non-normative architecture for the AI Internet Protocol (AIIP), centered on the "ai" URI scheme. It focuses on a direct connection model for robots, devices, and software agents to communicate natively over AIIP.

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 1 April 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	2
2. Terminology	2
3. AI Systems and Robots Direct Connection	3
4. Example Connection Flow	3
5. Security Considerations	4
6. IANA Considerations	4
7. Normative References	4
Author's Address	4

1. Introduction

This document introduces the architecture for the AI Internet Protocol (AIIP), which defines mechanisms and naming structures for AI-centric communication using the "ai" URI scheme. It outlines the motivation, scope, and structure of the protocol and provides a foundation for discussion and refinement within the IETF. The AIIP aims to enable standardized AI-native addressing and interaction models on top of existing Internet infrastructure.

AI-enabled systems such as robots, devices, and intelligent software agents connect to and interact with Internet resources by initiating interactions using "ai://" identifiers. Resources and capabilities are identified and resolved using the AIIP model, allowing a consistent, verifiable invocation flow across heterogeneous networks and deployments.

This work aligns with the URI framework defined in [RFC3986] and follows the URI scheme registration procedures in [RFC7595].

The key words **MUST**, **MUST NOT**, **REQUIRED**, **SHALL**, **SHALL NOT**, **SHOULD**, **SHOULD 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. Terminology

AIIP: The AI Internet Protocol described in this document.

AI System: A robot, device, or software agent that produces or consumes AI capabilities.

Manifest: Signed metadata describing identity, capabilities, and endpoints for invocation.

3. AI Systems and Robots Direct Connection

AI systems (robots, devices, agents) **SHOULD** initiate interactions using "ai://" identifiers. Implementations treat the identifier as the primary handle for discovery and invocation and avoid embedding deployment-specific hostnames in application logic. This improves portability and enables consistent policy enforcement and auditing.

AIIP operates at the application or protocol layer and is transport-agnostic. Implementations **MAY** use any suitable underlying network transport, including wired broadband, 4G or 5G cellular, satellite constellations, or private networks, provided the chosen transport satisfies deployment requirements for latency, throughput, and security.

Devices **SHOULD** embed a resolver client, maintain trust anchors for verifying manifest signatures, and implement local policy for consent when capabilities imply risk (for example, payments, physical actuation, or signing). Implementations **SHOULD** define canonicalization and percent-encoding rules to avoid confusion during signature verification and policy checks.

4. Example Connection Flow

The following illustrative flow shows direct connection and resolution:

AI System (robot123)	Resolution Service	Service Endpoint
-----	-----	-----
1) ai://robot123.factory/global-task		
2) Resolve -> Signed Manifest		
3) Invoke per manifest		

Figure 1

Example identifier (non-normative):

ai://robot123.factory/global-task

Figure 2

Resolution yields a signed manifest describing publisher identity, verification keys, capabilities (for example, "actuate", "inspect"), and one or more endpoints (for example, HTTPS URL or MQTT topic). The client verifies the manifest, selects an endpoint, and proceeds with invocation according to local policy.

5. Security Considerations

Manifests **SHOULD** be integrity-protected and attributable (for example, JOSE or COSE signatures). Clients maintain trust anchors, perform expiry and revocation checks, and log resolutions and invocations for audit where appropriate. User interfaces for human-in-the-loop scenarios **SHOULD** present verified publisher identity and capability prompts when risk is material.

Implementers **SHOULD** avoid embedding personal data in clear-text identifier components and prefer passing sensitive attributes within protected session channels. Canonicalization rules **SHOULD** be defined to mitigate spoofing and confusion attacks.

6. IANA Considerations

This document makes no requests of IANA.

7. Normative References

- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", RFC 3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC7595] Thaler, D., Hansen, T., and T. Hardie, "Guidelines and Registration Procedures for URI Schemes", RFC 7595, June 2015, <<https://www.rfc-editor.org/info/rfc7595>>.
- [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>>.

Author's Address

Aram Sogomonian
Artificial Intelligence Internet Foundation (AIIF)
United States of America
Email: waterbottling@icloud.com