Network Working Group
Internet-Draft
Intended status: Informational
Expires: 18 April 2026

J. Mozley
N. Williams
Infoblox, Inc.
B. Sarikaya
Unaffiliated
R. Schott
Deutsche Telekom
15 October 2025

AI Agent Discovery (AID) Problem Statement
draft-mozley-aidiscovery-00

Abstract

With the deployment of AI agents comes a need for mechanisms to support agent-to-agent discovery. This document presents requirements and considerations for agent-to-agent discovery.

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 18 April 2026.

Table of Contents

1. Introduction

AI Agents play a crucial role in modern telecommunications by enabling intelligent automation, decision-making, and adaptive network management. These agents are software-driven entities that leverage artificial intelligence, including machine learning and natural language processing, to interact with users, applications, and network components.

Organisations require robust automatable mechanisms for the secure advertisement and discovery of AI agents on public and private networks.

2. Conventions and Definitions

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.

3. Scope of Discovery Processes

The scope of the agent-to-agent discovery could include:

* Dynamic registration, deregistration, and advertisement of agent capability

* Context-aware discovery by other agents (domain, geography, operational constraints like batch versus real-time)

* Standardisation of the discovery process with privacy preservation

  - Having a known entry point such as an index of an organisation’s agents

* Standard schemas to describe agent capabilities

  - Versioning and lifecycle management

* Resiliency in advertising agents

* Ensuring trust and verification mechanisms built into the discovery process

* Using the same mechanisms on both public and private networks

* Interoperability and integration with existing technologies and datasets (backward compatability)

Considerations and constraints on these processes may include:

* Organisations needing control over the process of advertising their agents

* Scalability owing to the large number of agents

* Energy and resource efficiency

* Adversarial use of agents (i.e. using inputs from an index to modify agent reasoning)

* Agent decision variance based on temperature, context, or model state

* Requirements for organisations and agents to operate independently

* Legislation and regulatory frameworks that could include constraints such as sovereignty of data and operational considerations (human in the loop governance)

* Legal issues with ownership of names and brands, or having mechanisms to deal with disputes

* Ethical concerns with respect to agent capabilities, access, and/or advertised metadata

4. Autonomy and Governance

Due to the wide range of services agents could be used for, organisations will need autonomy when advertising agents and may need to work within regulatory frameworks that stipulate such things as operational requirements and sovereignty of data, or such regulations may not be acceptable to some organisations. Discovery mechanisms will need to allow each organisation to publish its own agents' capabilities independently of others. A centralised directory of all possible public agents does not facilitate this. An adversarial centralized directory is also able to stifle competitor advertisement capabilities. The needed autonomy ensures that discovery remains resilient to governance disputes, competitive interference, and jurisdictional constraints.

Given that organisations will need to be autonomous but still facilitate agents discovering other agents and being able to communicate securely with them, a well known entry point is needed. This further allows organizations to delegate authority to others for specific jobs or capabilities.

5. Agent Registration Advertisement and Discovery

Organisations will need to advertise agents and their capabilities to other entities via a predictable entry point. This will facilitate any entity's agents discovering the organisation's agent capabilities, and subsequently performing any verification, tests prior to initiating service provided by the agent. To achieve this, the entry point MUST be based on a mechanism that is both ubiquitous and interoperable across public and private networks.

It is likely this will mean a directory service at or referenced by the entry point that an organisation uses to describe agent capabilities based on a standard schema.

Once an agent discovers another agent e.g. via a directory service, it will communicate with it using protocols such as (A2A, MCP, ACP, etc.) via properties it has discovered from the index service.

All communication will need to be secured through public key infrastructure mechanisms. While in-process agent communication avoids network concerns (co-located agents), any external, cross-host agent interaction still routes through existing network infrastructure requiring the entry point.

6. Regulation and Compliance Features

Services provided by agents may be subject to legislation and regulation because they are time critical, relate to health care, financial in nature, etc. The agent discovery process, protocols and communication will need to take this into account as otherwise organizations will not be able to deploy them.

Compliance considerations include, but are not limited to, availability, auditability, and security. Discovery protocols MUST provide mechanisms to ensure high availability, as service outages in regulated environments can have severe legal and operational consequences. Similarly, auditability MUST be supported through verifiable logging of discovery and registration events, enabling organizations to demonstrate compliance during regulatory reviews or incident investigations. These should be considered during the choices of using existing protocols or developing new ones for agent discovery to avoid barriers to implementation.

7. Scalability

Given the large number of agents, volume of transactions, legal and regulatory considerations, the discovery mechanisms need to be both highly scalable and devolved to the organizations advertising agent capabilities. Any solution that introduces centralized bottlenecks or single points of failure is inherently unsuitable for this purpose, which should also be weighed for how a solution is developed.

8. Schema Evolution

Agents will advertise machine-readable capabilities that evolve over time. The discovery system MUST support explicit versioning, content negotiation, and backward compatible evolution of capability schemas. It MUST place constraints on payload size and complexity to ensure efficient transport and caching while allowing extensibility for domain-specific metadata. Mechanisms for deprecation, sunset notices, and compatibility matrices are needed to prevent negotiation failures at runtime.

9. Abuse

Open discovery surfaces are targets for abuse including amplification attacks, scraping, and denial of service. The system MUST support rate limiting, request authentication when appropriate, abuse detection, and response shaping to minimize amplification potential. Economic considerations, such as preventing cost shifting to responders and preserving fairness across tenants, SHOULD be

incorporated into protocol and operational guidance.

10. Connectivity

Discovery alone is insufficient without practical reachability. The problem statement SHOULD acknowledge challenges introduced by NATs, carrier-grade NAT, firewalls, split-horizon networks, and dual-stack IPv4/IPv6 environments. Agents may exist behind egress-only boundaries or at the edge with intermittent connectivity. The system MUST not assume stable, symmetric, or publicly routable paths and SHOULD define expectations for relay, proxy, or rendezvous patterns where direct connectivity is infeasible.

11. Attestation and Provenance

Agents may be required to prove properties about their software and runtime environment. The discovery layer SHOULD carry or reference attestations such as software bill of materials, model/version identifiers, supply-chain provenance, or trusted-execution environment evidence. Safety classifications, permissible-use statements, and risk tiers can be part of discovery metadata to enable policy-aware client behavior.

12. Challenges with Existing Proposed Mechanisms

Several alternative approaches to agent discovery have been proposed, including centralized registries, blockchain-based naming systems, and proprietary service directories. While these models may offer novel features, they introduce significant architectural, operational, and governance challenges. Centralized registries concentrate control and introduce single points of failure, limiting resilience and scalability.

Blockchain-based systems often lack integration with existing internet infrastructure, suffer from latency and cost constraints, and present jurisdictional ambiguity due to decentralized governance. Proprietary directories fragment the discovery ecosystem and inhibit interoperability across agent frameworks and administrative domains. Furthermore, organizations securely delegating authority to AI solution providers requires interfacing between this central authority as opposed to business-to-business which can be used to suppress innovation.

The introduction of a novel discovery protocol would require the establishment of new governance structures, security models, and operational tooling. Early implementations would likely be centralized among a limited set of stakeholders, introducing bottlenecks in protocol evolution and risk mitigation, and could result in centralization.

13. Security Considerations

All agent discovery mechanisms need to be secured through public key infrastructure.

All agent discovery mechanisms need to have thread detection.

All agent discovery mechanisms and security mechanisms need to work in real-time.

All agent discovery mechanisms need to use standardized interfaces or APIs in case external security instances are used.

Protection needed against * DoS attacks * Unauthorized access and control * Secure communication framework e.g., TLS 1.3 etc. * Continuous monitoring of the agent discovery mechanism

14. IANA Considerations

This document has no IANA actions.

15. 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/rfc/rfc2119>.

[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/rfc/rfc8174>.

Acknowledgments

TODO acknowledge.

Authors' Addresses

Jim Mozley
Infoblox, Inc.
Email: jmozley@infoblox.com

Nic Williams
Infoblox, Inc.
Email: nic@infoblox.com

Behcet Sarikaya
Unaffiliated
Email: sarikaya@ieee.org

Roland Schott
Deutsche Telekom
Email: roland.schott@telekom.de