Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 21 November 2026

H. Yu
China Internet Network Information Center

20 May 2026

IPv6 Networking Considerations for AI Agent Communication
draft-yu-ai-agent-ipv6-networking-considerations-00

Abstract

AI agents are increasingly expected to communicate across platforms, organizations, clouds, edge environments, and administrative domains. Ongoing work on agent protocols has started to address agent description, discovery, authentication, authorization, invocation, delegation, and collaboration. These functions are essential, but they generally treat the IP network as a transparent connectivity substrate.

This document discusses IPv6 networking considerations for AI agent communication. It focuses on networking support after an agent has been discovered by an application-layer or control-plane mechanism. It discusses how stable Agent identifiers can be associated with IPv6 locators or Agent Gateways, how agent communication requirements can be mapped to IPv6 or SRv6 forwarding policies, how limited agent-related network context can be carried within controlled deployment environments, and how network-layer telemetry can complement application-layer audit records.

This document does not define a new agent discovery protocol, a new application-layer agent collaboration protocol, or a new authentication or authorization mechanism. It is intended to be complementary to existing and emerging agent discovery, identity, authorization, auditing, and collaboration mechanisms.

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/.

Copyright Notice

Table of Contents

1. Conventions Used in This Document

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. Introduction

AI agents are moving from isolated applications toward large-scale, cross-domain collaboration. An agent may invoke tools, communicate with other agents, migrate across cloud or edge nodes, act on behalf of a user or organization, and interact through multiple application-layer protocols. In such environments, agent communication is not only an application-layer issue. It also raises questions of network reachability, locator binding, path selection, traffic isolation, observability, and operational governance.

Several ongoing discussions and individual Internet-Drafts explore agent registration, discovery, identity, capability advertisement, authorization, delegation, auditing, and invocation. Such work is important for enabling agents to describe themselves, discover each other, and select suitable application-layer interaction mechanisms. However, after an agent has been discovered and an application-layer protocol has been selected, the resulting communication still needs to be carried by the underlying network.

In large-scale deployments, the following questions remain relevant at the networking layer:

* How is a stable Agent identifier associated with a changing IPv6 locator or Agent Gateway?

* How can a discovered agent endpoint be connected to IPv6 reachability and network policy enforcement?

* How can different agent communication flows be steered through different network paths or service chains?

* How can agent communication be isolated, observed, and audited at the network layer?

* How can application-level intent be translated into structured and enforceable IPv6 or SRv6 policies without requiring ordinary routers to understand agent semantics?

This document discusses IPv6 networking considerations for AI agent communication. It uses IPv6 as the basic addressing and forwarding substrate. In deployments where Segment Routing over IPv6 is available, SRv6 may be used to support policy-based path steering, service function chaining, and domain-local network programming. The IPv6 Segment Routing Header (SRH) is defined in [RFC8754], and SRv6 Network Programming is defined in [RFC8986].

This document does not assume that ordinary Internet routers understand agent identities, task semantics, natural language intent, or application payloads. Any processing of agent-related network context is limited to explicitly participating endpoints, Agent Gateways, policy enforcement points, or SRv6-capable nodes within controlled domains.

3. Scope and Non-Goals

This document focuses on:

* the separation of stable Agent identifiers from IPv6 locators;

* the binding of Agent identifiers to IPv6 locators or Agent Gateways;

* IPv6 reachability for discovered agents;

* mapping agent communication requirements to IPv6 or SRv6 policies;

* limited carriage of agent-related network context in controlled deployments;

* network-layer telemetry that can complement application-layer audit records;

* operational considerations for deployment in endpoints, Agent Gateways, enterprise networks, data centers, operator networks, and SRv6 domains.

This document does not define:

* a new agent discovery protocol;

* a new application-layer agent collaboration protocol;

* a new authentication, authorization, or delegation protocol;

* a new replacement for A2A, MCP, ARDP, Agent Directory, OAuth, OIDC, DID, SD-JWT, WIMSE, RATS, SCITT, or OpenTelemetry mechanisms;

* a general-purpose semantic routing protocol;

* a requirement that public Internet routers process agent-specific metadata;

* a requirement that IPv6 extension headers be usable across the public Internet.

4. Terminology

Agent:
: A software entity that can perceive context, reason, plan, invoke tools, communicate with other agents, and perform tasks on behalf of a user, organization, or system.

Agent Identifier:
: A stable identifier used to identify an agent independently from its current network location. This document uses the term Agent-ID when referring to such an identifier in examples.

IPv6 Locator:
: An IPv6 address or prefix used to reach an agent instance, an Agent Gateway, or another network element representing the agent.

Agent Gateway:
: A network or application gateway that represents one or more agents and provides controlled access, address mapping, policy enforcement, protocol adaptation, or telemetry collection.

Agent Context:
: A limited set of structured metadata related to agent communication, such as an Agent-ID reference, Session-ID, Trace-ID, Policy-ID, Trust-Level, or Audit-Flag. Agent Context in this document does not include natural language prompts, user private data, or application payloads.

Agent Domain:
An administrative or operational domain that explicitly supports agent-related networking functions, such as Agent-ID to locator binding, Agent Gateway enforcement, SRv6 policy selection, or network-layer telemetry collection.

Agent-ID Registration Data Service:
A trusted query mechanism that can provide registration data for an Agent-ID, such as the responsible entity, lifecycle state, associated domains, IPv6 locators, Agent Gateways, public key references, and revocation status. An RDAP-based profile is one possible realization, but is not specified by this document.

5. Problem Statement

Agent discovery and collaboration mechanisms can help an application identify an agent, understand its capabilities, and select an application-layer protocol for interaction. However, those mechanisms do not by themselves define how the resulting traffic should be handled by the IPv6 network. This document identifies four networking problems that become important in large-scale and cross-domain deployments.

5.1. Separation of Agent Identity and Network Location

An agent may be replicated, migrated, hidden behind a gateway, or deployed across multiple network domains. Its identity should remain stable, while the IPv6 locator used to reach it may change. A method is needed to associate Agent identifiers with IPv6 locators or Agent Gateways in a verifiable and operationally manageable way.

5.2. Network Reachability After Discovery

An agent discovery mechanism may return an endpoint, but that result does not necessarily prove that the endpoint is reachable, authorized to represent the Agent-ID, or suitable for a specific network policy. Agent communication therefore needs a way to connect discovery results with IPv6 reachability, locator selection, and policy enforcement.

5.3. Path Control and Service Chaining

Different agent communication flows may have different network requirements. Some flows may require low latency, some may require security inspection, some may require a compliance-aware path, and some may require audit collection. A method is needed to map structured communication requirements into IPv6 or SRv6 forwarding behavior without requiring the data plane to interpret application

semantics.

5.4. Network-Layer Evidence for Audit

Application-layer logs may record task execution, authorization decisions, or API invocations. Network operators may also need independently collected flow-level evidence, such as source and destination locators, Agent Gateway traversal, SRv6 policy identifiers, timestamps, and telemetry records. Such evidence can complement application-layer audit records, but does not replace them.

6. Design Requirements

The following requirements guide the considerations discussed in this document. They are intended to constrain Agent-related use of IPv6 and SRv6 mechanisms. They do not define a new application-layer agent protocol.

REQ-1:
Agent identity and network location MUST be decoupled. A stable Agent identifier MUST NOT depend on a single IPv6 locator.

REQ-2:
Agent-ID to IPv6 locator or Agent Gateway bindings MUST be verifiable by an authorized mechanism.

REQ-3:
This document MUST NOT require ordinary Internet routers to understand agent identities, task semantics, natural language intent, or application payloads.

REQ-4:
Metadata exposed at the network layer MUST be minimized and MUST NOT contain sensitive user data, natural language prompts, private application payloads, or long semantic descriptions.

REQ-5:
Deployments MUST be possible without relying on IPv6 extension headers being forwarded across the public Internet.

REQ-6:
Use of SRv6-specific mechanisms MUST be limited to SRv6-capable domains and MUST follow the operational and security constraints of SRv6.

REQ-7:
The considerations in this document MUST remain complementary to application-layer discovery, identity, authorization, auditing, and collaboration mechanisms.

REQ-8:
Network-layer telemetry SHOULD be able to correlate with application-layer audit records through structured identifiers such as Trace-ID, Policy-ID, or Gateway-ID.

REQ-9:
Deployments MUST define where agent-related network context is processed, who is authorized to process it, and how such processing is protected and audited.

7. Networking Considerations Overview

The considerations discussed in this document can be organized into five logical functions. These functions are logical and may be implemented by different systems in different deployments.

7.1. Agent Discovery Function

This function discovers agent capabilities and endpoints through existing or emerging mechanisms, such as an Agent Directory, ARDP, A2A AgentCard, MCP-related directories, DNS-based indirection, or enterprise service discovery systems. This document does not replace this function.

7.2. Agent Identity and Binding Function

This function maintains or queries the binding between an Agent-ID and networking information. A binding record may include:

* Agent-ID;
* responsible entity;
* IPv6 locator;
* IPv6 prefix;
* Agent Gateway;
* certificate or public key reference;
* binding status;

* validity period;

* revocation status.

The binding function may use an Agent-ID Registration Data Service, an Agent Registry, DNS-based indirection, an enterprise identity system, an operator-managed database, or a controller. This document does not require a single global registry.

7.3. Agent Locator Resolution Function

This function selects the IPv6 locator or Agent Gateway that should be used for a specific communication. Selection may depend on network reachability, policy, trust level, service class, administrative domain, or deployment location.

7.4. Network Policy Function

This function maps structured communication requirements to network policies. Such policies may include default IPv6 forwarding, SRv6 path steering, service-chain insertion, security gateway traversal, compliance-aware routing, audit collection, traffic isolation, or low-latency forwarding.

7.5. Forwarding and Telemetry Function

This function carries agent communication using IPv6 and, where applicable, SRv6. It may also collect telemetry information that can be correlated with application-layer audit records.

8. Agent-ID to IPv6 Locator Binding

This document separates Agent-ID from IPv6 locator. The Agent-ID identifies the agent as a stable entity. The IPv6 locator identifies the current network location of an agent instance, an Agent Gateway, or another network element that represents the agent.

An Agent-ID may be associated with one or more IPv6 locators. The relationship may be one-to-one, one-to-many, many-to-one, or dynamic:

* A fixed agent instance may map to one IPv6 locator.

* A replicated agent service may map to multiple IPv6 locators.

* Many agents may be represented by a single Agent Gateway IPv6 locator.

* A mobile or migrated agent may change its IPv6 locator over time.

Binding information SHOULD be authenticated and SHOULD include a validity period and revocation status. If an Agent Gateway represents an Agent-ID, the gateway authorization SHOULD also be verifiable.

9. Use of IPv6 and SRv6

IPv6 is used as the basic network-layer substrate. In simple cases, agent communication can use ordinary IPv6 forwarding. In more advanced cases, SRv6 may be used within an SRv6-capable domain to support path steering, service chaining, security inspection, or audit path selection.

SRv6 may be useful when agent communication needs to pass through specific network functions, such as:

* Agent Gateway;
* authentication gateway;
* policy enforcement point;
* security inspection node;
* audit collection node;
* data compliance gateway;
* inference gateway;
* target agent service endpoint.

This document avoids defining new IPv6 extension headers. [RFC8200] notes that defining new IPv6 extension headers is not recommended unless existing extension headers cannot be used, and also notes operational concerns around Hop-by-Hop options. Therefore, this document primarily considers existing mechanisms and deployment profiles.

10. Agent Context Carriage

Limited Agent Context may be carried in network-layer or network-adjacent mechanisms. Candidate fields include:

* Agent-ID reference;
* Session-ID;

* Trace-ID;
* Task-ID reference;
* Policy-ID;
* Trust-Level;
* Audit-Flag;
* Gateway-ID;
* Validity-Time;
* Service-Class.

Agent Context MUST be minimized. Large semantic descriptions, natural language prompts, private user data, or sensitive application payloads MUST NOT be carried in IPv6 extension headers or other network-visible metadata fields.

10.1. Destination Options

IPv6 Destination Options may be considered when the context is intended for the final destination or an intermediate destination identified by a Routing Header. [RFC8200] specifies that a Destination Options header can appear before a Routing Header or before the upper-layer header.

Destination Options may be suitable only in deployment environments where participating endpoints or gateways are explicitly configured to process such options. This document does not assume that Destination Options are reliably forwarded or processed across the public Internet.

10.2. SRH TLVs

SRH TLVs may be considered inside an SRv6-capable domain when the information is related to SRv6 policy, service-chain context, audit path selection, or domain-local agent policy enforcement. Such TLVs are not intended to carry user prompts, full Agent identities, or sensitive payload data.

10.3. Fallback Mechanisms

Because IPv6 extension header support may be inconsistent across public networks, deployments SHOULD provide fallback mechanisms. Such mechanisms may include application-layer headers, TLS-protected metadata, Agent Gateway encapsulation, controller-side policy mapping, or out-of-band policy association.

11. Semantic Routing Considerations

This document does not require routers to understand natural language, prompts, complex semantic descriptions, or agent reasoning processes. Semantic routing, if used, is understood as an indirect process:

1. An upper-layer agent system, registry, controller, or gateway interprets task intent and converts it into a structured network policy identifier.

2. The IPv6 or SRv6 network uses the structured policy identifier to select a path, service chain, security domain, telemetry behavior, or audit policy.

```
User or task intent
   -> Agent application / Registry / Gateway
   -> structured Policy-ID
   -> IPv6 or SRv6 policy
   -> forwarding path or service chain
```

This design keeps semantic interpretation out of the forwarding plane and limits network-layer behavior to structured and verifiable policy identifiers.

12. Network Telemetry and Audit Support

IPv6 and SRv6 deployments can provide network-layer evidence that complements application-layer audit records. Examples include:

* source and destination IPv6 locators;

* Agent Gateway traversal records;

* SRv6 policy identifiers;

* service-chain identifiers;

* Trace-ID or Policy-ID correlation fields;

* timestamps and flow-level telemetry;

* policy enforcement events;

* path or domain transition records.

This document does not define a complete audit record format, legal compliance framework, or non-repudiation mechanism. Such functions may be provided by application-layer audit architectures, logging systems, transparency services, or other mechanisms. This document focuses on the network-layer evidence that can be correlated with those systems.

13. Deployment Models

The considerations in this document can be applied incrementally. This document identifies three non-exclusive deployment models.

13.1. Endpoint and Gateway Mode

In this model, agent-related network context is processed only by agent endpoints and Agent Gateways. The public Internet or transit network forwards ordinary IPv6 traffic and does not process agent-specific metadata. This model is suitable when extension header processing is not available or when functions are implemented primarily at the edge.

13.2. Controlled-Domain Mode

In this model, agent-related networking functions are deployed within an enterprise, campus, data center, cloud, operator network, or dedicated interconnection environment. Participating nodes may process limited Agent Context according to local policy. This model is suitable for environments where administrative control, security policy, and extension header behavior can be managed.

13.3. SRv6-Domain Mode

In this model, SRv6 policies are used to steer traffic through specific service functions or network domains. Agent-related context may be associated with SRv6 policies, service chains, or domain-local telemetry collection. This model is limited to SRv6-capable domains and must follow SRv6 operational and security practices.

14. Relationship to Existing Work

The considerations in this document are complementary to existing and emerging work on agent protocols. Agent discovery mechanisms can identify agents, advertise capabilities, and resolve endpoints. A2A and MCP focus on application-layer interaction, tool invocation, and collaboration. OAuth, OIDC, DID, SD-JWT, WIMSE, RATS, and related mechanisms may be used for identity, authorization, attestation, selective disclosure, or workload identity. OpenTelemetry and related systems may provide application-layer tracing and observability.

This document focuses on a different layer: IPv6-based networking support for agent communication after discovery. It can provide locator binding, network policy selection, Agent Gateway traversal, SRv6 service chaining, and network-layer telemetry that can be correlated with higher-layer identity, authorization, and audit mechanisms.

A common limitation of many agent discovery and collaboration mechanisms is that they stop at the point where an endpoint, protocol, credential, or capability description has been selected. They generally do not specify how the selected interaction is bound to an IPv6 locator or Agent Gateway, how the traffic is steered through an operator controlled path, or how network-layer evidence is correlated with application-layer audit records. This document addresses that gap.

This document also avoids placing agent semantics in ordinary routers. Any interpretation of user intent, task descriptions, or agent capabilities is expected to occur in an application, registry, controller, or gateway. The network is expected to operate on structured identifiers, locators, and policies.

15. Security Considerations

This document raises several security considerations.

First, Agent-ID to IPv6 locator or Agent Gateway binding must be authenticated. An attacker must not be able to claim an Agent-ID or bind it to an unauthorized locator.

Second, Agent Gateway authorization must be verifiable. A gateway must not represent an agent unless it has been authorized by the responsible entity or administrative domain.

Third, Agent Context carried in network-visible mechanisms must be minimized and integrity-protected where necessary. Sensitive data must not be exposed in cleartext network headers.

Fourth, replay attacks must be considered. Context fields such as Policy-ID, Trace-ID, or Agent-ID references may need validity time, nonce, session binding, or cryptographic protection depending on the deployment.

Fifth, policy identifiers must not be used to bypass access control, security inspection, or authorization checks. Network policy selection must be integrated with existing security policy enforcement.

Sixth, this document does not assume that IPv6 extension headers are available across the public Internet. Deployment profiles must define where such mechanisms are allowed and what fallback mechanisms are used.

Seventh, SRv6 usage must follow SRv6 security considerations. SRv6 policies and SIDs used for agent communication must be protected against unauthorized insertion, modification, or misuse.

16. Privacy Considerations

Agent-related metadata may reveal information about agent identity, task class, policy class, organizational relationships, or network path choices. Deployments must minimize exposed metadata and avoid carrying sensitive user data in network-layer headers or network-visible metadata fields.

Agent-ID references may be pseudonymous, scoped to a domain, or represented by opaque references. Detailed identity information should be obtained through authorized query mechanisms, rather than exposed in every packet.

Network telemetry used for audit support may contain sensitive operational information. Access to telemetry records should be controlled, retained only as long as necessary, and protected against unauthorized disclosure or correlation.

17. IANA Considerations

This initial version does not request any IANA action.

Future versions may request code points for specific Destination Options, SRH TLVs, or other identifiers if concrete encodings are standardized. Such requests are out of scope for this considerations document.

18. 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>.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <https://www.rfc-editor.org/info/rfc8986>.

19. Informative References

[I-D.pioli-agent-discovery] Pioli, L., "Agent Registration and Discovery Protocol", <https://datatracker.ietf.org/doc/draft-pioli-agent-discovery/>.

[Agent2Agent-Archive] IETF, "IETF Agent2Agent Mailing List Archive", <https://mailarchive.ietf.org/arch/browse/agent2agent/>.

Author’s Address

Haisheng Yu
China Internet Network Information Center
Email: yuhaisheng@cnnic.cn