DMSC Working Group
Internet-Draft
Intended status: Standards Track
Expires: 9 August 2026

X. Li
China Telecom

5 February 2026

Multi-agent Collaboration Protocol Suite based on Agent Gateway
draft-li-dmsc-mcps-agw-00

Abstract

This document specifies a Multi-agent Collaboration Protocol Suite based on Agent Gateway, which enables scalable, secure, and semantically driven collaboration among distributed agents across heterogeneous networks. The protocol suite introduces Agent Gateways as control-plane entities responsible for agent registration, authentication, capability management, semantic routing and other functions, while preserving direct peer-to-peer semantic interactions among agents.

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 9 August 2026.

Table of Contents

1. Introduction

As multi-agent systems become increasingly distributed across heterogeneous networks and administrative domains, efficient, secure, and semantically meaningful collaboration among agents becomes a critical challenge. Traditional service-oriented or message-based interaction models are insufficient to capture agent-level capabilities, dynamic task decomposition, and semantic intent-driven communication.

This document specifies a Multi-agent Collaboration Protocol Suite based on Agent Gateway (AGW). The suite defines a set of coordinated protocols that enable agent registration, authentication, capability synchronization, semantic routing, task-based invocation, and peer-to-peer semantic interaction. The architecture leverages Agent Gateways as first-class network entities that mediate control, policy enforcement, and orchestration, while allowing agents to directly exchange semantic information once authorized.

The protocol suite is aligned with the architectural principles of control/forwarding plane separation, least-privilege authorization, and session-scoped semantic communication.

2. Conventions used in this document

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 [RFC2119] .

3. Terminology

The following terms are defined in this draft:

* AGW: Agent Gateway. A network-resident control and forwarding entity responsible for agent registration, local binding, capability management, semantic routing, and policy enforcement.

* Agent: An autonomous software entity capable of perception, planning, decision-making, and execution.

* Semantic Routing: The process of routing an Agent request based on the meaning or intent of the request, rather than solely on a pre-defined address or identifier.

* Central Auth: Central Authentication Service: A logically centralized authority that performs identity verification and authorization decisions for agents and gateways.

4. Multi-Agent Collaboration Protocol Suite Overview

The Multi-Agent Collaboration Protocol Suite based on Agent Gateway defines a set of coordinated protocols as shown in figure 1 that collectively enable secure agent onboarding, ,distributed capability visibility, semantic request resolution, peer-to-peer semantic interaction, and task-oriented multi-agent orchestration. Rather than operating independently, these protocols are designed to be executed in a tightly coupled manner along the agent lifecycle and collaboration workflows. The Agent Gateway (AGW) serves as the anchoring point for control-plane coordination, while semantic interactions are progressively delegated to agents once resolution and authorization are completed.

The protocol suite consists of the following functional components::

* Agent Registration Protocol (ARP) and Agent Authentication and Authorization Protocol (AAAP), which jointly establish agent identity, trust, and operational scope.

* Capability Synchronization Protocol (CSP), which maintains distributed visibility of agent capability digest across gateways.

* Semantic Resolution and Routing Protocol (SRRP), which enables semantic request discovery and routing across gateway domains.

* Task-based Invocation Protocol (TIP), which extends semantic routing to multi-agent task decomposition and orchestration.

Each protocol operates at a specific phase of the collaboration lifecycle and may be invoked independently or in combination with others. The following sections describe how these protocols are integrated into coherent operational flows.

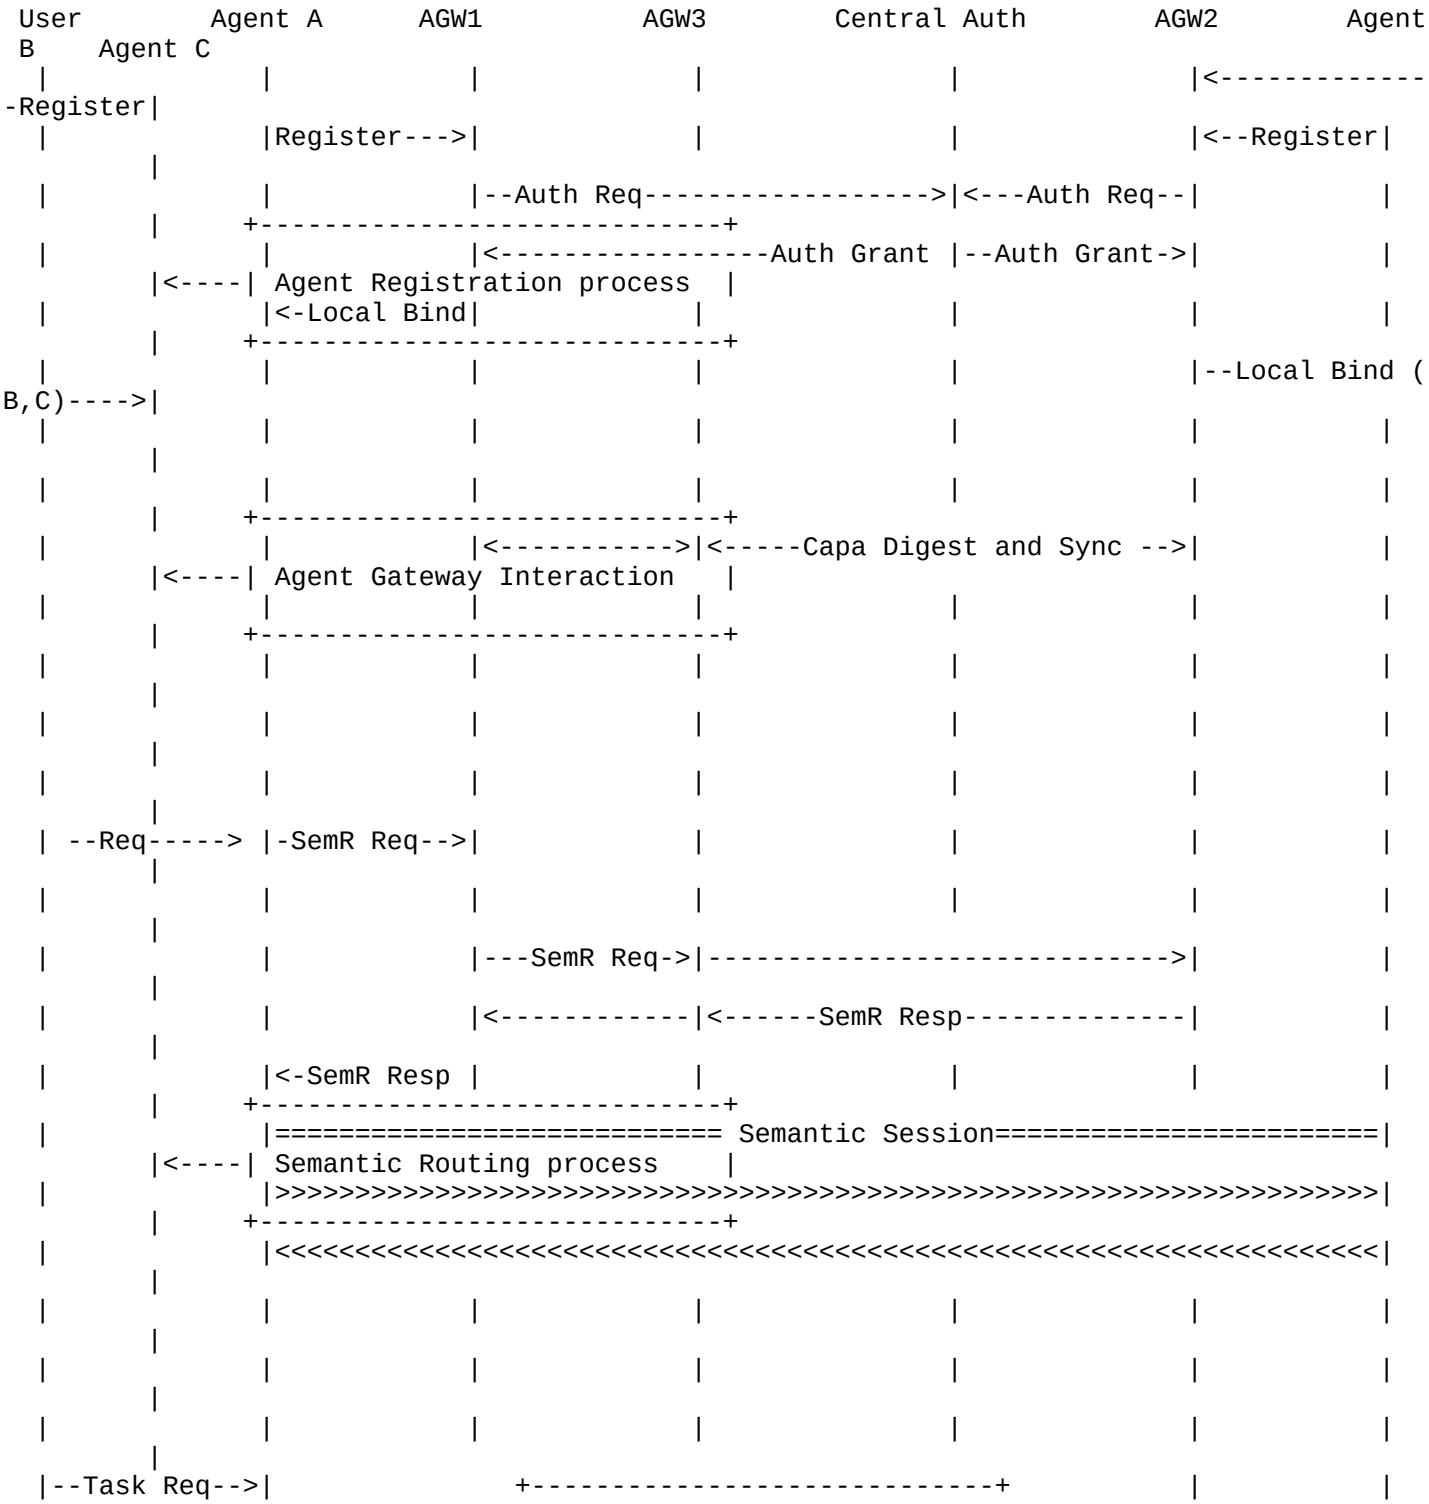

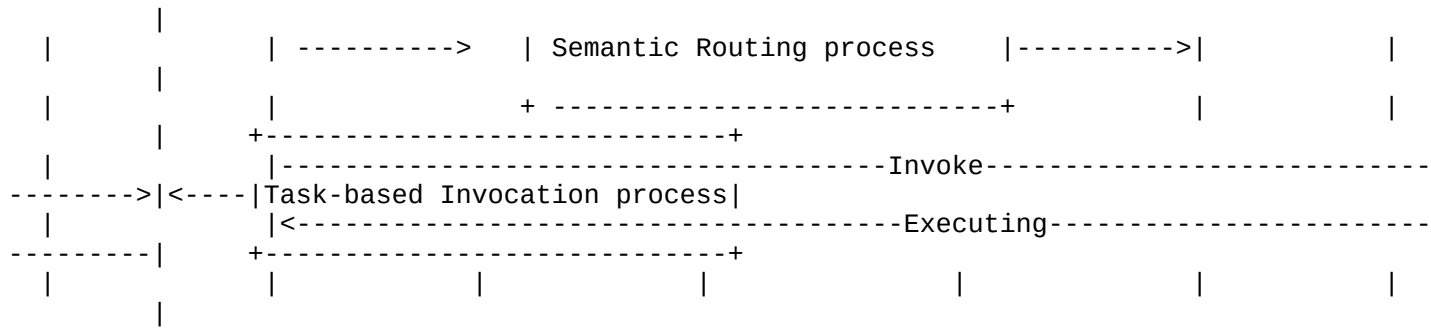

Figure 1 The overall sequence diagram of MCPS-AGW

4.1. Agent Registration and Authorization Process

An agent MUST register with its locally attached Agent Gateway before participating in any collaboration. This process is governed jointly by ARP and AAAP and establishes the agent's identity, trust status, and capability binding. Upon receiving a registration request, the Agent Gateway performs preliminary validation of the agent's identity attributes and initial capability description, and creates a provisional local binding. The gateway then initiates an authentication and authorization request to the Central Authentication Service, conveying the agent identity, gateway identity, and requested operational scope.

The Central Authentication Service evaluates the request and returns an authorization grant or denial. Upon successful authorization, the Agent Gateway finalizes the registration by assigning the agent a globally unique Agent Identifier and Capability Identifier(s). An agent MUST NOT be considered active, discoverable, or invocable until this process completes successfully. This combined registration and authorization procedure ensures that all subsequent semantic routing and task invocation operate on authenticated identities and policy-approved capability scopes.

4.2. Capability Digest and Synchronization Process

Each Agent Gateway maintains detailed capability information only for its locally registered agents. Gateways do not synchronize full agent capability states with each other. Instead, to support inter-gateway semantic resolution, gateways exchange capability digests using the Capability Synchronization Protocol (CSP). A capability digest is a locally generated, abstract summary of available capabilities, designed solely to indicate what kinds of capabilities exist behind a gateway, rather than how those capabilities are internally implemented or executed by agents.The structure and semantics of capability digests are intentionally decoupled from agent-internal capability representations, allowing gateways to evolve local capability models without impacting inter-gateway interoperability.

CSP distributes these capability digests incrementally. An initial exchange establishes basic inter-gateway visibility, while subsequent updates convey only digest changes, such as newly advertised capabilities, capability updates, or withdrawals. Digest updates are versioned and acknowledged to support consistency and conflict resolution. Through this digest-based mechanism, gateways maintain a scalable and privacy-preserving view of distributed agent capabilities without requiring centralized directories or full capability replication.

4.3. Semantic Resolution and Routing Process

When a user issues a request to an agent (e.g., Agent A), the agent abstracts the request into a semantic request and submits it to its locally attached Agent Gateway (AGW1). This interaction is governed by the Semantic Resolution and Routing Protocol (SRRP). Upon receiving the semantic request, AGW1 performs semantic parsing and normalization and consults its local capability directory. If no matching capability identifier is found, AGW1 forwards the semantic request to a peer or upstream gateway (e.g., AGW3), which repeats the same resolution procedure. If the request remains unresolved, it is further forwarded to another gateway (e.g., AGW2).

When a gateway identifies a matching capability in its local directory, it generates a semantic resolution response containing the resolved capability identifier and the corresponding target agent information. This response is propagated hop-by-hop back to the originating gateway and ultimately delivered to Agent A.

Following successful resolution, Agent A and the target agent (e.g., Agent B) directly establish a semantic session. During the lifetime of this session, semantic data is exchanged directly between agents in a peer-to-peer manner, while gateways remain responsible for resolution, authorization scope enforcement, and security policy application during session establishment.

4.4. Task-based Multi-Agent Invocation Process

Task-based collaboration extends semantic resolution to scenarios requiring multiple agents and coordinated execution, as defined by the Task-based Invocation Protocol (TIP). When a user initiates a task request, the request is delivered to Agent A, which performs semantic understanding of the task and decomposes it into one or more sub-tasks along with the required capabilities. If Agent A does not possess task decomposition capabilities, its attached Agent Gateway MAY act as a proxy to analyze and decompose the task on behalf of the agent.

For each sub-task, Agent A submits a semantic request to its local gateway, triggering the same multi-hop semantic resolution process defined by SRRP. Unlike pure point-to-point semantic communication, gateways additionally apply task-level constraints, policy considerations, and capability selection logic to identify suitable target agents.

The resolved results are returned to Agent A, which then directly invokes the selected agents and establishes the necessary semantic sessions for execution. Through this mechanism, multiple agents can

be dynamically selected and coordinated to collaboratively execute complex tasks, while maintaining consistent authorization and security enforcement through gateway-mediated control-plane functions.

5. Conclusion

By explicitly separating control-plane functions from semantic interaction flows, and leveraging gateways as control-plane coordination points, the proposed protocol suite enables scalable and secure multi-agent collaboration without compromising agent autonomy.

6. IANA Considerations

TBD

7. Acknowledgement

TBD

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>.

Author’s Address

Xueting Li
China Telecom
Beiqijia Town, Changping District
Beijing
Beijing, 102209
China
Email: lixt2@foxmail.com