INTERNATIONAL TELECOMMUNICATION UNION
TELECOMMUNICATION STANDARDIZATION SECTOR
STUDY PERIOD 2025-2028

SG11-TD1069/GEN
STUDY GROUP 11
Original: English

Question(s): 1/11 Geneva, 3-11 March 2026

TD

Source: Editors

Title: Output - draft baseline for new work item QSTR.ACC "Signalling and protocol considerations on networking for AI-agent collaboration" (Geneva, 3-11 March 2026)

Contact:	Wei Wang China Telecom	Tel: +86-10-50902397 E-mail: wangw36@chinatelecom.cn
Contact:	Huan Deng China Telecom	Tel: +86-10-50902867 E-mail: denghuan@chinatelecom.com
Contact:	Miao Xue China Unicom	E-mail: xuemiao9@chinaunicom.cn
Contact:	Jinyou Dai CICT	E-mail: diy@fiberhome.com

Abstract: This document is the output of the baseline text of draft new Technical Report QSTR.ACC "Signalling and protocol considerations on networking for AI-agent collaboration", resulting from the discussion of Q1/11 meetings (Geneva, 3-11 March 2026).

C353-R1	China Telecommunications Corporation, China Unicom, China Information Communication Technologies Group	Proposal to start a new Technical Report - QSTR.ACC "Signalling and protocol considerations on networking for AI-agent collaboration"	Q1/11	Accepted with modifications

- Proposal of contribution

  This contribution proposes to initiate a new Technical Report on "Signalling and protocol considerations on networking for AI-agent collaboration".

- Meeting results

  The meeting agreed to accept this contribution with modifications:

  – Remove security-related contents.

  – Remove Section 10 "Considerations on Recommendations".

Table of Contents

Page

Draft new Technical Report ITU-T QSTR.ACC

Signalling and protocol considerations on networking for AI-agent collaboration

1. Scope

This draft Technical Report specifies the signalling and protocol considerations on networking for AI-agent collaboration, including:

- Overview;
- Existing AI-agent Communication Protocols
- Signalling architecture considerations on networking for AI-agent collaboration
- Signalling and protocol considerations on networking for AI-agent collaboration

2. References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published.

The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[ITU-T YSTR.ACN] *AI agent communication network in IMT-2020 networks and beyond*

[ITU-T YSTR.NAC] *Framework of networking for AI agent collaboration in future networks*

3. Definitions

3.1. Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 AI [b- ITU-T L.1022]: Artificial Intelligence.

3.2. Terms defined in this Recommendation

This Recommendation defines the following terms:

TBD

4. Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

5. Conventions

In this Recommendation:

The keywords "is required to" indicate a requirement which must be strictly followed and from which no deviation is permitted, if conformance to this Recommendation is to be claimed.

The keywords "is recommended" indicate a requirement which is recommended but which is not absolutely required. Thus, this requirement need not be present to claim conformance.

The keywords "can optionally" indicate an optional requirement which is permissible, without implying any sense of being recommended. This term is not intended to imply that the vendor's implementation must provide the option, and the feature can be optionally enabled by the network operator/service provider. Rather, it means the vendor may optionally provide the feature and still claim conformance with this Recommendation.

6. Overview

With the rapid advancement of artificial intelligence (AI) technologies, extensive research has been conducted specifically on large language models (LLMs) and multi-agent systems. AI-agents are evolving from "isolated intelligence" toward "collective intelligence". Under this trend, AI-agents are no longer merely standalone program modules; instead, they are becoming autonomous entities that must collaborate with one another, share information, and jointly accomplish complex tasks. Figure 1 shows a schematic diagram of AI-agent communication network.

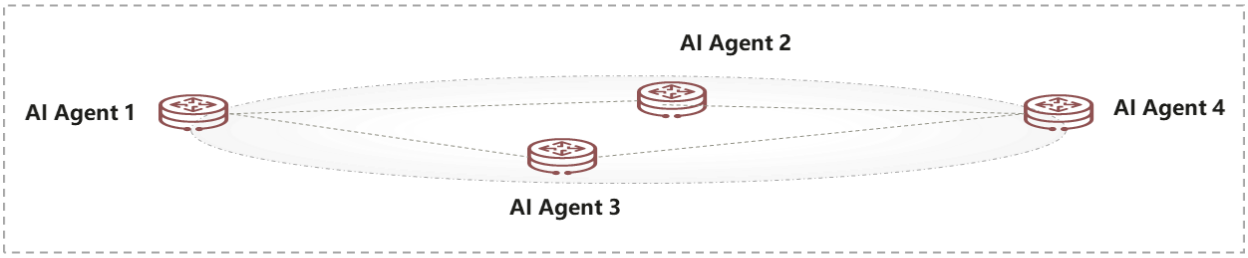

Figure 6-1 A schematic of AI-agent communication network

Within an AI-agent communication network, multiple agents cooperate to execute complex functionalities and business workflows.

Figure 2 shows the process of fault self-healing process based on AI-agents.

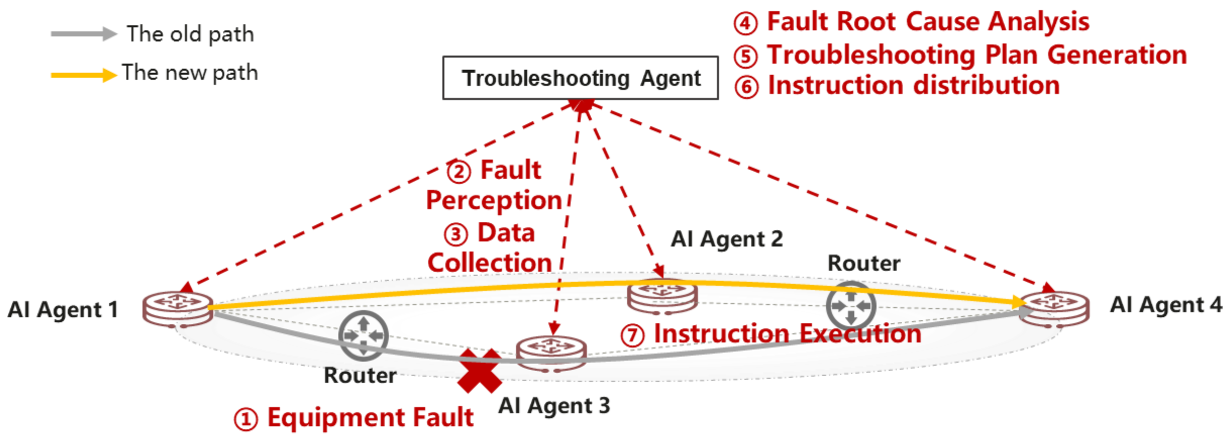

Figure 6-2 Process of fault self-healing process based on AI-agents

Originally, the traffic between AI-agent 1 and AI-agent 4 was transmitted via AI-agent 3 (along the grey path). When a fault occurs in AI-agent 3, the troubleshooting agent perceives the fault and collects relevant fault data. It analyses the root cause of the fault using AI technologies and related algorithms, generates a corresponding troubleshooting plan, and issues instructions to the AI-agents. After the AI-agents execute the received instructions, the traffic between AI-agent 1 and AI-agent 4 will be switched to the backup path (along the yellow path).

Figure 3 shows the process of traffic scheduling and optimization based on AI-agents.

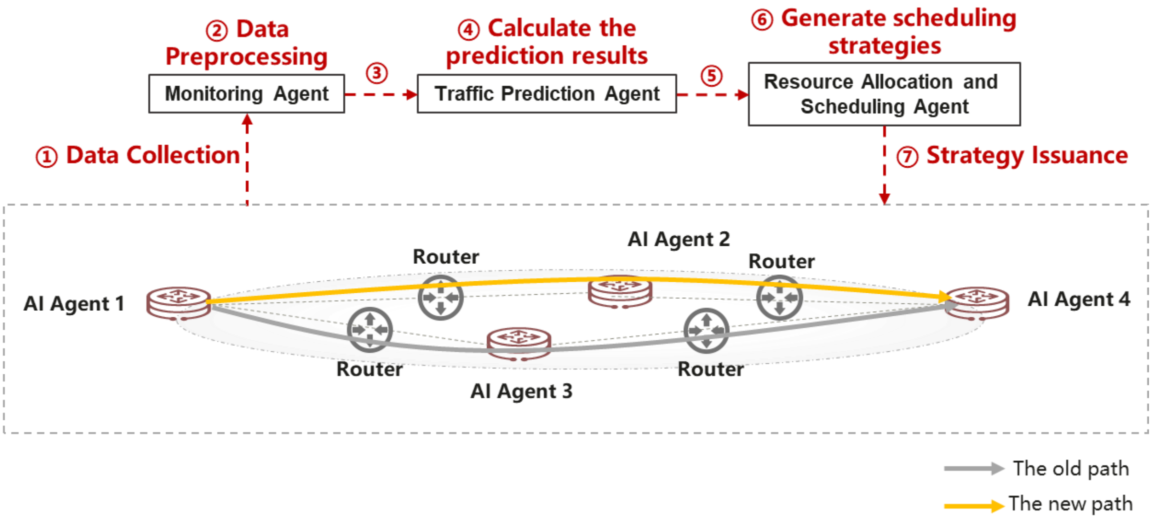

Figure 6-3 Process of traffic scheduling and optimization based on AI-agents

In this process, the monitoring agent continuously collects data in the network, preprocesses the collected data, and then transmits the processed data to the traffic prediction agent. Based on the prediction model, the traffic prediction agent calculates the prediction results and sends the results to the resource allocation and scheduling agent. The resource allocation and scheduling agent generates scheduling strategies according to the received prediction results and issues the strategies to the AI-agents in the network, which make the traffic on the old path (the grey path) switch to the new path (the yellow path).

However, the current AI-agent ecosystem is highly fragmented. AI-agents developed by different vendors, platforms, and frameworks often rely on proprietary communication mechanisms and lack unified standards. This results in poor interoperability, high integration costs, and low collaboration efficiency, which significantly hinder the scalable development and value realization of AI-agent networks.

To better leverage the advantages of AI-agent communication, protocols such as ACP, A2A, and MCP have been proposed. These protocols aim to establish a standardized, efficient, and scalable communication framework for heterogeneous AI-agents, enabling AI-agents to collaborate freely within an AI-agent communication network.

Currently, ITU-T SG13 has initiated several ongoing work items on AI-agents. [ITU-T Y.supp.NAC-roadmap] *"Networking for AI agent collaboration - Standardization status and roadmap"* focuses on describing the standardization status and roadmap of networking for AI-agent collaboration, collecting related Recommendations and Technical reports, and providing a roadmap to the possible standardization direction of networking for AI-agent collaboration. [ITU-T YSTR.NAC] *"Framework of networking for AI agent collaboration in future network"* studies the requirements and framework of networking for AI-agent collaboration in future networks (as shown in Figure 6-4). [ITU-T YSTR.ACN] *"AI agent communication network in IMT-2020 networks and beyond"* studies the AI-agent communication network in the context of IMT-2020 and beyond (as shown in Figure 6-5). ITU-T SG11 has initiated one ongoing Technical Report [ITU-T QSTR.SPACN] "Signalling and protocol for AI-agent communication network in IMT-2020 networks and beyond" studies the signalling and protocol for AI-agent communication network in IMT-2020 networks and beyond.

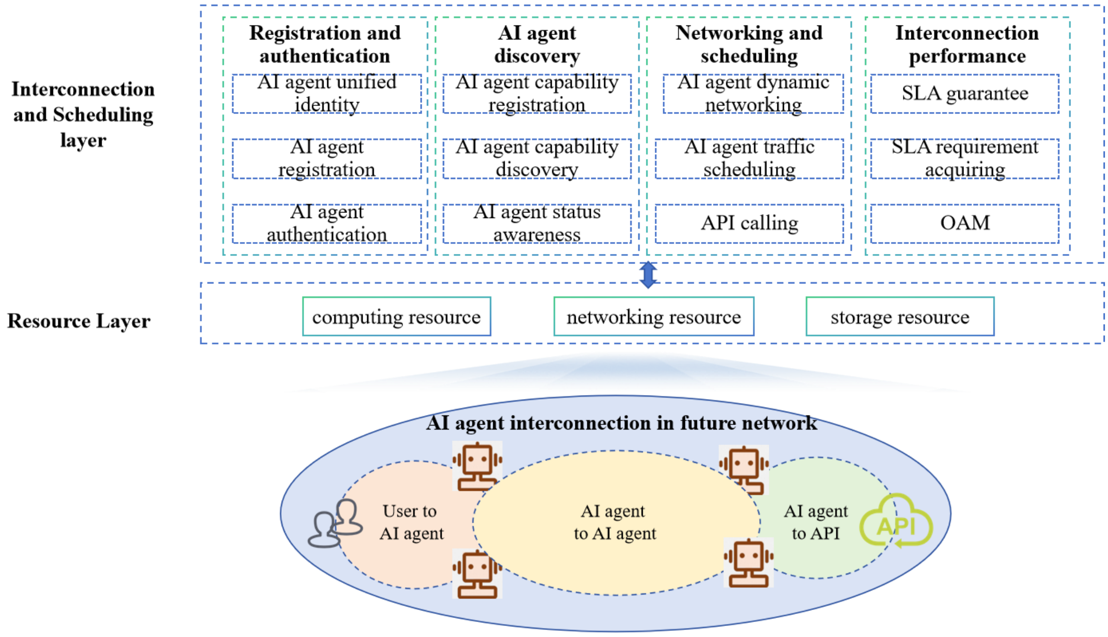

Figure 6-4 Overview of networking for AI-agent collaboration

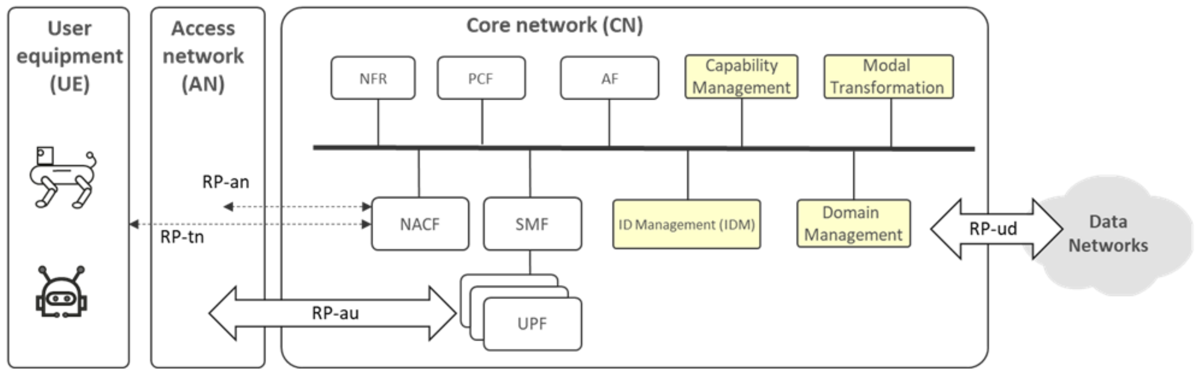

Figure 6-5 Potential framework of AI-agent communication network in IMT-2020 networks and beyond

Based on these ongoing work items, this Technical Report will investigate signalling and protocol considerations on networking for AI-agent collaboration, including a review of current protocols and an analysis of potential requirements for future AI-agent communication protocols.

7. Existing AI-Agent Communication Protocols

[Editor's Note] This clause specifies the existing AI-agent communication protocols.

7.1. Agent Discovery Protocols

Agent discovery protocols aim to enable an agent or control system to dynamically identify, authenticate, and obtain information about other agents' capabilities, interfaces, and access methods, serving as a prerequisite for building scalable multi-agent networks.

Current agent discovery protocols are evolving along three main technical approaches:

1) Static declarative discovery

Static declarative discovery is a lightweight, practical, and easily deployable approach among current agent discovery protocols. Its core idea is to enable service providers to proactively declare, in a standardized and machine-readable format, interface information that is friendly to AI-agents.

A representative protocol of static declarative discovery is agents.json. Websites host an agents.json file in their root directory to declare AI-agent-friendly API endpoints, authentication methods, rate limits, and supported operations.

2) Registry-based centralized discovery

Registry-based centralized discovery is a structured, manageable, and widely adopted pattern in agent discovery protocols, particularly suitable for enterprise environments or closed ecosystems. Its core idea is that all agents actively register their information with a centralized registry upon startup or when their capabilities change. Other agents or orchestration systems then discover and select appropriate collaboration partners by querying this registry.

A representative protocol of registry-based centralized discovery is Agent Card. It allows agents describe their capabilities using a JSON-LD–formatted "digital business card" (Agent Card), which can be published to a curated central registry.

3) Decentralized dynamic autonomous discovery

Decentralized dynamic autonomous discovery is an emerging paradigm in agent discovery protocols, designed for open, heterogeneous, and trustless environments. Its core principle is that each agent possesses a self-sovereign digital identity, enabling it to dynamically broadcast its presence, discover other agents, and negotiate collaboration relationships directly within a peer-to-peer network, without relying on a centralized registry.

A representative protocol is ANP (Agent Network Protocol). Agents broadcast service descriptions via a DHT (Distributed Hash Table), supporting NAT traversal and low-latency networking.

7.2. Agent-to-Agent Communication Protocols

Agent-to-agent communication protocols aim to address how heterogeneous agents in multi-agent systems can efficiently, and semantically consistently exchange information, delegate tasks, and coordinate actions.

1) A2A (Agent-to-Agent Protocol)

A2A is an enterprise-level communication protocol designed for structured collaboration between agents. It adopts a five-layer architecture, with core innovations including Agent Cards (capability declarations), Tasks (standardized task lifecycle), and Artifacts (structured outputs). A2A supports context-aware multi-turn dialogues and task delegation. It emphasizes semantic consistency and interoperability, making it well-suited for closed-loop business workflows involving multiple agents within enterprise or cloud environments.

2) ACP (Agent Communication Protocol)

ACP is a name shared by multiple distinct implementations proposed by different teams. Currently, two mainstream versions exist:

- The AgentUnion version of ACP employs a centralized Access Point (AP) architecture, using AID-based identity and low-latency discovery mechanisms, which is ideal for localized, high-real-time scenarios such as industrial robot swarms.
- The IBM version of ACP extends JSON-RPC, is compatible with MCP, and supports streaming communication and asynchronous callbacks, focusing on standardized integration in enterprise multi-agent systems.

3) MCP (Model Context Protocol)

MCP is a lightweight protocol specifically designed to connect large language models (LLMs) with external tools and data sources. Based on JSON-RPC 2.0 and HTTP/SSE, it enables LLMs to invoke tools, inject contextual information, and retrieve real-time data in a standardized manner. MCP's primary value lies in providing AI applications with a unified "capability interface", significantly simplifying integration between agents and enterprise systems.

7.3. Human-Agent Interaction & Explainability Protocols

As AI-agents increasingly participate in high-risk, high-responsibility scenarios, such as medical diagnosis, financial decision-making, and judicial assistance, human users not only need to use these agents but also must understand their reasoning processes, intervene in their actions, and trace accountability.

1) AG-UI Protocol

AG-UI is an event-driven protocol designed for real-time human-agent collaboration experiences, aiming to standardize communication between front-end interfaces and agent runtimes. Its core value lies in enhancing interaction transparency and user control, making it well-suited for scenarios such as general-purpose AI assistants and AI-powered coding tools.

2) PXP (Predict and eXplain Protocol)

PXP is a structured explainability protocol tailored for high-stakes professional domains. PXP enables bidirectional reasoning validation between human experts and AI-agents. By transforming vague "explanations" into actionable feedback commands, PXP can significantly improving the accuracy and trustworthiness of human-agent collaborative decision-making.

3) LOKA Protocol

LOKA is an ethics- and governance-oriented human-agent interaction protocol. Leveraging W3C Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs), LOKA binds each agent to an auditable identity and embeds ethical alignment checks at critical decision points to ensure behaviors conform to predefined values. Additionally, LOKA empowers users to trace operation logs and initiate accountability procedures, elevating human-agent interaction from mere functionality to compliance, transparency, and social responsibility. It is particularly suitable for highly regulated sectors such as finance, public administration, and education.

8. Signalling architecture considerations on networking for AI-agent collaboration

[Editor's Note] This clause specifies the signalling architecture considerations on networking for AI-agent collaboration.

This clause aims to give a potential signalling architecture of networking for AI-agent collaboration based on [ITU-T YSTR.NAC] and [ITU-T YSTR.ACN]. Once the functional architectures proposed in these two Technical Reports are stabilized, we will derive the corresponding signalling architecture based on the commonalities of these two functional architectures.

9. Signalling and protocol considerations on networking for AI-agent collaboration

[Editor's Note] This clause specifies the signalling considerations on networking for AI-agent collaboration.

The signaling and protocol serves as the common language for AI-agent communication, defining the structure, semantics, and exchange rules of messages.

The signalling and protocol considerations on networking for AI-agent collaboration are as follows:

- **Context Consistency:** All related messages must be bound to the same context identifier (e.g., contextId or sessionId) to ensure state persistence and prevent ambiguity across multi-turn interactions.
- **Observability and Auditability:** Each signaling message should carry traceable metadata (e.g., sender ID, timestamp, trace ID) to facilitate logging, debugging, and accountability attribution.
- **Interoperability:** Use open, standardized data formats and avoid proprietary encodings. Field names and semantics should follow community consensus (e.g., A2A Schema).
- **Transport Flexibility:** The protocols should be decoupled from the underlying transport layer and compatible with multiple channels such as HTTP/SSE, WebSocket, gRPC, and P2P.
- **Error Robustness:** Define standardized error codes and actionable error descriptions, enabling automated retries or graceful degradation.
- **Version Evolvability:** Support protocol version identification and negotiation mechanisms to ensure coexistence of legacy and new agents. New features must be backward-compatible.
- **Extensibility:** Allow custom fields or extension points to accommodate domain-specific needs without breaking the core specification.

Bibliography

[b-ITU-T L.1022] Recommendation ITU-T L.1022 (2019), *Circular economy: Definitions and concepts for material efficiency for information and communication technology.*

___