

Agent Communication Protocols (agentproto) — WG-forming BoF

IETF 126, Vienna · Thursday 23 July 2026 · 09:00–11:00 CEST · Grand Park Hall 3
Chairs: Leslie Daigle · Orie Steele | AD: Charles Eckel (ART)

This session is being recorded

Note Well

By participating in the IETF you agree to follow IETF processes and policies. This Note Well is a reminder of some of those policies. For a linked version of this text, please visit **www.ietf.org/note-well** or use the QR code below

- IETF participants are expected to behave in a professional manner and extend respect and courtesy to their colleagues at all times (see **RFC 7154: IETF Guidelines for Conduct and IETF Anti-Harassment Policy**). If you have any concerns about behavior, please contact the *Ombudsteam* who have a duty of confidentiality and extensive powers to act, as set out in **RFC 7776: IETF Anti-Harassment Procedures**.
- If you are aware that any IETF contribution (as defined in **RFC 5378: Rights Contributors Provide to the IETF Trust**) is covered by patents or patent applications that are owned or controlled by you, your employer or your sponsor, you must disclose that fact, or not participate in the discussion (see **RFC 8179: Intellectual Property Rights in IETF Technology**).
- For detailed process information consult **RFC 2026: Internet Standards Process** and **RFC 2418: IETF Working Group Guidelines and Procedures** and updates to those.
- The IETF routinely makes public written, audio, video, and photographic records of IETF activities, including your personal information as set out in the **IETF Privacy Statement**.



For advice, please talk to Working Group chairs or Area Directors.



This session is being recorded

IETF 126 Meeting Tips

In-person participants


- Make sure to sign into the session via Datatracker or the QR Code in this session.
- Use Meetecho (usually the "Meetecho lite") client to join the queue, show hands
- *Keep audio and video off if not using the onsite version.*

Remote participants

- Make sure your audio and video are off unless you are chairing or presenting during a session.
- Use of a headset is strongly recommended.

All participants

- State your name each time you begin speaking in the queue



Administrivia

- Blue Sheets — auto-generated when you log into Meetecho
- Zulip — <https://zulip.ietf.org/#narrow/stream/126-agentproto>
- Notes — <https://notes.ietf.org/notes-ietf-126-agentproto>
- Agenda & Materials — <https://datatracker.ietf.org/meeting/126/session/agentproto>
- Datatracker — <https://datatracker.ietf.org/group/agentproto/about>
- FAQ –
<https://github.com/sureshkrishnan/agentproto-faq/blob/main/agentproto-faq.md>
- At the mic: state your name + affiliation

Agenda (2 hours)

1. Introduction & administrivia — chairs (10)
2. How this BoF relates to work happening elsewhere — chairs (5)
3. Motivations: use cases → requirements — proponents (12) [clarifying Qs only]
4. Multimodal support & ties to session/transport — proponents (10)
5. Session: managing context across an interaction — proponents (10)
6. Framework: connecting to work in other groups — proponents (10)
7. Clarifying questions — chairs (13)
8. Open-mic discussion (scope & charter) — chairs (25)
9. WG-forming consensus calls — chairs (18)
10. Wrap-up & next steps — chairs (7)

This is a WG-forming BoF

- Today (and following up as appropriate) we decide whether to charter a working group — not to design protocols.
- Charter text is displayed throughout:
 - ***Detailed technical discussion is reserved for the WG, if one forms.***
- By the end we want a sense of the room on:
 - Is the problem clear, well-scoped, and worth solving?
 - Is the need for interoperability clear?
 - Is the IETF the right venue?
 - Is the proposed scope the right cut — one WG, as chartered?
 - Are enough people committed to do the work (review · edit)?

Related venues

- We focus on our own scope — the agent session + framework.
- Related conversations are deferred to their own venues:
 - IETF DAWN (WG forming BoF) — agent / workload / entity discovery
 - IETF DMSC (non-WG forming BoF) — dynamic multi-agent secured collaboration
 - IETF WIMSE (WG) — workload / agent identity
 - IETF OAuth (WG) — authorization
 - Linux Foundation — A2A and MCP
- We can reference and coordinate with this work — we cannot direct it.
- Please take detailed discussion of those topics to those venues.

Thumbnail sketch of where we are going

Quick flash of the charter segments we're going to motivate today –
it's not yet time to charter bash :-)

Agent Proto

"The Agent Communication Protocols (AgentProto) Working Group will work on defining protocol building blocks for enabling interoperability for AI Agent applications across the Internet. An AI agent is an autonomous, adaptive intelligent software system that uses AI models to complete a specific task. AI Agents often interact with users via chat or voice, performing tasks based on the flow of the conversation. To complete tasks on behalf of a human user or another AI agent, they can independently make decisions, execute actions, and interact with other AI agents and tools."

Motivations & Requirements

"Use Cases, Gap Analysis, and Requirements (Informational) – Foundational work will be documented through a set of informational Internet-Drafts covering:

- *Use cases focused on Agent-to-agent and Agent-to-tool communications, used to verify the suitability of existing protocols and the protocols being developed.*
- *Gap analysis and requirements based on examination of existing de facto standard protocols implemented in open-source projects, from which necessary protocol requirements are derived."*

Session

"AI Agent Session Protocol (Standards Track) – A standards-track protocol for creation and maintenance of communication sessions between AI agents, or between AI agents and tools. These sessions allow for the bidirectional exchange of data, including model context, tool call results, and chat messages."

Framework

"AI Agent Protocol Framework (Standards Track) – A standards-track framework document that identifies the key building blocks and defines the protocol suite for interoperable agent-to-agent and agent-to-tool communications. The document provides an architectural overview and highlights areas for subsequent protocol specification work."

Main presentation matter

Motivations & Requirements

"Use Cases, Gap Analysis, and Requirements (Informational) – Foundational work will be documented through a set of informational Internet-Drafts covering:

- *Use cases focused on Agent-to-agent and Agent-to-tool communications, used to verify the suitability of existing protocols and the protocols being developed.*
- *Gap analysis and requirements based on examination of existing de facto standard protocols implemented in open-source projects, from which necessary protocol requirements are derived."*

Motivations: From Use Cases to Requirements

Laurent Ciavaglia

Nokia

Yao Kehan

China Mobile

Lionel Morand

Huawei

Motivation for new work, The Industry Perspective

External ecosystem and industry landscape

- Agentic systems are moving from isolated assistant/application toward **network of agents and services interacting at Internet scale**
- Multiple ecosystem efforts are emerging, creating a **risk of fragmentation and interoperability** without shared protocol work

Motivation for new work, The Industry Perspective


External ecosystem and industry landscape

- Agentic systems are moving from isolated assistant/application toward **network of agents and services interacting at Internet scale**
- Multiple ecosystem efforts are emerging, creating a **risk of fragmentation and interoperability** without shared protocol work


Operator and implementer needs

- Operators need **deployable, interoperable approaches** applicable across diverse use cases, and that can work across administrative and trust boundaries
- Implementers need clearer expectations for functionality such as **discovery, identity, security, data exchange, lifecycle, session management, multi-modality support, observability, etc.**


Motivation for new work, The Standards Perspective (1)



Motivation for new work, The Standards Perspective (2)



Motivation for new work, The Standards Perspective (3)



Describe the key interaction **patterns** of agentic AI communication systems,
Then derive the **protocol requirements** those patterns imply
To guide IETF standardization of an agentic **protocol framework**.

Not scoped to one use case or specific deployment,
Requirements derived from behavior, i.e. what the protocol must do, not who deploys it,
Deployment-neutral → **reusable** and (aim at) applicability across various operators/vendors scenarios.

Motivation for new work, The Standards Perspective (3bis)

Key Interactions Patterns

1

**Simple
single-agent task**

one agent, tools behind it

2

**Orchestrator +
agents**

decompose & delegate subtasks

3

**Long-running +
authz checkpoint**

pause and ask before acting

4

**Peer multi-hop
problem solving**

dynamic delegation graph

5

**Cooperative
reasoning /
consensus**

many agents, one problem

6

**Tool / data / API
mediation**

a gateway in front of resources

Output and transition to next part of the BoF discussion

- Based on these motivations, use cases and gap analysis,
 - Which aspects does AGENTPROTO focus on?
 - What framework and protocols (or extensions, profiles of) should the IETF define?

Multimodal support & ties to session/transport

Suresh Krishnan

Cisco

Dapeng Liu

Alibaba Group

Cullen Jennings

Cisco

Some Real-World Multimodal Agent Interactions

Customer Support Agent

Voice conversation + screen share
+ real-time tool calls

Collaborative Coding

Streaming code + audio discussion
+ live preview rendering

Medical Triage Agent

Video of symptoms + voice Q&A
+ structured form data

Multi-Agent Orchestration

Agent A streams audio to Agent B
while exchanging JSON tasks

Example: Booking a business trip (Agentic version)




← MULTIMODAL – CONCURRENT STREAMS IN PARALLEL →



4 parties · 4 network hops · 3 org boundaries · **Multimodal (voice + video + chat + tools + API)**

What breaks?

“Cancel the hotel, keep the flight.”



How does **CANCEL propagate** across 4 parties and 3 network hops?

What happens to the **in-flight voice stream** between user and agent?

What about the **in-flight tool call** to the hotel API — abort mid-execution?

The hotel API already received **user preferences** — how to revoke shared context?

The airline agent already **charged the credit card** — how to reverse the transaction?

HTTP can't do this — it's stateless, no session lifecycle.


WebSocket can't do this — no multiplexing, no priority, no partial teardown.

MoQ is promising — but it doesn't know what an *agent* is.

Why do we need multimodal support?

- Diverse latency needs within same conversation
 - Realtime media (voice/video/screen share etc.)
 - Semi-realtime interactive text (chat, streaming LLM outputs etc.)
 - Non-realtime bulk data (context transfer/exchange, tool calls, file transfers etc.)
- Need to synchronize across the modalities
 - e.g. A session cancellation on a control channel needs to be able preempt a bulk transfer on a different channel
- Protocol agnostic reuse
 - Provide a substrate that can be reused across multiple application protocols such as MCP, A2A etc. instead of each protocol rolling up its own (non-interoperable) way of doing things.

What do we need to close this gap?



Session

"AI Agent Session Protocol (Standards Track) – A standards-track protocol for creation and maintenance of communication sessions between AI agents, or between AI agents and tools. These sessions allow for the bidirectional exchange of data, including model context, tool call results, and chat messages."

Why Agentic Networks Need a Session Layer to manage Context?

Hesham Moussa

Huawei Technologies Canada

Arashmid Akhavain

Huawei Technologies Canada

❖ **Agentic Networks Involve Multiple Interacting Participants:**

- Human ↔ Agents, Agents ↔ Tools, and Agents ↔ Agents

❖ **Multimodal Communication**

- Text, Audio, Video, Tools, APIs, Sensor Streams

❖ **Nature of Interactions**

- Long-lived and short-lived, Stateful and chained, Parallel, simultaneous, across multiple domains, and fan-out capable

❖ **Conversation and Task Dynamics**

- Can be interrupted, modified, canceled, paused, or resumed
- Evolve over time based on context and environment


❖ **Networking Nature**

- Participants join, leave, act autonomously, and influence the environment
- Agents may fail (so do tools), hallucinate, recover, migrate, be replaced, or evolve
- Multiple frameworks and protocols coexist and adapt dynamically (e.g., A2A, MCP, HTTP)
- Run on top of potentially multiple transports (e.g., MoQ, WebTransport, gRPC, MQTT)

THE PROBLEM: WHAT IS CONTEXT IN AGENTIC NETWORK?


2 / 5

- ❖ **Everything in agentic networks begins with a task**
- ❖ **Tasks create waves of interactions**
 - These waves intersect, overlap, and involve shared participants.
 - Agents exchange text, audio, video, tool calls, sensor data, and structured API responses.
- ❖ **Context is the glue that keeps these waves coherent...For example:**
 - It tells an agent which artifact belongs to which task or query.
 - It tells a coordinator which sub-task belongs to which parent task.
 - It tells the network which interactions are part of the same workflow.
- ❖ **Without PROPER context management and structure, handling tasks becomes challenging**
- ❖ **This is why a session layer is required. A session layer that imposes structure. It:**
 - Imposes order on otherwise a chaotic, multi-modal, and multi-agent workflows,
 - Anchors context and preserves continuity
 - And prevents cross-contamination between tasks



INTEGRATION: WHERE COULD THIS SESSION LAYER FIT?

3 / 5



❖ But why a standardized Session layer?

- Application frameworks such as MCP, A2A, AGNTCY, etc recognize the problem. However, each currently implements its own context and session handling and management mechanics. But, without a standard session layer, the result is fragmentation, not interop.

REQUIREMENTS: WHAT A SESSION LAYER NEEDS TO DO TO HANDLE CONTEXT?

4 / 5

❖ **Resilient against System-level failures**

- Support the "Hydration" Trick (agent context replenishing after crashing)
- Tool Timeout Isolation and Rollbacks
- Graceful Handshakes During Human/Network Drops (humans rejoining a session after dropping)
- Poison Pill Detection (Avoiding Infinite Loops)

❖ **Supports Multi-participant interactions**

- Hierarchical Session Tree (Scoped Context Passing)
- Preventing Race Conditions and outdated contexts
- Fault Isolation and "Blast Radii"
- Support and maintain context across chained interactions
- Provide means for isolated context and accessibility management (breakout rooms)

• **Handle concurrent multi-modal streams (as described earlier)**

- Spatial and Temporal Alignment (The "Look at This" Problem)
- Stream Interruption and Asynchronous Echoes
- Asymmetric Processing Speeds (Heavy vs. Light Modalities)
- Modality Fan-Out and Synthesis

- ❖ **In summary, IETF has protocols that provide**
 - A form of failure recover → QUIC
 - Stream Multiplexing → HTTP/2.
 - Priority handling → WebTransport
 - Low-latency media streaming and a form of fan-out→ MoQ.
- ❖ **However none natively manage agent context across turns, interrupts, reconnects, or handle coordination and multi-entity interactions — that's the gap.**
- ❖ **A new session layer is needed**
 - One that is agent-, tool-, and human- aware
 - Handles multi-modality and stream coordination
 - Provides means for managing context, recovering from system-level failures, and handling interrupts and barge-ins
 - agnostic to underlying protocols and over the top frameworks for seamless integration across the stack
 - and is operational across diverse domains and multi-hop connections

Framework

"AI Agent Protocol Framework (Standards Track) – A standards-track framework document that identifies the key building blocks and defines the protocol suite for interoperable agent-to-agent and agent-to-tool communications. The document provides an architectural overview and highlights areas for subsequent protocol specification work."

Framework for Agent Communication Protocol

Dapeng Liu


Alibaba Group

Hesham Moussa

Huawei

Agent communication needs coordinated building blocks


Framework defines how the building blocks compose into one unified agent system



FRAMEWORK ARCHITECTURE: What are the boundaries between the framework and other WGs?

2 / 5


The framework identifies building blocks (each owned by a home WG) **and defines the composition patterns**



The framework references building blocks & identify gaps.

Coordination with other WGs

Per charter — if we need changes to other WGs' protocols, we raise the issue; they decide.



COORDINATION WGs

Security

OAuth, webbotauth, WIMSE
identity, authorization

Transport

WebTransport, MoQ, QUIC, TSVWG
data transport and session management

Discovery & OPS

INT area, OPS area
agent discovery and operational considerations

The framework will:

Per charter : a standards-track document identifying building blocks and defining the protocol suite.

01 · AGENT COLLABORATE

Enable AI Agents to select and collaborate with other AI Agents/tools across domains and ecosystems, to execute simple or complex tasks.

02 · MULTI-MODAL SUPPORT

Allow multi-modal collaboration using text, images, video, audio, and structured data with exchange of multi-modal contexts.

03 · FUNCTIONAL BLOCKS

Describe functional blocks, relationships, and mechanisms for structured, semi-structured, and multi-modal information exchange.

04 · IDENTITY & AUTH & EVIDENCE

Enable independent identity, describe agent-specific authentication and authorization — how existing and emerging mechanisms are composed, including **confirmation and evidence** requirements.

05 · FINE-GRAINED & DYNAMIC BEHAVIORAL

Fine-grained behavior-driven scopes bound to specific operations, with **dynamic behavioral boundaries**.

06 · PROTOCOL SUITE

Identify the protocol suite covering session management, transport, security, and identity building blocks.

Evolving Work Item

- Iteratively integrates existing and missing building blocks in successive steps, until all core modules are fully incorporated.
- Proceeds in parallel with specific protocol deliverables.

Out of Scope

Per charter — the following topics are explicitly out of scope for this working group.

- 1 **Discovery of AI agents.**
- 2 **Definition of AI models, agent reasoning algorithms, or tool-specific business logic.**
- 3 **Standardization of agent behavior, decision-making, or planning semantics.**
- 4 **AI agent behavioral security** (e.g., preventing hallucinations, though mitigating impact via protocol-level user confirmation is in scope)
- 5 **Design of human-to-agent user interfaces, client application UX, or rendering of agent outputs on end-user devices.**

Clarifying questions & open-mic discussion

- Clarifying questions only on the presentation block — questions to understand the material.
- Open-mic (scope & charter):
- At the mic: name + affiliation.

WG Formation discussion

Current charter text

Is available here:

<https://github.com/ietf-artarea/charters/blob/main/agentproto/charter.md>

WG-forming consensus questions

1. Do you think that the problem statement is well-understood, solvable, and useful to solve?
2. Is the need for interoperability clear?
3. Do you think the IETF is the right place to do this work?
4. Do you think the initial scope from the charter we just discussed is correct?
5. Do you think the initial deliverables from the charter we just discussed are correct?
6. Do you think a WG should be formed on this topic, with a charter based on the draft charter we just discussed?
7. Are you willing to help write and/or review drafts in this potential WG?

Wrap-up & next steps

- Read back the consensus and the named contention points
- Where discussion continues:
 - agent2agent@ietf.org – unless/until more specific list set up
 - <https://datatracker.ietf.org/group/agentproto/about/>
- Thanks to our note-takers / minute-takers!

BACKUP slides

Agent Communication Protocols (agentproto) Proposed Charter

The Agent Communication Protocols (agentproto) Working Group will work on defining protocol building blocks for enabling interoperability for AI Agent applications across the Internet. An AI agent is an autonomous, adaptive intelligent software system that uses AI models to complete a specific task. AI Agents often interact with users via chat or voice, performing tasks based on the flow of the conversation. To complete tasks on behalf of a human user or another AI agent, they can independently make decisions, execute actions, and interact with other AI agents and tools.

With the expansion of communication over the Internet between AI Agents and external resources that can be tools or other AI Agents, reliable communications across platforms and vendors becomes increasingly important. The role of the agentproto working group is to facilitate such interoperability so that tools and agents can be provided by multiple vendors.

Key Considerations

There are several considerations that are unique to AI Agent applications that need to be addressed while working on developing the building blocks:

- AI Agents act as autonomous software entities that may need to be authenticated independently of the users they represent. Establishing verifiable agent identity that is distinct from user identity enables independent revocation of agent access, scoping of agent permissions to a subset of user permissions, and auditability of agent-initiated actions distinct from user-initiated actions.
- AI Agents possess unique and specialized functional capabilities which can be enhanced by collaboratively working with other agents or tools. This brings new considerations for how these specialized capabilities can be leveraged to select AI agents or tools for collaboration, initiate communication and maintain interactions, including across network boundaries.
- Interactions of AI Agents with users, other AI Agents, and tools can be long-lived, utilize significant amounts of context across various modes (text, audio, video), and require very low latency (including fast barge/interruption times). This introduces new considerations around reliability, transport session management, and data transport.

Charter text 3/9

- To protect data exchanged between AI Agents (and between AI Agents and tools) over potentially untrusted networks, particularly when handling sensitive information (such as personal data or conversational context), mechanisms are required to establish and verify identity, ensure confidentiality, integrity, authenticity of the exchanged data, and delegated authorization across AI Agent chains. This introduces new considerations around protocol-level security and privacy mechanisms.

The scope of the working group includes agent-to-agent and agent-to-tools communication protocols. The working group will document common use-cases to derive requirements for these protocols. Human-agent communication protocols — specifically the protocol-level mechanisms for establishing sessions, negotiating modalities, and exchanging multimodal data between a human user and an AI Agent — are also in scope.

Charter text 4/9

The session protocol will:

- Provide timed (short or long-lived) session management, enabling the establishment, update, context handling, and termination of the services of interacting agents and tools.
- Facilitate highly scalable and reliable session management, capable of surviving network and server failures while supporting graceful recovery.
- Support concurrent exchange of real-time data (such as voice and video), semi-real-time data (such as chat), and non-real-time data (such as tool call inputs and outputs).
- Supports point-to-point and point-to-multipoint communication topologies.

This protocol is expected to be a foundational building block on top of which additional protocols can be built. It is anticipated that the AI Agent session protocol will utilize modern IETF application transfer protocols, such as QUIC, WebTransport, WebRTC or MOQ, based on the anticipated use cases. The protocol must also be usable by other application layer protocols with the appropriate layering and extension points enabling its adoption by any application. Examples of protocols that can utilize this include the existing de facto standard agent communication protocols such as the MCP and A2A protocols being worked on by the Linux Foundation.

AI Agent Protocol Framework (Standards Track)

A standards-track framework that identifies the key building blocks and defines the protocol suite for interoperable agent-to-agent and agent-to-tool communications. The framework provides an architectural overview and highlights areas for subsequent protocol specification work.

This is an evolving work item that can proceed in parallel with the development of specific protocol deliverables associated with the identified architectural blocks. It will iteratively integrate both existing and currently missing protocol building blocks in successive steps, until all core modules are fully incorporated.

Charter text 6/9

The framework will:

- Enable AI Agents to select and collaborate with other AI Agents on the Internet or intranet, deployed in various interconnected domains and ecosystems, to execute simple or complex tasks.
- Allow multi-modal collaboration using varied data formats such as text, images, video, audio, and structured data with exchange of multi-modal contexts.
- Describe the functional blocks, their relationships, and the mechanisms for structured, semi-structured, and multi-modal information exchange to support collaborative tasks across domains.
- Describe agent-specific integration for agent authentication and authorization about how existing and emerging mechanisms are composed and applied in AI agent scenarios, including the confirmation and evidence requirements for AI agent operations.
- Enable an AI Agent to create an independent identity, obtain and exchange access tokens with fine-grained, behavior-driven scopes bound to the specific operations that it is permitted to perform on behalf of the user. Agent authorization needs to account for dynamic behavioral boundaries, including conditional and context-dependent privileges that may vary across interactions and provide a way of requesting confirmation for operations that are about to be performed by AI agents. Any extensions to OAuth protocol mechanisms required to support agent authorization are expected to be developed within the OAuth working group. Any extensions required for independent AI agent identity are expected to be developed within the wimse working group.
- Identify the protocol suite covering session management, transport, security, and identity building blocks.
- The framework may be delivered as multiple standards-track documents, as the working group determines based on the structure and maturity of the building blocks.

Use Cases, Gap Analysis, and Requirements (Informational)

Foundational work will be documented through a set of informational Internet-Drafts covering:

- **Use cases** focused on Agent-to-agent and Agent-to-tool communications, used to verify the suitability of existing protocols and the protocols being developed.
- **Gap analysis and requirements** based on examination of existing de facto standard protocols implemented in open-source projects, from which necessary protocol requirements are derived.

Coordination

This working group is expected to closely coordinate with other related IETF working groups:

- **Security:** Web Authorization Protocol (OAuth), webbotauth, WIMSE — on identity, authorization, and security considerations.
- **Transport:** WebTransport, MoQ, QUIC, TSVWG — on data transport and session management.
- **Discovery and Operations:** INT area, OPS area — on agent discovery and operational considerations.

If the working group needs any changes to or extensions of protocols specified by other working groups, those issues will be raised with the relevant working groups for decisions on how best to handle them. The group is also expected to maintain close communication with open-source projects running under the Linux Foundation.

Out of Scope

The following topics are explicitly out of scope for this working group:

- Implementation details of AI Agents, including definition of AI models, backend AI infrastructure network and protocols, agent reasoning algorithms, or tool-specific business logic.
- Standardization of agent behavior, decision-making, or planning semantics.
- AI agent behavioral security (e.g., preventing the AI model itself from hallucinating, though mitigating the impact of hallucinations via protocol-level user confirmation is in scope).
- The design of human to agent user interfaces, client application UX, or the rendering of agent outputs on end-user devices

Deliverables

The working group will produce the following standards track and informational documents. The work on these deliverables is expected to proceed in parallel.

AI Agent Session Protocol (Standards Track)

A standards-track protocol for creation and maintenance of communication sessions between AI agents, or between AI agents and tools. These sessions allow for the bidirectional exchange of data, including model context, tool call results, and chat messages.

Why session matters - context loss (backup)

- Example Case:

AI Agent



Online Edit Tool



1st Connection (stateful)

1. Open doc
2. Edit doc
3. Modify format

(multi-turn context stored only in this connection)



Connecton Break




2nd Connection

Previous context lost. Must restart from scratch !!!

Agent communication needs coordinated building blocks

Framework defines how the building blocks compose into one continuous agent interaction.



Identity, auth, session, transport
— maps to IETF WGs.


Binds building blocks into one agent scenario

FRAMEWORK

EXAMPLE: CONTEXT LOSS

4 / 9

TASK: "I need to plan a three-city business trip next month. Book travel, prepare meeting briefs, translate the documents for each city, and prepare voice summaries of previous meeting video recordings for a quick refresher when I arrive. Save meeting briefs to my laptop, and send the summaries to my phone."



NEW TASK: ““I changed my mind! Now only book two cities for the same period of time.”

❖ **This leads to many changes**

- Need to cancel flights and hotel for the third city
- Need to modify reservations for the other two cities
- Cancel translations that were needed for the third city

❖ **Discussed further later!**