The Federated Agent Secure Transport (FAST): A Secure Fabric for the Orchestration of Multi-Agent Communication

Keywords: Multi-Agent Systems, Decentralized AI, Secure Communication, Peer-to-Peer Networks, Collective Intelligence

Extended Abstract

To progress from isolated multi-agent simulations to a truly interconnected artificial intelligence, a paradigm shift from centralized control to decentralized communication is necessary. This paper introduces the Federated Agent Secure Transport (FAST), a novel, decentralized framework for Agent-to-Agent (A2A) communication that enables this shift. This architecture stands in contrast to prevailing multi-agent systems like Auto-GPT and CrewAI, which are confined to centralized, monolithic environments. Their reliance on inter-process calls within a single runtime inherently limits scalability, security, and interoperability, precluding the formation of a truly open agent ecosystem. The FAST framework synthesizes two core components: (1) a peer-to-peer protocol for secure, end-to-end encrypted message transport, and (2) a standardized Model Context Protocol (MCP) that serves as a semantic language for task delegation and data exchange. This paradigm, however, also introduces new attack surfaces at the semantic layer, where malicious agents can craft MCP payloads to perform sophisticated prompt-injection attacks on their counterparts. By decoupling the communication layer from the agent's core logic, this architecture allows agents from different entities to interact in a trustless manner. This fosters the conditions for an open market of specialized AI services and enables a new generation of scalable, resilient, and secure collaborative AI, laying the groundwork for a true Internet of Agents. The ability of LLMs to reason and plan has catalyzed a new wave of autonomous agent research. A foundational concept in this area is the "Reason and Act" paradigm, which demonstrates how LLMs can interleave reasoning with tool use to accomplish complex tasks [1]. Landmark research, such as the "Generative Agents" simulation from Stanford, demonstrated how dozens of LLM-powered agents could exhibit complex, emergent social behaviors within a simulated environment [2]. The prevailing monolithic architecture, while effective for contained simulations and single-user applications, presents significant barriers to creating a global, open, and resilient agent ecosystem. First, interoperability is lacking [3–4], as agents are often tightly bound to their native frameworks, resulting in isolated ecosystems that inhibit true composability—where users could flexibly choose the most suitable expert agent from a broader, open marketplace. FAST is built upon three foundational principles that directly address the short-comings of monolithic agentic systems. By removing central points of control, FAST provides a network that is inherently resilient to outages and resistant to censorship or unilateral control by a single entity. Security and privacy are not add-ons but are integral to the protocol. The architecture mandates end-to-end encryption for all substantive communication, ensuring that the network infrastructure itself cannot access or decipher the content of agent interactions. FAST enables agents to be treated as modular, interchangeable components. The broader vision for FAST extends beyond a mere technical specification. It is a blueprint for the essential communication backbone of a future where AI is not a collection of isolated, proprietary tools but a federated, interconnected network. By providing a censorship-resistant fabric for interaction, FAST is designed to foster an open market of AI

services, encouraging innovation and democratizing access to autonomous technology. This represents a crucial step on the journey towards true, emergent collective intelligence.

References

[1] Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., Cao, Y.: ReAct: Synergizing Reasoning and Acting in Language Models. arXiv preprint arXiv:2210.03629 (2022)
[2] Labrou, Y., Finin, T., Peng, Y.: Standardizing Agent Communication Languages: The Current Landscape. IEEE Intelligent Systems, vol. 14, no. 2, pp. 45-52 (1999)
[3] Dähling,S,et al..: Enabling scalable and fault-tolerant multi-agent systems by utilizing cloud-native computing. Autonomous Agents and Multi-Agent Systems, vol. 35, article 10 (2021)
[4] Kanj, H., et al.: A novel dynamic approach for risk analysis and simulation using multi-agents model. Applied Sciences, vol. 12, no. 10, article 5062 (2022)

Table 1. **Title.** Comparative Analysis of Multi-Agent Communication Architectures.

Metric /Attribute	FAST Protocol (Decentralized)	Centralized Orchestrator (Monolithic)
Scalability (Throughput)	**High**. Scales horizontally as more nodes join the network. Load is distributed, preventing a single bottleneck.	**Low**. Limited by the processing and bandwidth capacity of the single central server.
Resilience (Uptime)	**Very High**. No single point of failure. The network remains operational even if a significant portion of nodes fails.	**Low**. The central server is a single point of failure. If it goes down, the entire system is offline.
Latency (Single Message)	**Variable / Higher**. Message must propagate through an indeterminate number of gossips hops to reach its destination.	**Low / Predictable**. direct, two-hop round trip from client to server to client.
Interoperability & Composability	**High**. Any agent that speaks the protocol can join and interact permissionlessly.	**Very Low / None**. Integrating a new agent requires custom code and API changes.
Security Model	**Trustless by Design**. End-to-end encryption is mandatory. Network infrastructure cannot access message content.	**Requires Trust**. central server operator has full access to all unencrypted agent communication and data.
Development & Integration Cost	**Low (Ecosystem)**. Once an agent is FAST-compliant, it can interact with the entire network.	**High (Per Integration)**. Every new agent requires a bespoke, costly integration project with central system.

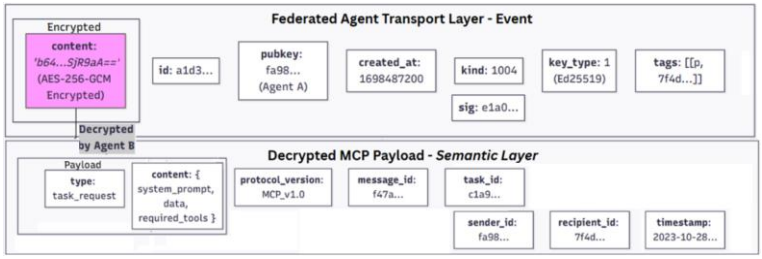

Figure 1. **Title.** The layered structure of a FAST event..

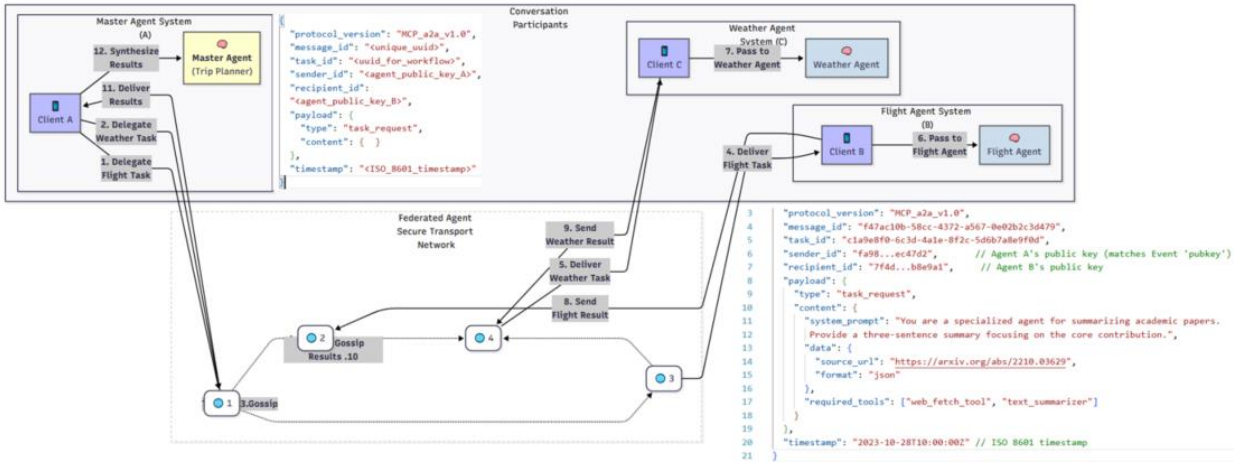

Figure 2. **Title.** Interaction between three autonomous agents and structure of the A2A Protocol (MCP_a2a_v1.0) payload.