On AI Agent Networking

Lixia zhang
UCLA Computer Science Department

IRTF Open, March 17, 2026

40 years of IETF: from networking hosts to networking AI agents

- Agent communication: a new front in networking
- Both about networking: what is not new
  - Communication with different patterns
  - Securing communications
- What is new in networking *AI agents*
- Identify new challenges

Networking: deliver packets, support different communication patterns

- More delivery patterns over time: $i \times j$, $where\ i,\ j \subseteq M \subseteq N$

Pattern	Example	AI Agent Use
1 × 1	TCP unicast, HTTP request/response	Single agent ↔ service call
1 × N	IP multicast, pub/sub broadcast	Orchestrator → *M* parallel agents (fan-out)
N × 1	Sensor aggregation, fan-in	*M* agents → aggregator / decision point
N × N	WebRTC mesh	Multi-agent collaboration, state sharing

Currently, all patterns are implemented as multiple unicast over HTTPS in general

An open research question: whether network layer should natively support $(\mathrm{i} \times \mathrm{j})$ communication

Securing communications, today

- TCP/IP started without security — not negligence, but the threat model was yet to be understood.
  - Connecting a small research community by IP addresses (all hosts got one)
  - Aiming to provide universal reachability
  - DNS deployed quickly, but only *servers* got names
- Security demands arose as the Internet commercialized
  - Must secure e-commerce transactions → server authentication and channel encryption quickly bolted onto TCP connections
  - Spam, spoofing, DDoS and more → added all kinds of blockage, without touching TCP/IP
  - Different solutions make use of different identifier spaces.
- Today: a set of fragmented security solutions at different layers
  - TLS+Web PKI with CT, OAuth/OIDC, Kerberos, DNSSEC, JWT — each uses its own defined identifier mechanism.

Security: a late addon → no unified namespace for users, machines, apps → fragmented solutions

Scale, dynamics, and multiparty interactions

1. Agents at *scale* ($10^9 \sim 10^{12}$?)
   - According to AI
     - ~5.5B users on the Internet
     - Total # of devices: 40~50B
     - We are already handling multibillions, what's new
2. *Dynamics* in agent lifetime span ($msec \sim days, months$?)
   - Not today's devices
3. $(i \times j)$ multi-agent *interactions*
   - None of current users nor devices
     - They all talk to the clous
     - They do not *directly* interact

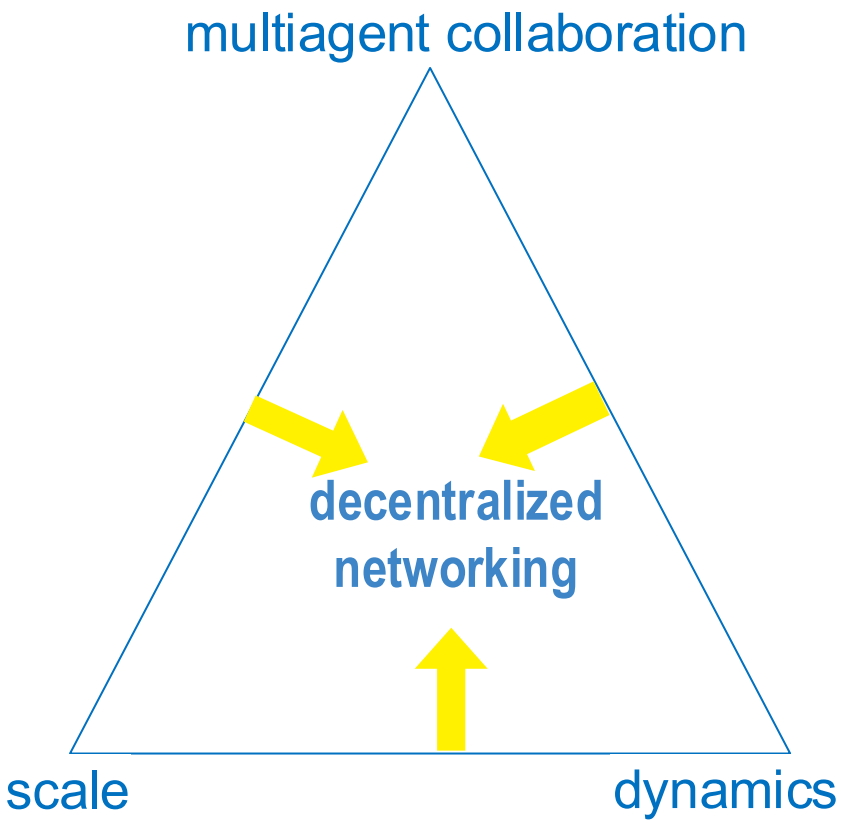

Agent networking patterns

Scale, dynamics, and interaction: security impact

1. Agents at *scale* (109 – 1012?)
2. *Dynamics* in agent lifetime span (mSec ~ days, months?)
3. $(i \times j)$ multiagent interaction
4. Delegation chains: act on behalf of users, agents spawn sub-agents, delegate tasks
   - Multi-hop delegation is a baseline requirement.
5. Autonomy: actions without human approval
   - Action without human approval: Higher security demands: actions are taken without per-action human sign-off.

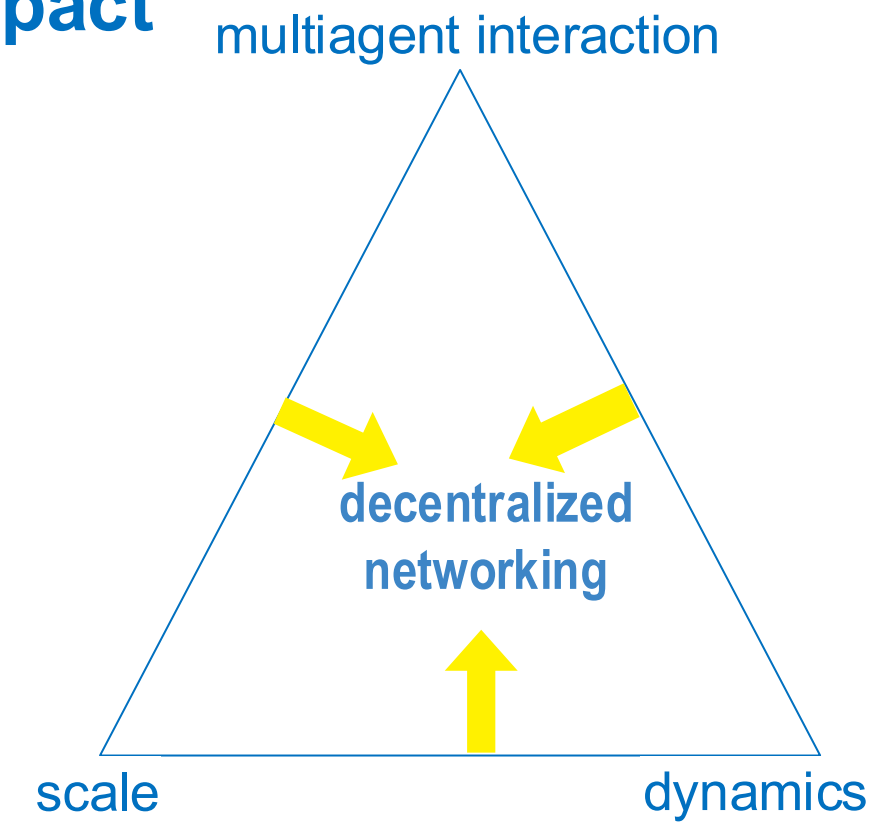

Agent networking patterns

Orders of magnitude beyond anything the current security solutions were designed for

Physical Agents: a specific example of the new agent frontier

- IETF agentic AI discussions so far: largely focused on software agents
  - Absence of attention to physical agents
- How physical agents differ from model agents
  - **Multiparty communication as the norm**
    - A warehouse robot fleet requires continuous coordination: position sharing, task assignment, collision avoidance — all simultaneously between ($i \times j$) participants.
  - **Real-time state synchronization**
    - Surgical robots, autonomous vehicles, drone swarms: shared world-state must be consistent across all agents with bounded latency.
    - A delayed communication can lead to physical harm.
  - **Direct physical interaction with human users** →Physical consequence of failure
    - A compromised software agent corrupts data.
    - A compromised physical agent causes physical harm.

Current Security Practice: Infeasible at Agent Scale

- No unified namespace
  - X.509 DNs, OAuth client IDs, SSH keys, API tokens — each system defines its own identity authority, operating in isolated domains.
    - No principled trust reasoning across boundaries.
- Certificate issuance
  - CA-signed certs assume slow, manual provisioning.
  - Billions of ephemeral agents cannot wait for CA issuance cycles.

  And let's skip the cert revocation challenges and failure consequences ....
- OAuth / bearer tokens
  - Designed for user-facing web flows.
  - Not suited for agent-to-agent delegation chains or scoped sub-agent authority.

Agents scale, dynamics, and multiparty interactions expose structural failures of existing security solutions — not just performance limits

Starting point: get naming right, then build security into agentic AI

- **DNS as a unifying namespace for cyberspace**
  - DNS: decentralized name management with TLD coordination and name delegation
  - DNS: offering *globally unique*, *semantically meaningful* names
  - Every entity — organization, user, agent, and service — have DNS names as primary identifiers
  - Semantic meaningfulness makes trust reasoning human-navigable
- **Identity = Name + Key**
  - Name provides semantic context; key provides cryptographic verifiability
  - Trust chains must be human-navigable; machine-verifiability alone is not adequate
- **Crypto protections anchor on local trust**
  - Global namespace, local trust
    - AI agents make this a scaling requirement, not an architecture option
  - Delegation chains are native infrastructure: scoped, verifiable, multi-hop

Tooling for scalable, decentralized networking and security

- **Developing namespace management support to enable**
  - Every organization, user, and agent manages their own namespace slice
  - Every entity can run their own DNS services
  - dynamic agent name lifecycle, per-entity DNS operation at scale.
- **Developing decentralized trust management to enable**
  - Every administrative entity manages their own trust anchors
  - Mutual authentication among peers through cross-certification
    - Scale via *trusted* intermediaries
- **Developing full support for 3 As**
  - Authentication: name ownership
  - Authorization: fine-grained security policies for who can do what, under what conditions, *with least privilege*
  - Audit for accountability

What the IRTF Community Should Do

- Resist the temptation to patch existing identity protocols for agents
  - Although may be needed as short term stopgap, patching inherits structural flaws and amplifies them at the agent scale
- Treat **DNS as the authoritative identity namespace** for all, including agent — not one option among many, but the foundation
- Adopt localized trust anchors built on DNS names as the model to decentralize trust, hence decentralize control power
- Design delegation chains as baseline infrastructure — **deep, dynamic, and verifiable by default**
- **Engage regulators early** — provide technical grounding for what to regulate before market force locks the decision in.

AI agent networking: Global Namespace, Local Trust

- What is new in networking: the requirement triangle
  - Scale (billions to trillions), dynamics (ephemeral lifetimes), and multiparty realtime interactions
- Today's fragmented identity stack is a structural failure
  - Absence of a unifying namespace — at agent scale, this failure becomes foundational.
- New solution starts with the right naming foundation
  - DNS as a unifying namespace; Name + key binding works for ephemeral agents. Localized trust with delegation chains handles dynamic agent spawning natively.
- Two research investment suggestions
  - Namespace management tooling for agent lifecycle at scale
  - Decentralized trust management for principled delegation and policy enforcement.

Agent communication is the new front in networking

The community has a rare chance to get the foundation right — naming, security, and trust at scale — before the patterns of the past repeat themselves at the agent scale.

Lixia zhang
UCLA, lixia@cs.ucla.edu