

Agent Interaction & Delegation Protocol (AIDP)

[draft-vandoulas-aidp-02](#)

Ioannis Vandoulas

Why AIDP?

- **Real-world Adoption:** Agents are already used in high-stakes sectors
- **The Compliance Gap:** New regulations (e.g., EU AI Act) mandate safety and auditability, but standardized enforcement tools are missing.
- **Intrinsic Risks:** LLM hallucinations are architectural, not incidental; they lead to non-deterministic execution.
- **Need for a "Safety Fuse":** We must decouple non-deterministic reasoning from deterministic system execution.

What AIDP Specifies

- *Intent Envelope (governed action request)*
- *Capability-based authority model*
- *Delegation chains with strict subsetting*
- *Execution Boundary as mandatory enforcement point*
- *Observation binding (execution → reasoning)*
- *Replay protection & revocation semantics*
- *Transport-agnostic (HTTP binding included)*

Relation to Existing Work

Builds on concepts from:

- GNAP
- OAuth RAR
- Capability systems (ZCAP / UCAN)
- SCIM

Does NOT replace them

Focuses on:

- agentic execution lifecycle
- delegation safety
- observation binding
- deterministic control loops

Hoped dispatching outcome

- **Primary Venue:** Dispatch to the **GNAP WG** (or a potential **Agentic Networking BoF**) to integrate AIDP's core architectural principles.
- **Alternative:** Inclusion of AIDP concepts into broader **IETF-wide discussions on AI/Agentic protocols** to prevent fragmented security models.
- **Goal:** Contribute the AIDP framework as a baseline for **governed agentic execution**, prioritizing architectural alignment over standalone protocol track.