WG Name:
Agentic AI Identity and Access Management Working Group (AI-IAM WG)

Purpose:
To define a standardized framework for **identity and access management for agentic AI systems** in enterprise environments. Existing SCIM Agent Schema drafts provide a foundation, but enterprises need extensions for delegation, continuous authorization, audit, and discovery to manage AI agents acting autonomously or on behalf of users, JIT for AI agents.

Initial Scope:

- **Profile the SCIM Agent Schema for enterprise use:** Today the SCIM Agent schema only define very few attributes of the AI agent. But it doesn't extend beyond these attributes. i.e. the Schema doesn't define attributes such as
    - On whose behalf the AI agent can act i.e., who is the delegated authority
    - Is there any ITSM ticket / change management reference
    - What is the privilege level of this agent
    - What is the scope of actions this agent can do

  The spec also defines any constraints on the expiry and the lifecycle state of the AI agent.

  Hence the idea is to profile SCIM to add additional attributes and define constraints and interoperability rules.

- **Define OAuth/OIDC extensions for agent delegation and chaining:** The gap is that there is no chaining or agent delegation model for enterprise AI.
    - The scope is to provide a standard delegation context beyond the token exchange claim set. Enterprises need richer context like who approved this delegation? For how long it is valid? ITSM ticket etc., There is no standard vocabulary for this.
    - Agent chaining supporting multiple hops, while maintaining the delegation chain without losing the full path. (like token chaining) Ex: user -> AI agent -> workflow agent -> service, the service should know the full path.

- **Specify mechanisms for continuous authorization and dynamic policy enforcement:**
    - Includes runtime assurance by continuously granting access, fine grained controls (step-up, revoke etc.,)

    - Discover additions such as agent introspection endpoint, event streams, policy modes, step up modes (CIBA, device flow), token sender constraints etc.,
    - Additional parameters in the token model like purpose, delegation chain, session IDtask ID etc.,
    - Introspection and heartbeats or per-action checks
    - Revocation or descoping of the scopes
    - Step-up during execution
    - Dynamic consents – Human in the loop
    - Cascading revocation in case of delegation chains
- **Create standard audit and compliance claims for agentic actions:** Standardized audit claims: Like who is the approver, approved time, policy id/ticket etc., Without standard audit claims, every vendor's logs look different, and you can't federate or correlate across IdPs.
- Profile **sender-constrained token usage** (mTLS, DPoP) for agents.
- Out of scope: AI ethics frameworks; human-only IAM flows.

(iv) Proposed Specifications / Deliverables:

- *Enterprise Profile of SCIM Agent Schema*.
- *OAuth/OIDC Extensions for Agent Delegation & Chaining*.
- *Audit & Logging Claims for Agent Actions*.
- Conformance test suite.

(v) Anticipated Audience / Users:

- IdPs (Entra ID, Okta, CyberArk Identity, Ping, Curity).
- Enterprises (finance, healthcare, government, SOCs, DevOps).
- Vendors (PAM, SIEM, ZTNA, API gateways, agent frameworks).
- Agent providers (OpenAI, Anthropic, Azure AI, AWS Bedrock).

(vi) Language:
English.

(vii) Method of Work:

- Public WG mailing list (specs@openid.net).

- GitHub Work Group repository.
- Monthly virtual meetings; optional face-to-face at OIDF/IETF events.

(viii) Basis for Completion:

- Two independent, interoperable implementations per deliverable.
- Consensus that scope has been achieved.
- Specifications approved through OIDF/IETF process.

(b) Background Information

(i) Related Work & Why Needed:

- **SCIM Agent Schema (IETF draft 2025):** Defines basic agent identity attributes (id, owner, lifecycle, expiry).
    - *Gap:* No enterprise extensions (delegation, compliance, policy enforcement).
- **OAuth/OIDC (Device Flow, CIBA, DPoP, mTLS):** Core building blocks exist.
    - *Gap:* No chaining or agent delegation model for enterprise AI.
- **CSA AI IAM Guidance:** High-level principles, no concrete protocol specs.

(ii) Proposers (≥5 OIDF Members):

- Monika Avalur (CyberArk) – Editor (monika.avalur@cyberark.com)
- Yuval Glasner (CyberArk) – Contributor(Yuval.glasner@cyberark.com)
- Wei Dai (CyberArk) - Contributor (Wei.Dai@cyberark.com)
- Abhijit Borgohain – Contributor (Abhijit.Borgohain@cyberark.com)
- Venkatesh Dunna – Contributor (Venkatesh.Dunna@cyberark.com)

(iii) Anticipated Contributions:

- Draft: *Enterprise Profile of SCIM Agent Schema*.
- Draft: *OAuth/OIDC Agent Delegation and Chaining*.
- Draft: *Agent Audit & Compliance Claims*.