The De-Provisioning Gap in Identity Management for Agentic AI

Fernanda Lauck Jablonski[1], Fernanda L. Müller[2], Julio Ribeiro Sant Ana[1], Yuri H. Lima Schoenardie[1], Kléber H. Cardoso da Silva[1], Frederico Schardong[2]

[1] Instituto Federal do Rio Grande do Sul (IFRS)

[2]Universidade Federal de Santa Catarina (UFSC)

{flj, jrsa, yhlls, khcs}@cesar.org.br
frederico.schardong@ufsc.br, muller.larissa@grad.ufsc.br

Abstract. *Various proposals have emerged to address challenges related to identity, authentication, and authorization in agentic systems. Given this context, this article compiles key proposals, such as OIDC-A, SPIFFE, A-JWT, A2A, and LDP, and compares them based on five criteria derived from the "Identity Management for Agentic AI" white paper. The results show that, although current proposals are robust in specific areas, none satisfies all analyzed criteria. The main limitation is the lack of rapid, verifiable de-provisioning, highlighting challenges in the reliable revocation of identities and permissions in agentic AI.*

1. Introduction

Most authentication and authorization protocols and frameworks were designed for contexts where behavior is predictable [Bandi et al. 2025]. Such protocols and frameworks prove inadequate when faced with Artificial Intelligence (AI) agents capable of interacting with multiple systems, delegating tasks, and operating across different trust domains. This technological advancement necessitates the evolution of current security systems.

This article analyzes some of the selected proposals for identity, authentication, and authorization in agentic systems. They are OpenID Connect for Agents (OIDC-A) [Nagabhushanaradhya 2025], the Agent-to-Agent (A2A) [Google 2026] protocol, Identity-Aware Protocol for Multi-Agent LLM Systems (LDP) [Prakash 2026], Agentic JWT (A-JWT) [Goswami 2025], and Secure Production Identity Framework for Everyone (SPIFFE) [SPIFFE 2026]. This paper is organized as follows. Section 2 introduces a short background, while Section 3 presents the method adopted in the analysis detailed in Section 4. Finally, Section 5 concludes the paper.

2. Identity Management for Agentic AI

The white paper "Identity Management for Agentic AI" [South et al. 2025] discusses current challenges regarding the authentication, authorization, and governance of AI agents. According to the authors, AI agents are systems capable of making decisions and performing actions to achieve specific goals. Unlike traditional programs, they can operate autonomously, use external tools, connect to Application Programming Interfaces (APIs), interact with other agents, and execute tasks on behalf of users.

The authors argued that current identity mechanisms were developed primarily for human users, rendering them insufficient to handle ecosystems composed of autonomous agents. The white paper identifies 43 challenges related to identity in agentic ecosystems. To make the comparative analysis feasible, these challenges were consolidated into five

broader analytical categories according to their common objectives and characteristics: (i) *Scalable governance* refers to the ability to manage agents at scale, including auditing, permission management, and lifecycle management; (ii) *Verifiable delegated authority (On-Behalf-Of)* is the ability to prove which user or entity an agent is acting on behalf of; (iii) *Scope attenuation* is the possibility of progressively reducing permissions along delegation chains; (iv) *Native and portable identity* is the ability to assign an agent a distinct identity that is confirmable and interoperable across different systems; and (v) *Verifiable and rapid de-provisioning* is the ability to revoke permissions reliably, immediately, and permanently, with propagation throughout the ecosystem.

3. Method

We analyzed proposals for identity management in agentic AI that the research group had previously studied and considered representative of current work on agent identity, authentication, and authorization. Each proposal was assessed against the five criteria introduced in Section 2 and classified as *Yes* when it explicitly addresses a criterion, *Partially* when it offers only a limited or incomplete solution, and *No* when the criterion is not addressed. The first, third, fourth, and fifth authors produced the initial assessment and resolved divergences by consensus through technical discussion; the second and sixth authors reviewed the outcome. Two threats to validity follow from this design: the lack of a systematic search protocol with explicit inclusion and exclusion criteria, which limits the coverage and replicability of the selection; and the small number of reviewers per proposal, which leaves the classification more exposed to individual judgment.

4. Comparative Analysis

No current proposal fully satisfies all criteria. While some solutions perform strongly in specific areas, all exhibit limitations in at least one aspect. The comparative analysis result is laid out in Table 1. Next, we explain the rationale behind each approach's *No* or *Partially* classification. We adopt the concept of agent identity proposed by [Garzon et al. 2025], *i.e.,* the ability of an agent to authenticate itself to other agents through the use of Decentralized Identifier (DID) and Verifiable Credentials (VCs).

Table 1. Comparative Analysis of Agentic AI Protocols and Frameworks.

Protocol/ Framework	Scalable Governance	Delegated Authority	Scope Attenuation	Native/Portable Identity	Rapid De-Provision
OIDC-A	Yes	Yes	Yes	Partially	No
A2A	Partially	Partially	No	Partially	No
LDP	Yes	Yes	No	Partially	No
A-JWT	Partially	Yes	No	Partially	No
SPIFFE	Partially	No	No	Yes	No

4.1. OIDC-A

Although the OIDC-A [Nagabhushanaradhya 2025] introduces agent-related claims, *Native and Portable Identity* remains strongly tied to contextual attributes such as model, user, and execution environment. Consequently, different agents performing distinct functions may still be represented through similar claims, limiting the notion of a persistent

and independent agent identity. Moreover, the OIDC-A does not satisfy *Verifiable and Rapid De-Provisioning*, as it does not define trust or reliability claims for revocation of trust, identity, or permissions in a verifiable manner.

4.2. A2A

The A2A [Google 2026] protocol provides a standardized framework for interoperability and communication among autonomous agents. However, core identity and security functions, including authentication, authorization, trust management, and credential lifecycle management, are intentionally delegated to external protocols and underlying security infrastructures. As a result, while A2A can support aspects of *Scalable Governance*, *Verifiable Delegated Authority*, and *Native and Portable Identity*, these capabilities are not natively implemented by the protocol. Furthermore, *Scope Attenuation* and *Verifiable De-Provisioning* are not explicitly addressed.

4.3. LDP

The LDP [Prakash 2026] does not satisfy *Scope Attenuation*. Although the protocol distributes tasks among agents according to their capabilities and the complexity of the requested task, it does not provide mechanisms for progressively reducing permissions throughout a delegation chain. Instead, delegation focuses on task allocation, and delegated agents may continue to operate under the same privileges. *Native and Portable Identity* is only partially satisfied. LDP provides native agent identities, including descriptive attributes that characterize agent capabilities and roles. However, these identities are not designed to be portable across heterogeneous systems and trust domains. Finally, LDP does not address *Verifiable and Rapid De-Provisioning*, as the proposal does not define mechanisms for credential revocation, credential rotation, permission invalidation, or verifiable propagation of de-provisioning decisions.

4.4. A-JWT

The A-JWT [Goswami 2025] protocol partially satisfies *Scalable Governance*. Each workflow must be individually registered with the Identity Provider (IdP), reducing scalability. *Scope Attenuation* is unmet because multi-agent clients must use a single coarse-grained client-level scope, even when agents require different privileges. *Native and Portable Identity* is partial because agent-specific identifiers and metadata remain tied to the underlying IdP and trust infrastructure, limiting portability across independent domains and describing agent context rather than an independent identity. *Verifiable and Rapid De-Provisioning* is not addressed, as the proposal does not define mechanisms for propagating revocation decisions or producing verifiable evidence that delegated permissions and identities have been invalidated.

4.5. SPIFFE

The SPIFFE [SPIFFE 2026] does not satisfy *Verifiable Delegated Authority* nor *Scope Attenuation*. SPIFFE focuses on workload identity and authentication rather than representing delegation relationships or enforcing progressive permission reduction across delegation chains. *Scalable Governance* is only partially satisfied, as the framework provides identity lifecycle management, credential issuance, rotation, and federation mechanisms, but does not natively address authorization governance, auditing, or delegation management. *Verifiable and Rapid De-Provisioning* is not addressed, since identity removal relies on expiration and renewal cycles rather than immediate, verifiable revocation.

5. Conclusion

This study analyzed proposals focused on identity and authorization in agentic systems. As noted throughout the article, despite the strengths of current proposals in various aspects, many still fail to meet the requirements outlined in the OpenID Foundation's white paper. Among the criteria analyzed, the unanimous shortcoming of analysed approaches is the lack of verifiable and rapid de-provisioning, as this requirement is not addressed in any of the current proposals. This limitation highlights that, notwithstanding advances in areas such as authentication and delegation, significant gaps remain in building truly secure and scalable agentic ecosystems. Future work should investigate mechanisms capable of providing verifiable revocation across trust domains, propagating revocation decisions through delegation chains, and producing auditable evidence that identities and permissions have been effectively invalidated throughout an agent ecosystem. Additionally, research is needed on supervisory mechanisms capable of understanding user intent and detecting deviations in agent behavior, enabling more reliable decisions regarding permission revocation and identity de-provisioning.

Acknowledgments

This work has been fully/partially funded by the project GIAAI - Identity and Access Management for Intelligent Agents - Phase 1 supported by Centro Integrado de Segurança em Sistemas Avançados (CISSA), with financial resources from the PPI IoT/Manufatura 4.0 of the MCTI grant number 08/2-24, signed with EMBRAPII.

References

Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., and Vilipala, S. V. (2025). The rise of agentic ai: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges. *Future Internet*, 17(9):404. https://doi.org/10.3390/fi17090404.

Garzon, S. R., Vaziry, A., Kuzu, E. M., Gehrmann, D. E., Varkan, B., Gaballa, A., and Küpper, A. (2025). Ai agents with decentralized identifiers and verifiable credentials. *arXiv preprint*. https://arxiv.org/abs/2511.02841.

Google (2026). Agent2Agent (A2A) Protocol. https://a2a-protocol.org/.

Goswami, A. (2025). Agentic jwt: A secure delegation protocol for autonomous ai agents. *arXiv preprint*. https://arxiv.org/abs/2509.13597.

Nagabhushanaradhya, S. (2025). Openid connect for agents (oidc-a) 1.0: A standard extension for llm-based agent identity and authorization. *arXiv preprint*. https://doi.org/10.48550/arXiv.2509.25974.

Prakash, S. (2026). Ldp: An identity-aware protocol for multi-agent llm systems. https://arxiv.org/abs/2603.08852.

South, T. et al. (2025). Identity management for agentic ai: The new frontier of authorization, authentication, and security for an ai agent world. *OpenID Foundation Whitepaper*. https://doi.org/10.48550/arXiv.2510.25819.

SPIFFE (2026). Spiffe documentation. https://spiffe.io/docs/latest/.