SPOTLIGHT

Sponsored by: Teleport

Model Context Protocol is quickly becoming the de facto standard for agentic AI. The market for MCP identity solutions will enter a rapid growth phase, and companies are looking to address least privilege.

Putting an "S" in MCP: Delivering Identity to Agentic AI Implementations Securely

October 2025

Written by: Frank Dickson, Group Vice President, Security and Trust

AT A GLANCE

KEY TAKEAWAY

Implementing identity within the MCP involves more than just authentication. It requires a comprehensive framework for delegation, verification, and auditability. The MCP identity extension (MCP-I) introduces these capabilities by layering cryptographic identity and delegation mechanisms on top of the base protocol.

Introduction

Model Context Protocol (MCP) has quickly become the de facto standard to facilitate communication with AI agents and their environments. It defines how agents share context, state, and intent with each other and with external systems. MCP is especially important in multiagent systems where coordination, memory, and shared understanding are vital.

Introduced by Anthropic in November 2024, MCP is an open standard protocol designed to enable smooth communication among AI models, tools, and data sources. MCP aims to standardize the integration of large language model (LLM) applications with external systems such as databases and APIs to eliminate the need for bespoke integrations. MCP quickly gained traction, with major players such as Microsoft, OpenAI, AWS, and Google adopting the protocol. Vendors released supported versions of MCP servers, while Anthropic provided reference implementations for tools such as PostgreSQL, Google Maps, GitHub, and Slack.

The protocol uses a client/server architecture and enables hosts (AI applications) to manage MCP clients that connect to multiple MCP servers and expose functionalities in a standardized way. The focus of MCP is on connectivity and communication; security was not emphasized in the initial standard and has instead been making its way in after MCP's wide adoption. IDC sometimes quips that the "S" in MCP stands for security, as MCP currently lacks robust security specifications, which presents a significant challenge for enterprise adoption. Without built-in mechanisms for authentication, authorization, and identity management, MCP-based systems are vulnerable to impersonation, unauthorized access, and data leakage.

The lack of identity standards in agentic AI is concerning. As AI agents become more autonomous and embedded in business workflows, identity becomes foundational for:

- **Authentication:** Verifying that an agent is who it claims to be
- **Authorization:** Ensuring agents only access resources they are permitted to

- **Auditing:** Tracking agent actions for compliance and security
- **Life-cycle management:** Managing creation, updates, and decommissioning of agents

Without strong identity frameworks, agentic systems risk becoming opaque and insecure, especially in enterprise environments where agents may act on behalf of users or other systems.

Benefits of an Identity Platform in Implementing MCP

As agentic AI systems become more autonomous and embedded in enterprise workflows, the question of identity moves from a technical detail to a foundational concern. MCP, often described as a "USB-C for AI," enables AI agents to interface with external systems such as databases, APIs, and cloud services. While this connectivity is powerful, it introduces a new class of nonhuman identities (NHIs) that require careful governance.

Unlike traditional software integrations, agentic AI systems act independently, and by design, make decisions and execute tasks without human oversight. This autonomy means that each agent must be treated as a distinct identity, one that holds credentials, accesses sensitive resources, and operates across environments. Without robust identity mechanisms, organizations risk unauthorized access, audit blind spots, and compliance violations.

MCP itself does not define identity ownership or life-cycle management. It relies on existing standards, such as OAuth 2.1, that were designed for human users. This mismatch creates gaps in visibility and control. For example, credentials may be hardcoded, privileges may be overly broad, and audit logs may fail to capture agent behavior. The result is a growing "identity explosion" in which AI agents proliferate without the guardrails traditionally applied to human users.

To address this, most companies have to reimagine identity architecture so that AI identities operate seamlessly with humans and nonhumans, depending on whether they're executing as human delegates or as workloads. Agents should have cryptographically verifiable identities, scoped permissions, and clear delegation chains. They should be subject to the same principles of least privilege, credential rotation, and behavioral monitoring that govern human access. In short, securing MCP isn't just about protecting data; it's about knowing who (or what) is acting on the organization's behalf and ensuring they have the authorization to do so.

Implementing identity within the MCP involves more than just authentication. It requires a comprehensive framework for delegation, verification, and auditability. The MCP identity extension (MCP-I) introduces these capabilities by layering cryptographic identity and delegation mechanisms on top of the base protocol.

At the heart of MCP-I is the concept of verifiable identity. AI agents must be able to prove who they are, who authorized them, and what they're allowed to do. This is achieved through cryptographically signed identifiers and verifiable credentials, which serve as tamperproof attestations of permissions. These credentials can be issued by users, organizations, or trusted intermediaries and are validated by services before granting access.

Delegation is another critical component. Rather than granting blanket access, users can delegate specific scopes of authority to agents. This delegation is explicit, time bound, and auditable, ensuring that agents operate within well-defined boundaries. The delegation chain — from user to agent to service — is transparent and traceable, reducing the risk of privilege escalation or misuse.

Audit mechanisms round out the identity framework. Every action an agent takes can be logged and reviewed. This visibility is essential for compliance, incident response, and trust. It also enables organizations to enforce policies such as least privilege and credential rotation, which are often overlooked in AI deployments.

Finally, interoperability is key. MCP-I works across platforms and vendors to enable consistent identity management in heterogeneous environments. Whether an agent is booking a flight, accessing health records, or filing taxes, MCP-I ensures secure and uniform handling of identity and authorization. In essence, it transforms AI agents from opaque executors into accountable participants in enterprise systems.

Considering Teleport

Teleport's infrastructure identity platform offers a compelling approach to securing identity in MCP implementations, particularly as organizations grapple with the rise of agentic AI. Rather than reinventing the wheel, Teleport adapts its existing identity-first infrastructure model to the challenges AI agents pose.

Teleport treats AI agents as first-class identities, much like human users or service accounts. This means agents are subject to the same governance policies, which include short-lived credentials, role-based access control, and audit logging. By integrating with MCP, Teleport aims to ensure that agents accessing enterprise resources can do so with verifiable identity and scoped permissions.

One of Teleport's key contributions is the elimination of static credentials. In traditional setups, agents often rely on hardcoded tokens or long-lived secrets, which pose significant security risks. Teleport replaces these with ephemeral certificates tied to identity, reducing the attack surface and simplifying credential rotation. This is particularly valuable in dynamic environments where agents spin up and down frequently.

Teleport also enforces least privilege by mapping agent roles to specific access scopes. This ensures that agents can only perform tasks they've been explicitly authorized to do, such as querying a database, invoking an API, or modifying cloud infrastructure. Combined with detailed audit trails, this approach provides full visibility into agent behavior, making it easier to detect anomalies and enforce compliance.

Finally, Teleport's integration with cloud-native platforms allows organizations to scale MCP securely across distributed environments. By embedding identity into the infrastructure layer, Teleport transforms MCP from a connectivity protocol into a secure, governed interface for AI agents. This pragmatic solution leverages existing security models to meet the demands of a rapidly evolving AI landscape.

Challenges

Implementing Teleport for MCP introduces several challenges that organizations must navigate to ensure a secure and scalable deployment. Agents often operate autonomously and require dynamic, short-lived credentials, which can necessitate people and process adaptations to the traditional identity approaches for teams that have not yet already deployed Teleport for other use cases. Additional use cases are developing and may require features such as capabilities delegation, which are more nuanced than straightforward machine access use cases.

Another challenge lies in integrating Teleport with existing cloud environments and identity providers. While Teleport supports a broad variety of technology types, there may be older legacy systems that cannot accommodate a certificate-based access model.

Change management is a concern as well. Teleport's security model emphasizes ephemeral access and audit logging, which will require process change for teams unfamiliar with just-in-time access or observability tooling. Ensuring that logs are properly leveraged for identifying anomalies or providing audit evidence may alter existing process; Teleport does offer identity security to eliminate integration needs with observability tooling.

Finally, governance and compliance frameworks may lag behind the technical capabilities. Defining policies for agent delegation and audit review requires cross-functional collaboration between security, IT, and legal teams. Without clear standards, organizations risk inconsistent enforcement, which can reduce trust in AI-driven operations.

Conclusion

Agentic AI is clearly the future. Thus IDC believes the market for MCP identity solutions will enter a rapid growth phase, and companies are looking to address least privilege within MCP implementations. As Teleport addresses the challenges this paper describes, the company has a significant opportunity for success.

About the Analyst

Frank Dickson, *Group Vice President, Security and Trust*

Frank Dickson is the group vice president for IDC's Security and Trust research practice. In this role, he leads the team that delivers compelling research in the areas of AI security; cybersecurity services; information and data security; endpoint security; trust; governance, risk, and compliance; identity and digital trust; network security; privacy and legal tech; and application security and fraud.

MESSAGE FROM THE SPONSOR

Teleport is the Infrastructure Identity Company, modernizing identity, access, and policy for infrastructure, improving engineering velocity and infrastructure resiliency against human factors and compromise. The Teleport Infrastructure Identity Platform integrates access management, zero trust networking, identity governance, and identity security into a single platform, enabling organizations to access and protect both AI and classic infrastructure, including servers, databases, cloud consoles, GitHub, Windows desktops, Kubernetes, web applications, and MCP servers. Teleport secures MCP to govern interactions between AI/LLMs and your databases/MCP servers, delivering the access control, RBAC/ABAC, and auditability you need to secure your data. Visit us for more information.

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
blogs.idc.com
www.idc.com

IDC Custom Solutions produced this publication. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis that IDC independently conducted and published, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. This IDC material is licensed for external use, and in no way does the use or publication of IDC research indicate IDC's endorsement of the sponsor's or licensee's products or strategies.

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives.