descope

Auth Best Practices for AI Agents / MCP Servers

Current State of MCP Auth

MCP Architecture Overview

descope

A “universal remote” to help LLMs connect in a standardized manner with external tools over APIs

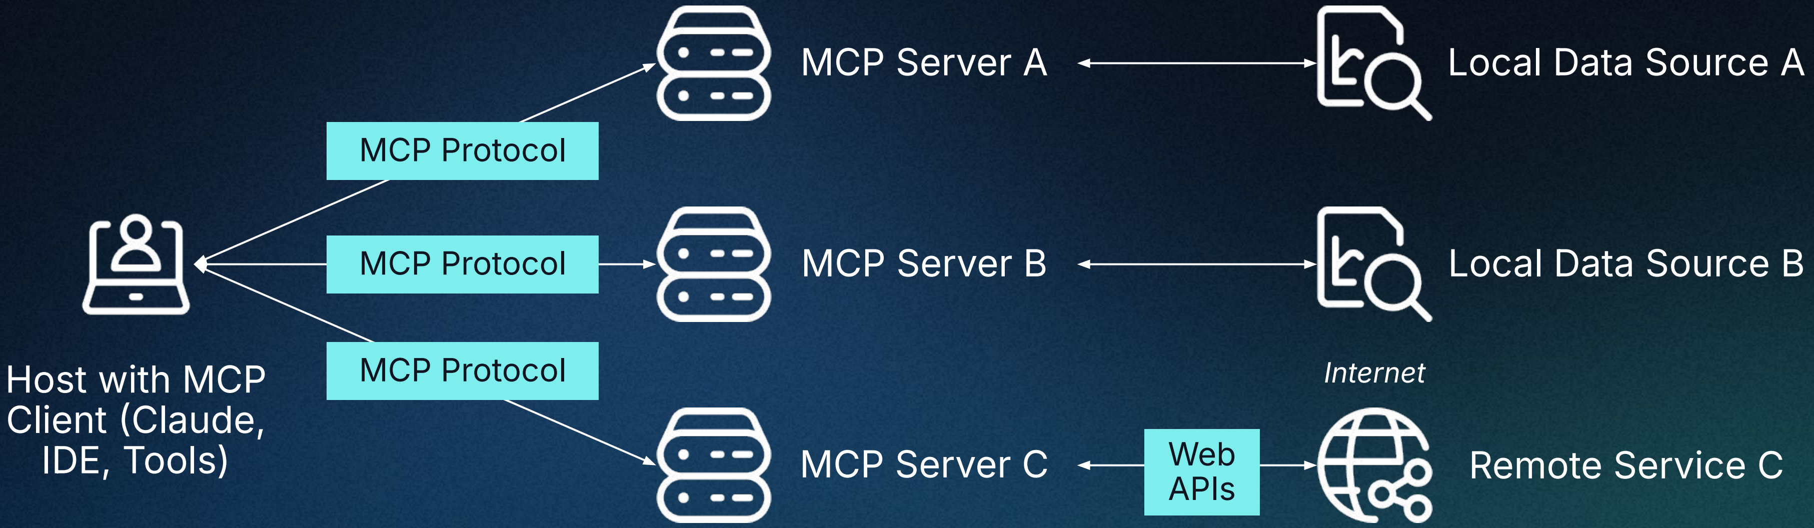

MCP in Numbers

descope

7000+

MCP server ecosystem

(Source: PulseMCP)

~2000

MCP servers with no security

(Source: Dark Reading)

13%

of remote MCP servers being deployed

(Source: Descope)

MCP Use Case 1: Internal MCP Servers

descope

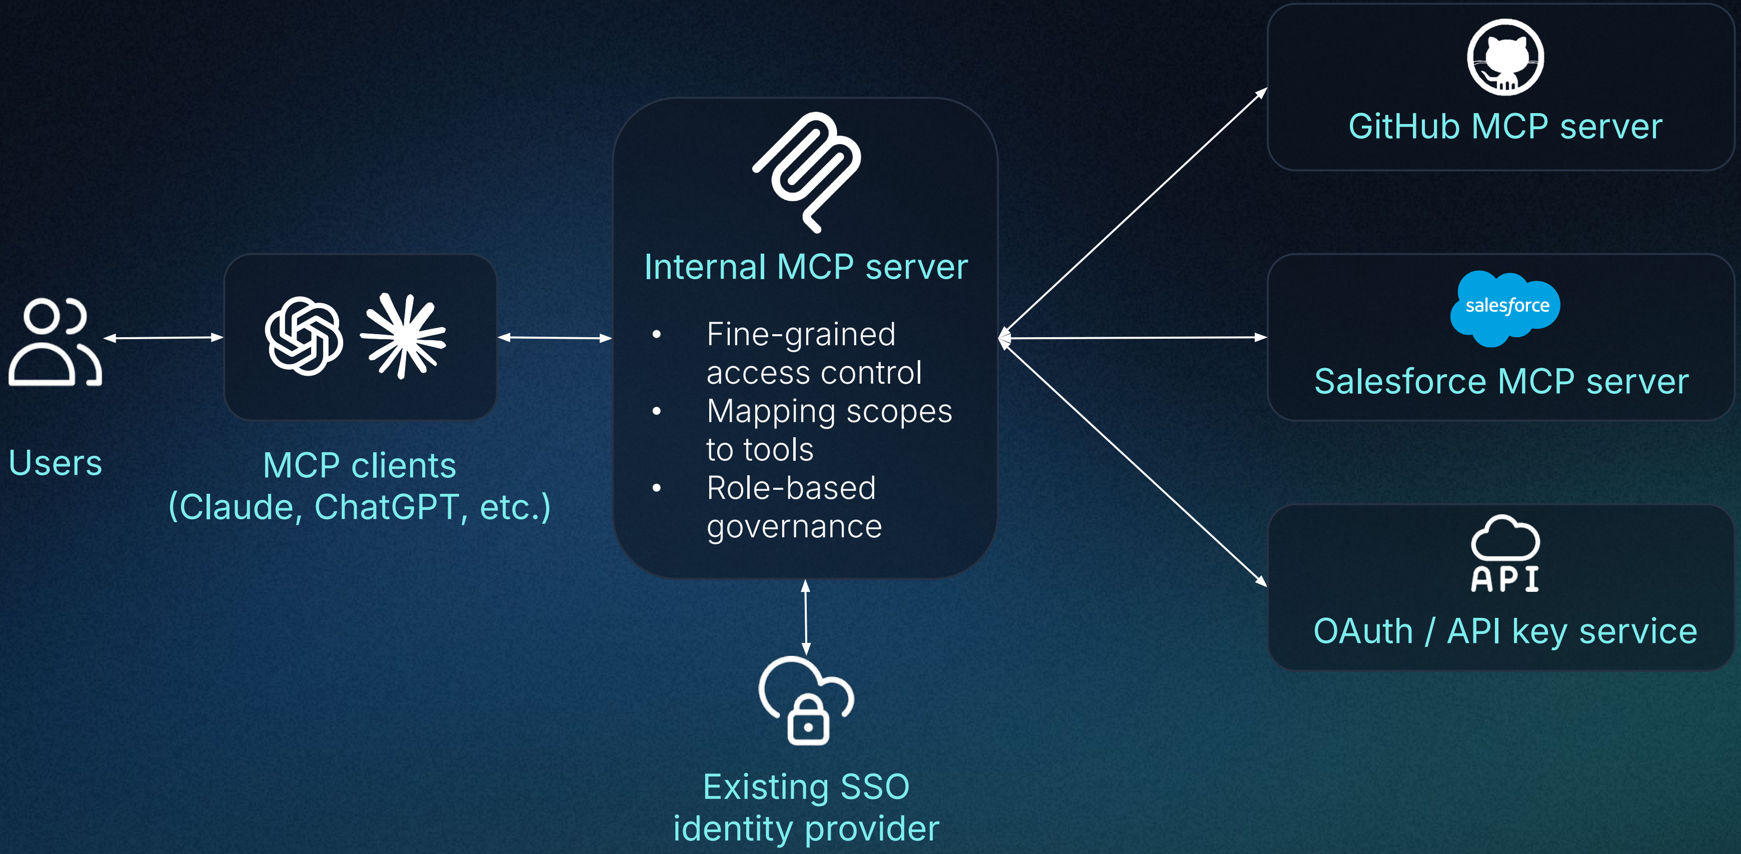

Users

MCP clients (Claude, ChatGPT, etc.)

Internal MCP server

- Fine-grained access control
- Mapping scopes to tools
- Role-based governance

Existing SSO identity provider

GitHub MCP server

Salesforce MCP server

OAuth / API key service

MCP Use Case 2: Customer-Facing MCP Servers

descope

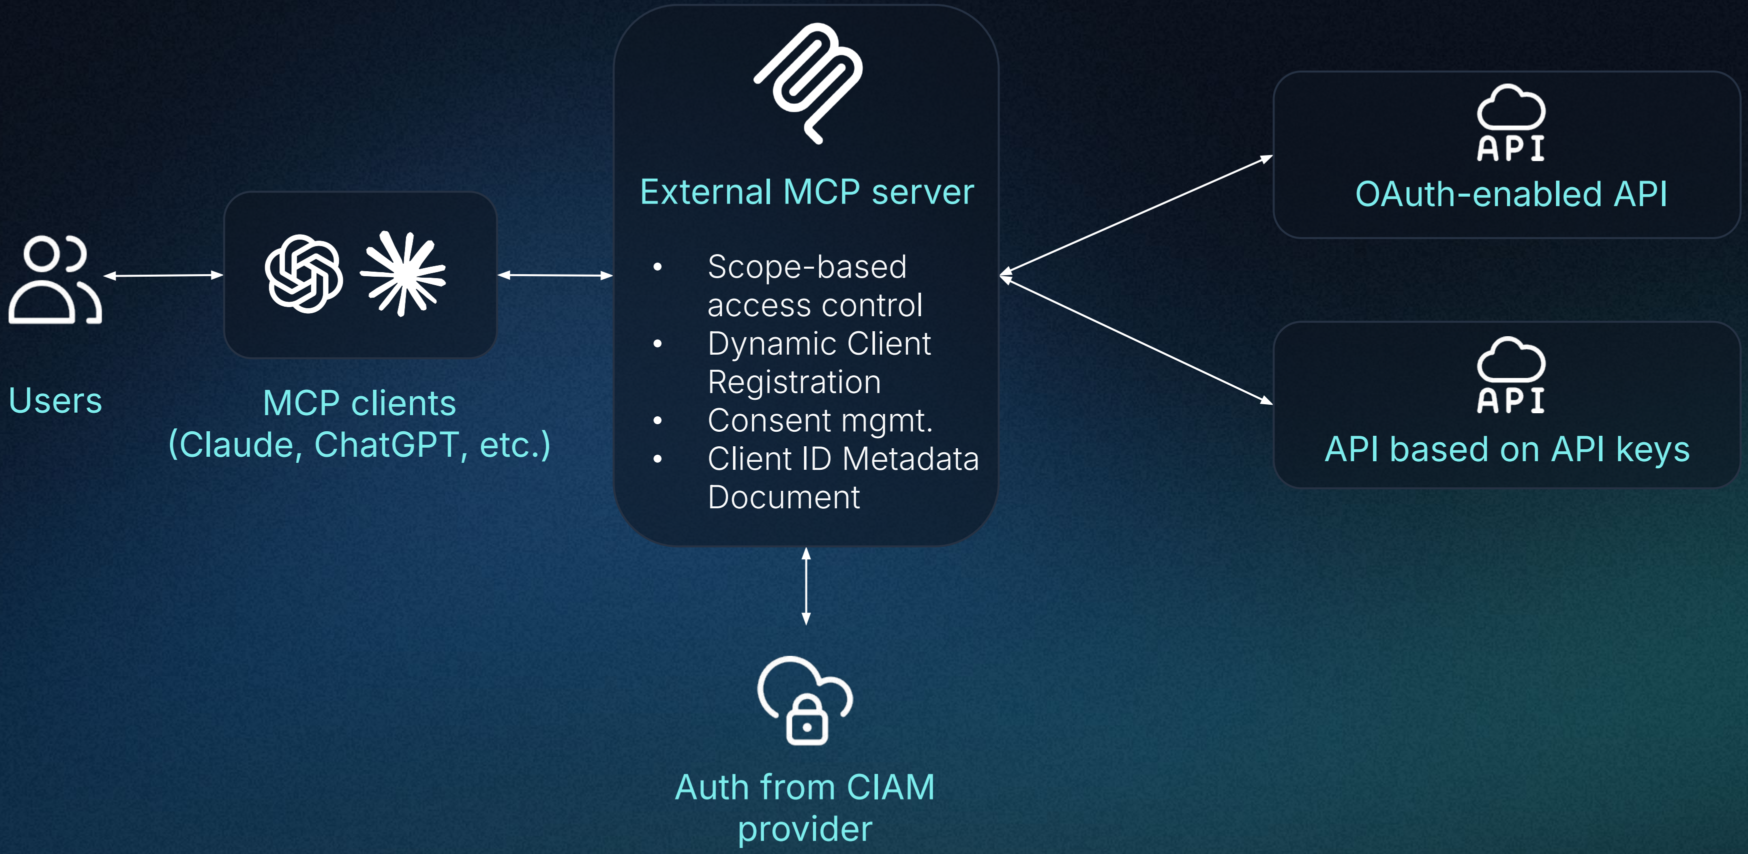

Users

MCP clients
(Claude, ChatGPT, etc.)

External MCP server

- Scope-based access control
- Dynamic Client Registration
- Consent mgmt.
- Client ID Metadata Document

OAuth-enabled API

API based on API keys

Auth from CIAM provider

Emerging OAuth Enhancements

Client Credentials Flow (RFC 6749)

A user goes to the UI, generates client ID / secret, manually configures these in MCP client

Dynamic Client Registration (RFC 7591)

Third-party clients automatically register with authorization servers without manual pre-reg

Client ID Metadata Document (SEP-991)

Client hosts a ‘metadata document’ on an HTTPS URL to identify itself to AuthZ server

MCP Identity Challenges

Developer Challenge: The Agentic Auth Rabbit Hole

descope

Developers are hired to build AI systems...

- API authentication
- OAuth 2.1
- PKCE
- DCR / CIMD

- Token mgmt. / storage
- Fine-grained AuthZ
- Implementing MCP spec
- User consent

...but get stuck building authentication and access control

88% of organizations use or are planning to use AI agents, but only 37% have moved past pilots.
Identity is the roadblock between playground and production.

(source: Descope)

Identity Tips for MCP / AI Agents

MCP Authorization

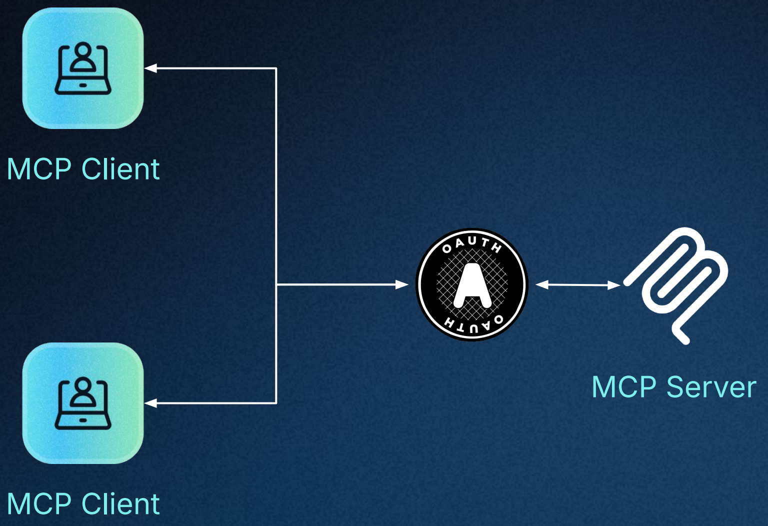

- Implement OAuth 2.1 with Proof Key for Code Exchange (PKCE) as recommended by the MCP specification
- Separate authorization servers from resource servers (your MCP server should only be a resource server)
- Implement CIMD or DCR with hardening controls
- Protect MCP servers with user auth and SSO (mapping groups / roles with existing IdPs)

Scope-Based Access Control

- Implement function or task level scopes for better AI agent experience (may not map 1:1 with APIs)

- Provide only the minimum required scopes to AI agents to complete their current task

- Have mechanisms for AI agents to progressively request more scopes for new tasks

- Restrict AI agent access to specific scopes / tools based on their trust level (verified, DCR, public) and the bound user (role, JWT claims)

Token Management and Storage

descope

- Institute OAuth or API-based connections between your MCP server / AI agent and third-party tools
- Never pass through tokens received from MCP client to upstream tools
- Provide support for storing tokens at both a user and tenant level

User Consent Management

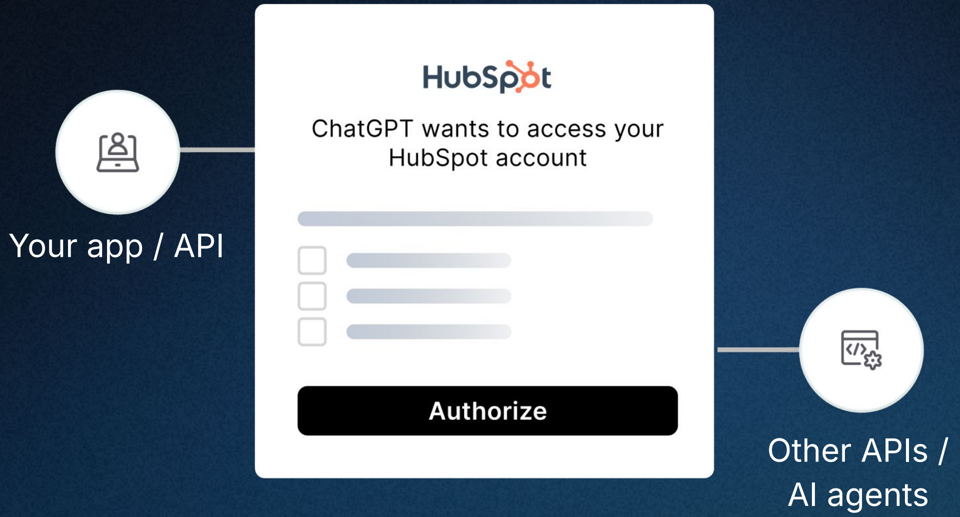

ChatGPT wants to access your HubSpot account

- Implement consent management flows to provide users visibility into what actions & data are being requested on their behalf
- Display mandatory and optional scopes to users to give them control over which scopes to authorize
- Initiate step-up consent flows when AI agents need elevated scopes to perform specific actions

AI Agent Identity Lifecycle Management

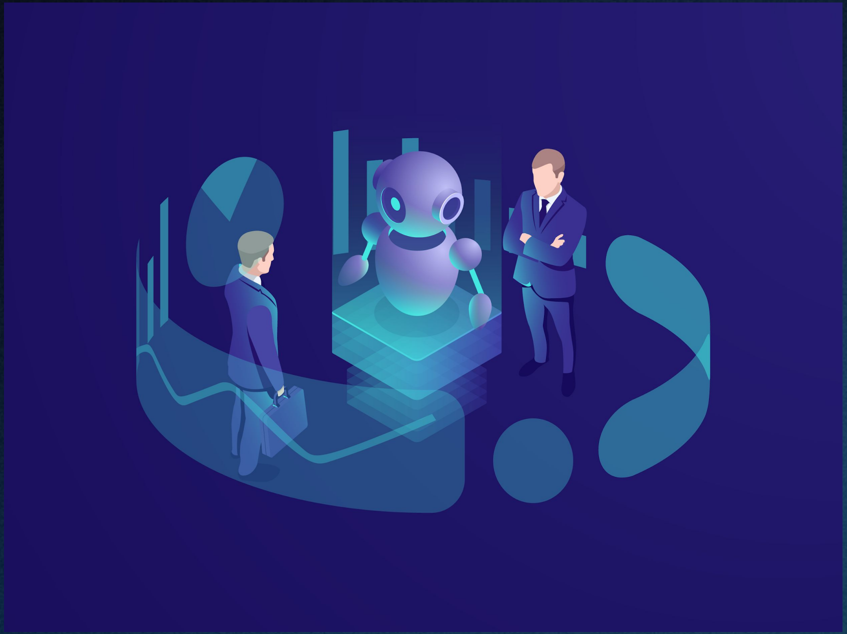

Provisioning

- Provide static and dynamic (DCR) AI agent identity provisioning
- Assess risk of DCR requests (e.g. IP address)
- Have different security controls for different types of AI agents (public, confidential, etc.)

Visibility and Auditing

- Institute detailed audit logs for AI agent registrations, type, scopes granted, etc.
- Audit AI agent identities alongside their bound user identities

descope

Thank You!

Learn more at

www.descope.com/use-cases/ai