MCP SECURITY PROXY

Control what AI agents do in your third-party apps

The Obsidian MCP Security Proxy sits between AI clients and vendor MCP servers. It maps each session to the person and app tenant behind it, and allows or blocks each tool call before it reaches the app.

Control agent actions at runtime

Allow or block each tool call before it reaches the app.

Know who is acting and where

Tie each session to the person and the app tenant it reaches.

Keep app tokens out of agents

The AI client never holds the app's own access token.

THE PROBLEM

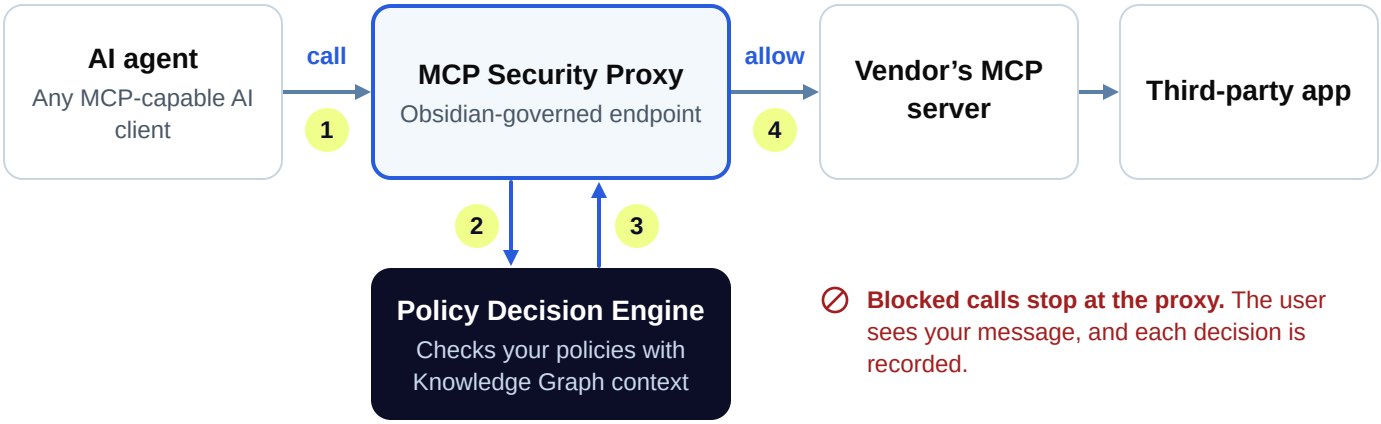
Each MCP connection gives an AI agent a direct path to your app data, outside your security team's controls

Employees can connect AI agents to apps like Salesforce, Snowflake, and Slack through MCP in minutes. On a direct connection, an agent can send customer records out through an MCP tool no one reviewed.

- **Any tool, any call.** Agents can call any tool the vendor's MCP server offers, including deletes and exports.
- **Tokens in the client.** The app's access token sits in the AI client, outside your team's controls.
- **No person or tenant.** A request looks like any valid session, with no sign of who is behind it.
- **Permission, not policy.** The app checks that the account can act, not whether the action fits your policy.

HOW IT WORKS

The proxy checks each tool call before it reaches the app



1. The agent calls a tool through the governed endpoint that replaced the vendor's MCP URL.

2. The proxy maps the session to the person and tenant, and checks for matching policies.

3. The engine checks your policies with Knowledge Graph context.

4. Allowed calls reach the app. Blocked calls stop with your message. Each request is recorded.

CAPABILITIES

How the MCP Security Proxy controls agent activity in your apps

Runtime control

Allow or block each tool call before it's forwarded to the vendor's MCP server.

User and tenant mapping

Route each session to the right org, workspace, or instance, and see who started it.

Token isolation

Agents get a proxy-issued token, and no shared service account carries everyone's access.

Activity record

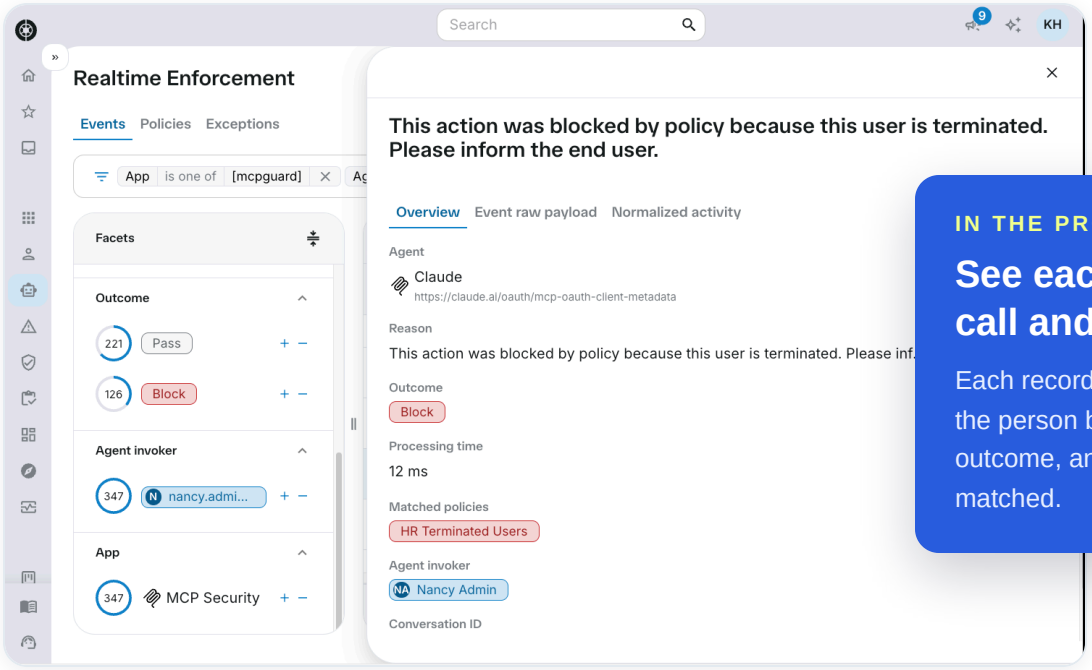
See each session and tool call, the person behind it, and whether it was allowed.

IN ACTION

A departed employee's agent is still connected. The proxy blocks its next tool call.

THE POLICY **Block agent tool calls for anyone HR marks as terminated**

1	Nancy connects Claude to company apps through governed endpoints.	Governed endpoint
2	Nancy leaves the company, and HR records her termination.	HR status change
3	Her agent makes another tool call, and the engine sees her HR status.	Terminated, from HR HR Terminated Users
4	The call never reaches the app, and the agent shows your policy message.	Blocked



IN THE PRODUCT

See each blocked call and why

Each record shows the agent, the person behind it, the outcome, and the policy that matched.

See the MCP Security Proxy in action

Talk to us about securing agent risk.

Request a Demo