

Say Yes To AI Agents Without Handing Them Everything.

Your teams want AI agents connected to the apps they use every day, and each agent arrives with its user's full access.

Lumos lets you safely approve those requests by assigning owners to NHI and agents, fixing risky access, and most importantly: checking every agent action before it runs.

95%
of non-human
identities have no
assigned owner

The risk is in what agents do, not just what they can reach

MCP servers, the connectors that give agents their tools, usually grant more access than any one task needs. Agents act at machine speed, and most of what they do never shows up in an access log. The non-human identities underneath, such as service accounts and API keys, are no better: 95% have no assigned owner.

What you can do **with Lumos**

MCP GOVERNANCE

Say yes to agent integrations, safely.

A team wants their agent connected to a sensitive app, but the access behind it is broader than you'd give a new hire. MCP Governance gives you control to allow the tools they need while blocking risky calls, like merging code or changing a database.

NHI INVENTORY

Put an owner on every key your agents use.

A service account with admin rights hasn't been used in months, and nobody remembers creating it. NHI Inventory shows it with its activity, its access, and the owner you assign.

IDENTITY AGENT FORCE

Fix risky access before someone uses it.

Over-privileged and dormant identities pile up faster than anyone can review them. Identity Agent Force finds them, shows the evidence, and fixes them once you approve.

Saying "yes" to agents requires control over what they do and who owns them. **Lumos gives you both in one platform.**

What makes it **work**

Covers more than MCP.

Bash commands, file edits, and browser use are captured alongside MCP calls.

Policies in plain language.

Albus, Lumos's AI agent, drafts rules from the servers and tools your teams actually use.

A hook, not a gateway.

Each call is checked before it runs, with a target of under 50ms and no traffic rerouting.

Findings with a fix.


Identity Agent Force flags risky identities with evidence and a recommended fix, and remediates once you approve.

One identity record behind it all

Your people. Lumos already knows what each person can access.

Their agents. MCP Governance checks every tool call against that access before it runs.

The identities agents run on. NHI Inventory gives every service account, key, and token an owner, and Identity Agent Force finds the risky ones and fixes them once you approve.



What you get

- MCP Governance for Claude Code and Codex, deployed centrally with no setup for users
- NHI Inventory for Microsoft Entra ID, AWS, and GitHub
- Identity Agent Force findings across the 300+ apps and directories Lumos connects to
- Full tool-call history from the first day

Find out what your agents can reach.

Book a demo: lumos.com/demo

Questions to ask any agent governance vendor

1. Does it see tool calls that don't go through an MCP server, like bash commands and file edits?
2. Does it require routing agent traffic through a gateway?
3. Are decisions based on each user's real access, or on a separate rulebook?
4. Can it show who owns the keys your agents use, and fix the risky ones?