



Deploy AI With Control Over Human, Machine, And Agent Access.

Every AI agent you deploy acts with access someone already granted: an employee's account, a service account key, an API token.

Lumos manages the access people, machines, and agents hold, controls what agents do with it, and gives your identity team its own agents to do the work.

More identities than you can govern, and agents that act on their own

Every agent runs on access your team already struggles to govern. Machine identities outnumber people at least 10 to 1¹, and in customer environments 95% have no owner². Quarterly reviews weren't built for that.

Agents add a second problem: they pick their own path to a goal. In testing, agent models with reduced safeguards have bypassed restrictions and run code on third-party servers. The right access isn't enough. You also have to control the action.

Platform Overview

AGENT ACTIONS

Set rules for the tools agents call, and enforce them before each call runs.

ACCESS REQUEST AGENT

ACCESS REVIEW AGENT

MACHINE IDENTITIES

Find every service account, key, and token, map it to an owner, and fix risky access.

NHI INVENTORY

IDENTITY SECURITY POSTURE

HUMAN ACCESS

Give every employee the right access for the right amount of time, and prove it to auditors.

APPSTORE

LIFECYCLE MANAGEMENT (JML)

ACCESS REVIEWS



80%

less standing access



98%

faster time to access



67%

less privileged access

Manage The Access. Control The Action.

Lumos does both For Every Identity Your AI Rollout Depends On.

Every employee has the right access, only as long as they need it. Access Reviews flag outliers and turn every decision into audit evidence. AppStore grants just-in-time access that expires on its own. Lifecycle Management adjusts access automatically when people join, move, or leave.

Every machine identity has an owner. Lumos inventories every non-human identity (NHI), from service accounts to API keys, with its owner, permissions, and activity. Identity Security Posture then ranks the risks and removes stale or excess access.

Every agent action is checked before it runs. MCP Governance sees every MCP (Model Context Protocol) server and tool call in Claude Code and Codex, and enforces your rules in under 50ms per call.

IDENTITY AGENT FORCE

Agents do the work. Your team sets the rules.

Your team can't hire its way to machine speed. Lumos agents run the routine work continuously, and your people set policy and approve exceptions.

Access Review Agent

Runs access reviews end to end

Access Request Agent

Grants access only as long as it's needed

Role Mining Agent

Learns how teams actually use access

Agent Ownership Finder

Catalogs every agent and NHI in your business

Entitlement Analyst

Translates any permission into plain English

NHI Threat Hunter

Monitors every NHI for risky configuration

MCP Governance Agent

Approves and monitors MCP servers

Agent Marketplace

30+ more agents

What you get

Lumos connects to 300+ apps and directories across HRIS, identity providers, SaaS, cloud, and on-prem, with an SDK for custom apps. Start where the pressure is (an agent rollout, a full IGA program, or one process you want fixed) and expand from there.

See how Lumos governs every human, machine, and agent identity.

[Get a Demo →](#)

Questions to ask any identity vendor

- Can you show every machine identity with a named owner?
- Do you see fine-grained entitlements, or only app-level access?
- Can you stop an agent's tool call before it runs, or only log it afterward?
- Do human and machine identities share one review and one audit trail?