GOOD AGENTS MANAGER

HR for your AI agents.

Good Agents Manager is an appliance on your own network that onboards every AI agent your organization runs, gives each one a role, reviews its work, and **stops dangerous actions before they happen**. Your people stay in charge; your data stays on your premises.

NOT JUST ANOTHER AI GATEWAY

A gateway connects agents. Good Agents manages them.

	TYPICAL AI GATEWAY	GOOD AGENTS MANAGER
Models	Routes requests	**Grants each agent only the models its role allows**
Tools and actions	Not seen	**Every tool call (MCP) checked before it runs: allow, deny, or ask a person**
Dangerous actions	Passed through	**Held for approval, with a briefing that cites the agent's record**
Agent quality	Token counts	**Performance reviews and a trust score per agent**
New agents	Allowed with a key	**Start on probation with nothing granted**
Incidents	Revoke keys, search logs	**Suspend in one step; a tamper-evident record of every action**
Your data	Often in a cloud	**On an appliance you own**

WHAT IT DOES

- **Onboard:** each agent gets a badge and a role: which models, tools, files and sites it may use.
- **Supervise:** model and tool gateways apply the role, budgets and approval rules to every call.
- **Review:** the Agent Steward, a model you choose, reviews a sample of each agent's work; trust grows with a good record.
- **Act:** people approve risky actions, and suspend, reinstate or offboard agents at once.
- **Record:** a hash-chained ledger of every action and decision, shipped to your SIEM; alerts to e-mail, Slack, Teams, ServiceNow.

YOUR AI JOURNEY

Exploring
See every agent and model, what they cost and do.

Adopting
Roles and budgets; people approve the risky few.

Scaling
More independence for agents with a good record.

Governing
Show auditors who did what, under which rules.

HELPS YOU COMPLY

- **AI inventory** of agents, roles, models and owners.
- **Human oversight:** approvals, probation, suspension.
- **Record-keeping** nobody can quietly change.
- **Least privilege** for agents; privacy by default.

Supports NIST AI RMF, ISO/IEC 42001, OWASP LLM Top 10 and Agentic AI, MITRE ATLAS, and SOC 2 and ISO/IEC 27001 controls, and the record-keeping and human oversight that rules such as the EU AI Act ask for.

STANDARDS AND CERTIFICATIONS

FIPS 140-3 validated cryptography (CMVP #5247)	Meets
Memory-safe language	Meets
Vulnerability disclosure policy, safe harbor	Meets
OWASP ASVS 5.0 Level 2	Self-assessed
NIST SSDF (SP 800-218)	Self-assessed
SLSA build provenance	Level 2
CISA SSDF attestation, Secure by Design pledge	In progress
Independent penetration test	Planned
OS FIPS (Ubuntu Pro), SLSA Level 3	Planned
SOC 2 Type II, ISO/IEC 27001, 42001	Planned

Self-assessed: assessed by Good Agents, not yet an independent auditor. Status as of November 2026.

Secure by default: two-factor sign-in for every administrator, no default passwords, HTTPS only, encrypted at rest and bound to the TPM, signed and verified updates.

See it on your own agents: request a demo