See who uses AI. Control its access.

Know which agents act for whom. Control access to models and enterprise systems, review exceptions, and stop new actions from one workspace.

Authorize every agent task

Issue expiring task keys. Limit models, tools and business actions. Keep enterprise credentials in the gateway.

See who is using what

Follow owners, tasks, LLM calls and tool activity. Review decisions, failures, tokens, costs and rate alerts.

Investigate and act

Mask or block sensitive input per app. Approve exceptions, revoke task access and export audit evidence to your SIEM.

From visibility to control

Authorize → Task identity & tool limits

Observe → Task activity & audit trail

Act → Input protection & task controls

↓

Owner → Agent → Task → System

Signed identity | Access decisions | Audit export

Connect MCP, REST/OpenAPI and A2A
Approve exact actions above business limits
Stop subsequent requests when needed

Enable input protection per application. Existing applications and supported coding accounts remain visible.

Testing preview. Controls cover connected traffic. Input protection is opt-in and requires detection services; model responses are not scanned. Rate alerts are rule-based, not ML anomaly detection.

Built with maintained open-source projects. Component credits and licenses are included in distribution notices.

Review your AI access.

agentdatamcp.com