Securing AI Access to Production Infrastructure

Control how AI agents and machine identities access data, systems, and production environments

hoop.dev

AI Productivity. Full Control.

- Access Guardrails
- Data Masking
- Action Level Controls
- Full Session Recordings

Data Protection for Secure AI Development

Prevent Data Exfiltration

Data is redacted in transit so AI agents only receive masked values, even when prompted directly.

Block Unsafe Actions

Block risky write operations using guardrails and fine-grained, action-level approvals.

Log and Record Everything

Every action taken by an agent is logged and recored for compliance export or playback.

End Standing Privileges

Access is granted with one click and revoked automatically to eliminate persistent attack surface.

What Traditional Access Governance Can't Protect

82% of enterprises report AI agents accessing sensitive data on a recurring basis.	**40%** of companies report data privacy incidents tied to AI usage.
72% of generative tools in enterprise environments are rated high or critical risk.	**200+** monthly instances of sensitive data shared with AI tools across large enterprises.

Developer Local Agent Loops

When developers give credentials to AI coding assistants. You can't control which models they use or where data goes.

Hoop's Solution: Redact sensitive data before it reaches the agent so only masked data leaves your perimeter.

Server-Side Agent Deployments

Cloud-based agents run 24/7 with privileged access, running hundreds of loops per developer, with production credentials.

Hoop's Solution: Guardrails on all agent identities that block destructive operations, enforce read-only access, and route privileged actions for approvals.

Services to Models Integrations

AI features send customer data to model providers. Prompt injection enables internet-based data exfiltration.

Hoop's Solution: Block sensitive data before it enters training sets or prompts without changing ETL pipelines or application code.

How it Works: Network Security and Data Visibility for AI

Granular Control

Layer 7 Tools (IAM policies, database roles, RBAC) can see your data but require writing millions of policy lines across unique protocols. Teams grant overly permissive access just to stay productive.

Broad Coverage

Layer 3 Tools (network segmentation, zero-trust gateways) deploy transparently across thousands of resources but lack visibility into what's actually flowing through the system.

New Controls for AI Driven Access. AI multiplies identities, access paths, and execution speed. Security models designed for human access don't scale when every developer runs multiple local agents and hundreds of server side agent loops.

The Hoop.dev Approach: Enforce Controls on Infrastructure Traffic

Hoop.dev runs directly on infrastructure protocols, unwrapping database traffic, SSH sessions, and Kubernetes APIs in real time to identify sensitive data using ML inference and enforce controls before data is accessed, modified, or exported.

No modifying schemas, rewriting applications, or changing developer workflows.

Refreshingly Easy Data Protection & AI Security

Securing AI Access Across the World's Largest Beverage Supply Chain

A Fortune500 Beverage Company deployed hoop.dev in front of 5,000 databases in less than two weeks. The results were powerful and immediate.

98% data coverage with zero configuration.

Sensitive data is inferred at the protocol level with no schema cataloging, allowlists, or per–database policies

80% reduction in IAM ticket volume

Manual policy writing was replaced with automated, in–transit data redaction

ROI in developer productivity.

Access delays were eliminated, context switching dropped, and incident response workflows were streamlined

Protecting Data, Not Just Access, Made AI Safe at Scale

Hoop.dev succeeded where no other solution could because controls were enforced directly on infrastructure traffic, before data reached developers, agents, or external AI tools. As a result, AI tools could interact with real data without exposing sensitive fields or creating new approval bottlenecks.