zentera™

Secure Your AI Agents. Protect Your IP.

Full visibility and control over every AI agent across your enterprise.

The Challenge

Your engineers are encouraged to adopt AI coding agents for business productivity reasons faster than your security team can track them. Claude Code, Cursor, GitHub Copilot, and others may already be running on endpoints everywhere and can access your most sensitive IP - writing code, refactoring services, accessing internal repositories, and interacting with production systems.

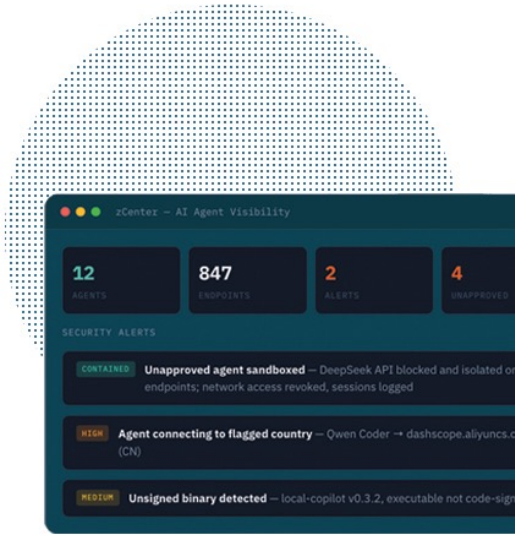

These agents operate autonomously. They choose what to share with LLMs, which tools to invoke, and which data to include in prompts. They can use personal API keys with no central oversight. And today, most CISOs cannot answer a basic question: what AI is running in my environment, who is running it, and is my IP safe?

IAM can't govern shadow AI or assets it isn't aware of. Firewalls block network destinations, not agent behaviors. Endpoint security detects malware, not legitimate AI tools. Nothing in your current stack was designed to directly and effectively address this challenge.

The Solution

Ensage™ AI is the only security platform that is purpose-built for enterprise governance of AI agents with end-to-end monitoring and control. It discovers every agent, enforces what each one is allowed to do, inspects every session inline, and maintains a complete audit trail - all on-premises, all without modifying your agents or network environments.

It extends Zentera's Zero Trust model with a new concept: while Virtual Chambers protect your trusted assets, Ensage AI Sandboxes contain untrusted agents and prevent them from reaching anything they are not authorized to access.

How It Works

zLink Endpoint Sensor. Endpoint sensor that detects agents via behavioral fingerprinting. Creates per-agent network sandboxes. Tunnels traffic for inline inspection. Reads agent configs to surface installed MCP servers and tools.

AI Session Controller (ASC). Inline proxy that terminates TLS to inspect every agent request and response. Enforces policies. Swaps personal API keys for enterprise credentials. Injects compliance prompts. Logs everything.

zCenter Orchestrator. On-premises management plane with policy engine, agent inventory, approval workflows, audit dashboards. Integrates with your IdP and delivers Zentera Zero Trust controls, including Virtual Chambers.

These key components are powered by Zentera Labs, Zentera's dedicated agent security research service. The Zentera Labs team continuously profiles the AI agent ecosystem - fingerprinting agents, tracking changes across releases, and evaluating them to help admins make informed decisions on which agents, plugins, and tools to trust. Zentera Labs allows your agentic security to keep pace with agents that evolve daily.

Key Capabilities

Agent Discovery. Discover and identify every AI agent across your endpoints. Know who is running what, where, and what it connects to.

Configuration Risk Detection. zLink monitors agent configuration, using trust scores powered by Zentera Labs Intelligence to manage access to MCP servers and tools.

Agent Sandboxing. Agents are placed in a network sandbox for access policy enforcement. Unapproved agents or runaway agents are quarantined to keep them isolated.

Policy Enforcement. Enforce policies by user, project, team, or agent type. Restrict LLM endpoints, MCP servers, tools, and token budgets. Integrate with Chamber policies for Zero Trust controls for critical asset protection.

Compliance Guardrails. Prepend system prompts into every agent session, reinforcing your data handling policies.

Project Isolation. Zero Trust controls prevent one project environment from accessing resources belonging to another, enforcing separation across teams, business units, or compliance boundaries.

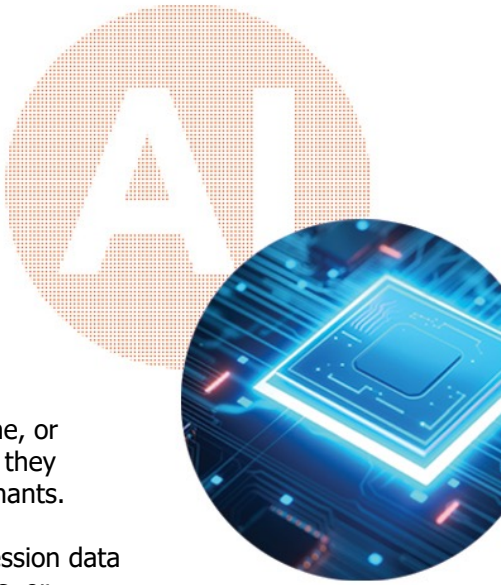

API Key Management. Give engineers proxy keys - real API keys are always shielded away from all endpoints. Every request can be attributed to a user and a project.

Chargeback-Ready. Track token usage by user, project, team, and LLM provider. Set quotas. Give project leads visibility into AI spend.

Multi-Tenant Ready. Separate by business unit, product line, or customer engagement. Project-scoped admins see only what they should. The workloads are isolated and segmented across tenants.

Comprehensive Audit. All management, inspection, and session data kept entirely within your environment. No prompts, responses, or sensitive data are shared with 3rd party clouds.

Extensible. Deploy custom scripts on ASC to extract internal metadata from agent requests and feed policy decisions, supporting your internal tooling workflows.

Zero-Touch Deployment. The entire Ensage AI platform is running overlay, on top of your existing physical network. There is no need to change any network topology or drill a pinhole in your firewall. Therefore, the solution onboarding is seamless and fast.

Getting Started

Ensage AI lets you adopt AI coding agents with confidence, without slowing down your engineering teams.

Whether you're just getting started with agentic AI or well on your way, Zentera can help you secure your AI use with Zero Trust right away.

Contact us today to schedule a technical briefing.

(408) 436-4811 • sales@zentera.net • www.zentera.net/ensage-ai

Scan for more info!