zentera™

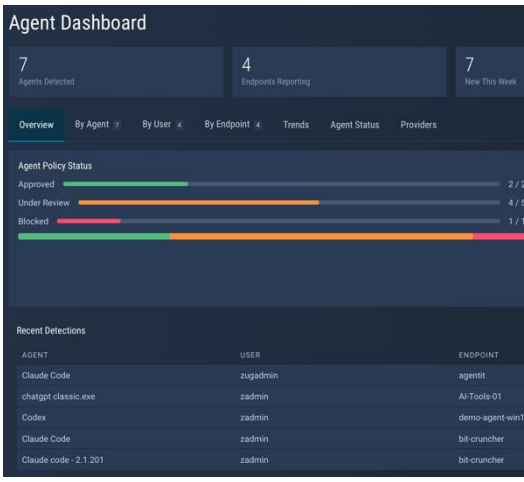
Ensage™ AI

See, control, and audit enterprise AI agents.

Overview

AI agents now write code, run tools, and query internal systems in design, research, and engineering environments. An agent usually runs with the access of the person who started it. It can send any data it can read, including design IP, to a model provider. Some agents run without the security team's knowledge.

Zentera Ensage AI is an AI agent security platform for these environments. It detects AI agents on enrolled endpoints and servers, including agents that your team did not deploy. It places each agent in a sandbox inside its project's enclave and inspects the agent's sessions with AI models. The evidence stays in your environment. Ensage AI deploys as an overlay on your existing network.



The Risks

Shadow agents. Engineers install coding assistants, command-line agents, and local models. Some agents use personal API keys.

Inherited access. An agent can open the same files and reach the same systems as its user, including data from other projects.

Uncontrolled prompts. The agent decides what goes into each prompt. Source code, design data, and customer records can go to any model that the agent can reach.

Exposed keys. API keys sit in configuration files and environment variables. Other processes that run as the same user can read and copy them.

Runaway agents. An agent in a loop can run up large token costs before anyone notices.

Key Benefits

Keep source code and regulated data away from unapproved models. Policy sets which AI providers each agent can reach. The AI Session Controller redacts sensitive data in prompts before the prompts leave your environment.

Keep each agent inside its project. Each project is an enclave. Resources outside the enclave are not reachable from inside it.

See every AI agent, whether you deployed it or not. zLink, the Zentera endpoint sensor, identifies AI agents and coding tools by their processes. zCenter keeps the inventory.

Keep enterprise API keys off the endpoint. Agents hold substitute credentials. The AI Session Controller injects the real key at the enforcement point.

Keep the evidence in your environment. Policy, decisions, and session records stay under your custody. Zentera does not host customer data.

Built For

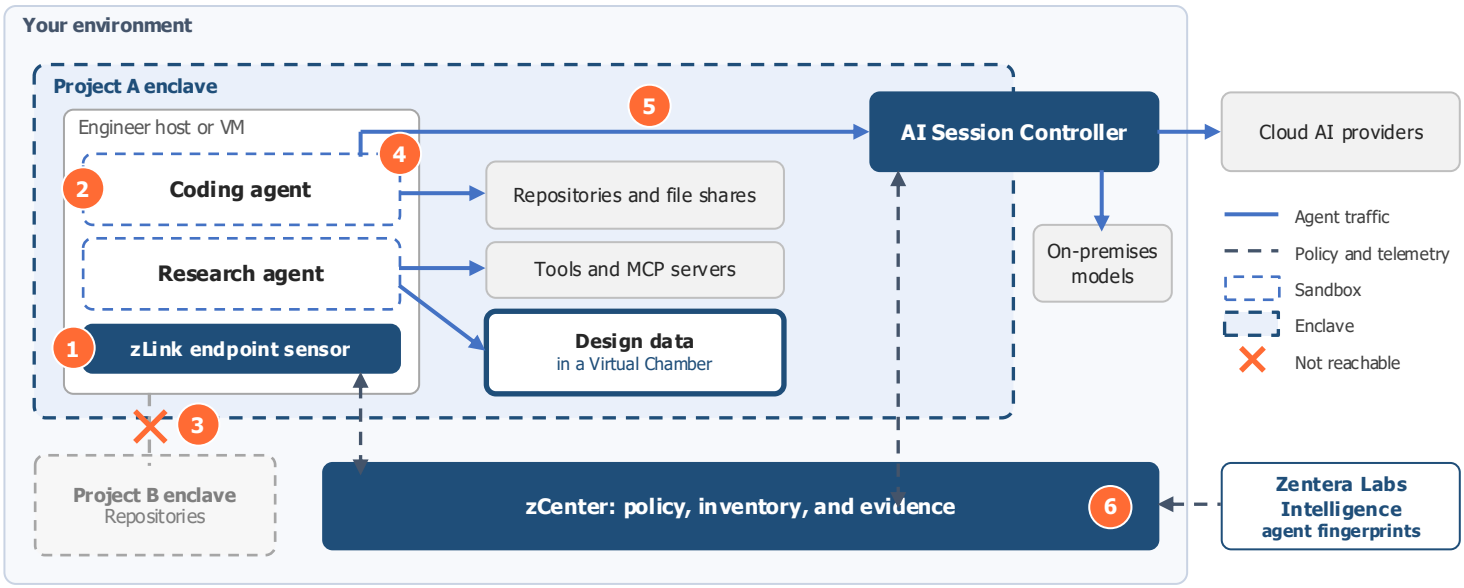
Enterprises where AI agents work with high-value IP or regulated data, including semiconductor and EDA design teams, pharmaceutical research and clinical operations, and financial services.

Copyright ©2026 Zentera Systems, Inc. All rights reserved.

zentera™

How It Works

Ensage AI places three components in your environment. zLink runs on each endpoint or server. The AI Session Controller sits on the path to AI providers. zCenter holds policy and the evidence record. Zentera Labs Intelligence, a Zentera cloud service, supplies current agent fingerprints to zCenter. Steps 1 to 6 follow one coding agent.



1 Detect. A developer starts a coding agent. zLink identifies the agent by its process and reports it to zCenter, whether or not your team deployed it.

2 Assign. Ensage places the agent in a sandbox in the enclave that an administrator assigned. If the agent has no assignment, Ensage can contain it automatically.

3 Reach. The agent reaches its project's repositories, tools, and approved MCP servers. Resources in other projects are not reachable from the enclave.

4 Open files. Data Access Control checks each file that the agent or its tools try to open or execute. The check runs on the agent's machine, before the agent reads any content.

5 Call a model. Calls to AI providers go through the AI Session Controller. It inspects the prompt and the response and redacts sensitive data. It swaps the substitute key for the real key.

6 Record. zCenter records the agent, its configuration, the policy version, each decision, and the session records in your environment.

Three Boundaries

Three boundaries control what each agent can reach. The enclave and the Virtual Chamber come from the CoIP Platform. The sandbox is new in Ensage AI.

Enclave. The boundary around a project. An enclave holds the agents, tools, data, and people for one project. Resources outside an enclave are not reachable from inside it, and enclave resources are not reachable from outside it.

Sandbox. The boundary around each agent. The sandbox limits the agent to the destinations, files, and AI providers that policy allows. Ensage can quarantine one agent without isolating the machine that it runs on.

Virtual Chamber. The boundary around a high-value asset inside the enclave, such as a design database or a research data store. A Virtual Chamber admits only identified users and software, over declared methods.

Copyright ©2026 Zentera Systems, Inc. All rights reserved.

Capabilities

Ensage AI capabilities follow the life of an agent: Discover, Authorize, Contain, Observe, and Maintain.

Discover

Agent detection. zLink detects AI agents and coding tools at the process level on enrolled endpoints and servers. Detection does not need cooperation from the agent. Zentera Labs Intelligence keeps the fingerprints current as agents change from release to release.

Local model discovery. zLink discovers local inference runtimes on enrolled endpoints and identifies model artifacts by hash. The inventory records the exact artifact in service rather than a filename.

Configuration inventory. zLink reads agent configurations to list the MCP servers, tools, and skills that each agent can use. Administrators use trust scores from Zentera Labs Intelligence to decide which ones to allow.

Authorize

Project assignment. Administrators assign each agent to the enclave of its project. The assignment sets what the agent can reach.

Agent policy. Policies apply by user, project, team, or agent type. They set the AI providers, MCP servers, and tools that an agent can use, and its token budget.

Configuration governance. Ensage records approved agent configurations, including skills and MCP servers, and reconciles deployed state against them. When a configuration drifts, Ensage alerts and restores the approved configuration.

Credential substitution. Agents hold substitute credentials. The AI Session Controller injects the real key at the enforcement point, so no enterprise AI credential remains on the endpoint to steal. A substitute credential copied off the host does not work against the provider directly.

Contain

Per-agent sandbox. Each agent reaches only the destinations and AI providers that its project and its policy allow. Inside the sandbox, Data Access Control limits the files that the agent can open.

Data Access Control. Policies set which files and folders each agent can open or execute, by filename or path. Tools that run as child processes of the agent inherit the same limits. Ensage enforces each decision on the machine that runs the agent, before the agent reads any content. Policies cover local files and network storage such as NFS.

Virtual Chambers. Virtual Chambers wrap individual high-value assets inside the enclave. They admit only identified users and software, over declared methods.

Automatic containment. When zLink finds an agent with no project assignment, Ensage can contain it automatically.

Circuit breaker. The AI Session Controller pauses an agent automatically when the agent's token spend runs far above its baseline. The pause limits the cost of a runaway agent.

Kill switch. Administrators can quarantine one agent at the boundary. The quarantine persists across process restarts on that endpoint until an administrator releases it. The machine stays on the network.

Fleet-wide block. To stop an agent type, version, or model, administrators block that item across the fleet in one policy action. The block applies to current and future instances.

Copyright ©2026 Zentera Systems, Inc. All rights reserved.

zentera™

Capabilities

Observe

Inline session inspection. The AI Session Controller inspects agent sessions inline as a TLS-terminating proxy inside your environment. It reads prompts, responses, and tool calls.

Redaction and DLP. The AI Session Controller applies content controls to prompts and responses on premises. It detects and redacts keyword matches and personal data. It applies policy to labeled content that it can decrypt, including Microsoft sensitivity labels.

Compliance prompts. The AI Session Controller adds your data-handling instructions to each agent session.

Token usage. Ensage reports token usage by user, project, team, and AI provider. Administrators set quotas, and project leads see their own AI spend.

Maintain

Evidence record. zCenter holds policy and the evidence record: agent identity, artifact hash, configuration and policy versions, decisions, and session records. The record stays under your custody.

Scoped administration. You can scope an administrator to a business unit, a project, or a customer engagement. A scoped administrator sees only that scope.

Agent lifecycle. zCenter tracks each agent from discovery through approval to retirement. Removing an agent from its enclave removes its access to the project.

Extensibility. Custom scripts on the AI Session Controller pass internal metadata from agent requests into policy decisions.

At a Glance

Components	zLink endpoint sensor; AI Session Controller; zCenter, which integrates with your identity provider; Zentera Labs Intelligence
Deployment	Overlay on your existing network, with no change to network topology and no new inbound firewall openings. zLink, the AI Session Controller, and zCenter run inside your environment, on premises or in your own cloud account.
Data custody	Policy, decisions, and session records stay in your environment. Zentera does not host customer data. Agents under Ensage policy send prompts only to the AI providers that you approve.
Coverage	AI agents, coding tools, and local models on endpoints and servers that run zLink, and agent sessions through the AI Session Controller. Agents that run entirely in a SaaS vendor's cloud are out of scope.

Getting Started

Start with one project.

Zentera can run a scoped validation in your environment with one project team, its agents, and its AI providers. During the validation, your team sees which agents run and what they reach before you set policy.

Contact us to schedule a technical briefing.

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

Scan for more information

Copyright ©2026 Zentera Systems, Inc. All rights reserved.