zentera™

Governing AI Agents at the Network Layer

A Framework and Platform Reference for Enterprise Agentic Security

Executive Summary

AI agents have moved from proof-of-concept to production infrastructure in a matter of months. They are non-human identities operating autonomously, often with the same access privileges as the engineers who deployed them. Yet the governance frameworks enterprises rely on - identity and access management, privileged access management, perimeter security - were designed for human users, not for software that takes actions on its own. Ensage AI addresses this gap with a structured, five-stage governance model and a three-component platform that enforces policy at the network layer. This whitepaper details both: the conceptual framework for maturing an organization's agentic security posture and the specific platform components that make enforcement possible without requiring infrastructure redesign.

The Framework

Agentic Security Governance in Five Stages

Ensage AI treats AI agents as non-human identities operating with human-level access. Its governance model is structured as a progressive, five-stage maturity path. Organizations can begin with pure visibility and scale to full network-layer enforcement at their own pace. Each stage builds on the last, and no stage requires ripping out existing infrastructure.

1 Discovery

Find Every Agent Before You Write a Single Policy

The foundation of any governance program is an accurate inventory. Ensage AI's zLink Endpoint Sensor passively fingerprints every AI agent running in the environment through process monitoring and network behavior analysis. The sensor operates at the OS level and requires no cooperation from the agents themselves. Each detected agent is cross-referenced against Zentera Labs Intelligence, a continuously updated catalog of more than 2,500 known MCP servers, VS Code extensions, and agentic tools, so security teams know exactly what is running - including shadow AI - before any policy is applied.

2 Authorize

Tie Identity to Access Before Anything Runs

Once agents are discovered, each one requires an explicit authorization decision. In the Authorize stage, agents are assigned to network enclaves with defined access boundaries. Identity is bound to access: nothing runs without a deliberate decision from the security or platform team. This mirrors the zero-trust principle of "never trust, always verify" but applies it specifically to non-human software identities operating autonomously within the enterprise.

3 Contain

Enforce the Declared Intent

The Contain stage turns what Authorize declared as policy into wire behavior, enforcing the enclave's reachability map through sandboxes and network controls that agents cannot opt out of, inspecting agentic traffic with an AI Session Controller that agents cannot bypass, and enforcing enterprise policies in ways agents cannot route around. When zLink encounters an agent that was never authorized (a shadow agent, a tampered binary, an unknown process), it enforces corporate intent guardrails at the OS and network layers without agent participation. Containment is a property of the architecture, not a reaction to behavior.

4 Observe

Build the Audit Record in Real Time

Authorization alone is insufficient; governance requires a continuous record of what agents actually do. The Observe stage captures every session, every prompt, and every tool call in real time through the AI Session Controller, a transparent policy proxy that logs interactions without requiring code changes in the agents themselves. This audit trail is the foundation for compliance reporting, forensic investigation, and ongoing behavioral analysis.

5 Maintain

Govern the Full Agent Lifecycle

Governance does not end at deployment. The Maintain stage provides continuous lifecycle management: tracking agent inventory as new agents are deployed and old ones decommissioned, rotating credentials on a defined schedule, and ensuring clean offboarding when projects end. This stage closes the governance loop and prevents "ghost agents" - tools that were once authorized but have since been abandoned without proper decommissioning - from accumulating as a long-term risk.

Taken together, the five stages describe a complete governance arc: from not knowing what is running, to knowing, authorizing, watching, controlling, and continuously maintaining a clean and auditable agent estate. Organizations do not need to implement all five stages at once. Starting with Discovery alone delivers immediate value by eliminating the blind spot that most enterprises currently have around AI agent activity.

The Platform

Three Components, One Control Plane

Ensage AI's five-stage governance model is operationalized through three purpose-built platform components. Each component addresses a distinct layer of the agentic security challenge: endpoint detection, session-level policy enforcement, and continuous intelligence. All three are managed through Zentera's existing zCenter Orchestrator, keeping AI agent governance in the same control plane as the broader Zero Trust policy rather than introducing a separate console.

Importantly, none of the three components requires infrastructure redesign. Ensage AI runs on Zentera's CoIP Zero Trust overlay, the same architecture already deployed in production by enterprise customers for network segmentation and ZTNA. Organizations that already use CoIP activate agentic governance on existing infrastructure.

Component	Key Capabilities
zLink Endpoint Sensor Detection & Sandboxing	• Deployed at the OS level - no agent cooperation required • Detects agents by process fingerprint and network behavior • Sandboxes unauthorized agents by cutting network access • Routes agent traffic to the AI Session Controller • Detects unauthorized modifications to agent binaries
AI Session Controller Policy Proxy	• Operates as a transparent proxy between agents and the resources they reach • Logs every session, prompt, and tool call for full auditability • Injects enterprise system prompts to enforce behavioral guardrails • Enforces per-user, per-project, and per-tool permissions • Provides auditable LLM and tool usage records for compliance
Zentera Labs Intelligence Agent Intelligence	• Continuously updated catalog of 2,495+ MCP servers and 7,859+ VS Code extensions • Process signatures and network behavior profiles for each known agent • Risk scoring based on known behavioral data and external threat intelligence • Continuously synchronized - not a static snapshot • Powers the Discovery stage by cross-referencing detected agents against known profiles

Enforcement Without Agent Cooperation

A critical design principle across all three components is that enforcement does not require the agent's participation. Third-party agents, shadow AI tools, and custom-built agents all operate under the same policy regime, because enforcement happens at the network layer and the OS level rather than through agent-side SDKs or instrumentation. An agent cannot opt out of a sandbox or bypass the policy proxy by ignoring an API call; the infrastructure itself enforces the boundary.

This is the core advantage of building agentic governance on a Zero Trust overlay rather than on agent-side tooling. The overlay is the enforcement point, and it operates independently of whatever the agent is or is not doing.

A Single Control Plane

Ensage AI's three components are managed through the same zCenter Orchestrator that governs the rest of Zentera's Zero Trust policy. Security teams see AI agent activity - active agents, alerts, unapproved tools, behavioral anomalies - in the same dashboard alongside their existing network segmentation and ZTNA policy. There is no separate console to learn, no additional data silo to correlate, and no new infrastructure to deploy for organizations already running on the CoIP overlay.

About Ensage AI

Ensage AI is Zentera's agentic security product, built on the CoIP Zero Trust overlay. It provides network-layer discovery, authorization, observation, containment, and lifecycle maintenance for AI agents operating in enterprise environments - without requiring infrastructure redesign or agent cooperation.

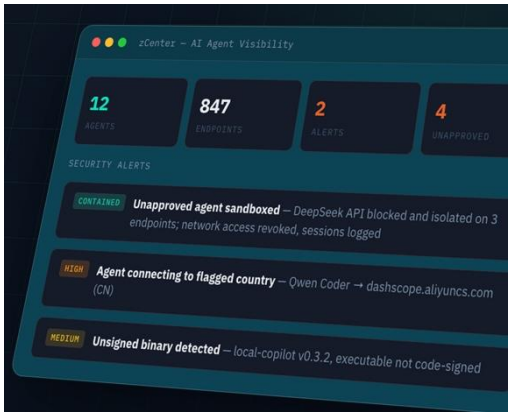

Getting Started

Whether you're just getting started with agentic AI for design 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