xage SECURITY

Xage Zero Trust for Agentic AI

Secure autonomous AI with deep visibility and deterministic control. Block rogue AI behavior, data leakage, and privilege escalation across cloud, SaaS, data center, OT, and edge.

AI agents adapt their behavior based on context, goals, tools, and data. As enterprises connect agents to APIs, SaaS apps, databases, cloud services, and operational systems, agents can access data, automate workflows, and take real-world actions. But most organizations still lack visibility and control over what agents actually do.

Most AI security solutions focus on prompts and outputs, not the actions agents take across enterprise systems. Without runtime visibility and control, agents can expose sensitive data, invoke unauthorized APIs, misuse delegated privileges, or escalate access across workflows. The result: organizations either expose critical systems to unacceptable risk or limit AI to sandboxed pilots.

Agentic AI Challenges

- **Limited Visibility:** Security teams lack insight into the specific actions agents perform. Visibility is often restricted to just the initial prompt and final output, hiding the actual steps the agent executed.
- **Fragmented Enforcement:** Without AI governance, agents inherit permissions directly from user or service accounts. This lack of privilege boundaries results in inconsistent security controls and limited oversight.
- **Over-privileged Access:** Agents are frequently granted expansive access to tools and data rather than scoped, time-bound permissions, significantly increasing the potential blast radius if a compromise occurs or an agent goes rogue.
- **Lack of Unique Identity:** Because agents often operate under shared service accounts or API keys, it becomes nearly impossible to accurately track ownership, audit actions, or revoke access, or identify which agent or user was responsible for which action.
- **Unpredictable Behavior:** Since agents reason and act dynamically at runtime, their actions cannot be fully predicted and bounded absent additional independent monitoring and controls.

xage SECURITY

Zero Trust for Agentic AI Architecture

The Xage Zero Trust for AI platform secures agents, LLMs, data, tools, APIs, SaaS applications and other resources:

- The **Xage Resource Gateway** authenticates and authorizes access to critical resources while logging AI-driven actions.
- The **Xage Agent Sentry** protects each AI agent by controlling communications, limiting permitted actions, and maintaining detailed activity records.

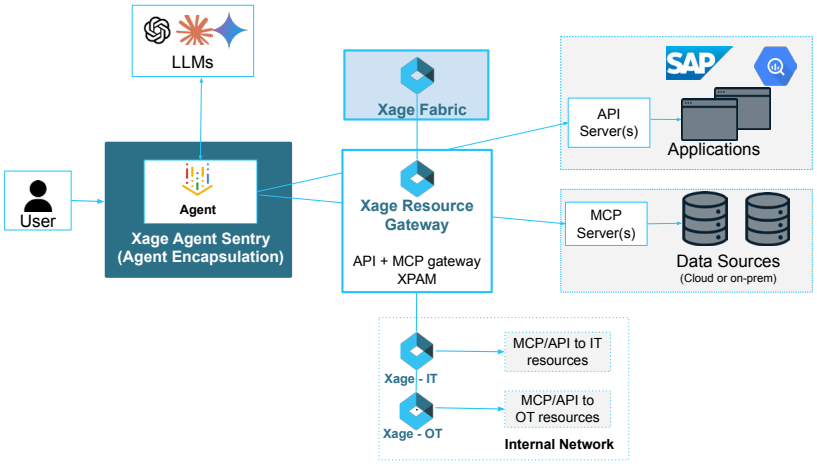

Feature	Description	Benefit	Business Outcome
End-to-End Visibility	Captures every prompt, output, AI agent action, tool call, API interaction, and resource access.	Stronger Governance	Creates accountable, auditable records of agent actions and activity.
Resource Protection	Brokers resource access with jailbreak-proof & deterministic policy enforcement, plus credential protection.	Improved Security Posture	Enables AI use in high-stakes environments where jailbreaks would be unacceptable. Reduces attack surface through least-privilege access. Blocks unauthorized data exfiltration.
Agent Protection	Inspects agent inputs, outputs and actions to prevent misuse, prompt injection, and unauthorized behavior.	Reduced AI Risk	Blocks harmful agent behavior before it affects the business.
Agent Identity Management	Provides centralized identity, registration, credential rotation, and lifecycle control for AI agents.	Simplified Credential Control	Eliminates hardcoded secrets and streamlines credential rotation.
Resource Management	Secures MCP servers, APIs, enterprise resources, SaaS applications and privileged credentials.	Consistent Protection	Applies security controls across hybrid, cloud, and edge environments.
Agent Entitlements	Enforces granular, identity-based permissions for each agent.	Stronger Access Control	Ensures agents can only access the resources and actions they are authorized to use.
Unified Console	Centralizes policy, visibility, governance, agent management, and resource onboarding.	Lower Complexity	Provides “single pane of glass” view of AI activity and management.
Distributed Security Fabric	Extends security controls across cloud, on-premises, multi-cloud, and edge environments.	Greater Trust in AI	Enables safer deployment of agentic AI with tamperproof controls and transparency.

xage SECURITY