xage SECURITY

Unified Zero Trust for LLMs, AI Agents and Applications

As AI reshapes business applications and operations, Xage Security enforces identity-based access to LLMs, agents, data, and infrastructure—with jailbreak-proof protection against data leakage, rogue AI risks, over-permissioned access, and insider threats across complex, hybrid environments.

AI is transforming business—but breaking traditional security. Modern ecosystems of LLMs, agents, APIs, and data pipelines span on-prem, cloud, and edge, interacting dynamically. This creates massive risk: over-permissioned access, data leakage, API abuse, agents going rogue, and insider threats. Enterprises struggle to enable the best use of AI while protecting sensitive and proprietary data across distributed environments and managing overlapping access, compliance pressures, and growing operational complexity.

Unlock AI Innovation—Without Compromising Security or Compliance

Unleash AI With Confidence

Unlock organization-wide AI with confidence—secure, governed, jailbreak-proof and responsible, while avoiding burdensome data classification and tagging.

Simplify Compliance and Operations

Centralize compliance, security, UX, and admin with a single scalable platform spanning the entire infrastructure and AI ecosystem.

Harden the AI Attack Surface

Use rigorous network-level least-privilege access control to block data leakage and stop a rogue AI agent from being able to inflict damage.

xage SECURITY

Xage Secures AI from Edge to Agent

Xage Fabric delivers unified, identity-first Zero Trust protection across the entire AI stack: from infrastructure and data to LLMs and agents. The Xage Fabric overlays dynamic policy enforcement, microsegmentation, and identity-based access control across fragmented environments—without requiring data tagging, reclassification, or infrastructure changes. Even better, and in contrast to so-called "LLM Firewall" or prompt/output/retrieval guardrails, the Fabric delivers jailbreak-proof and granular access control at the network protocol level to block data leakage and rogue AI risks outright.

Whether in large enterprise deployments, sovereign datacenters, or multi-party AI workflows, Xage enforces consistent security and eliminates vulnerabilities.

Architectural Overview

Xage's Zero Trust distributed architecture provides a foundational layer of security for AI-driven workflows by acting as a secure controller between users, AI agents, LLMs, and data sources. As shown in the diagram below, by proxying and inspecting all interactions—whether initiated by a human user or an autonomous agent—Xage enforces fine-grained, identity-based access controls to ensure that only authorized users and agents can access specific data or systems. This prevents unauthorized data exposure and lateral movement within the environment.

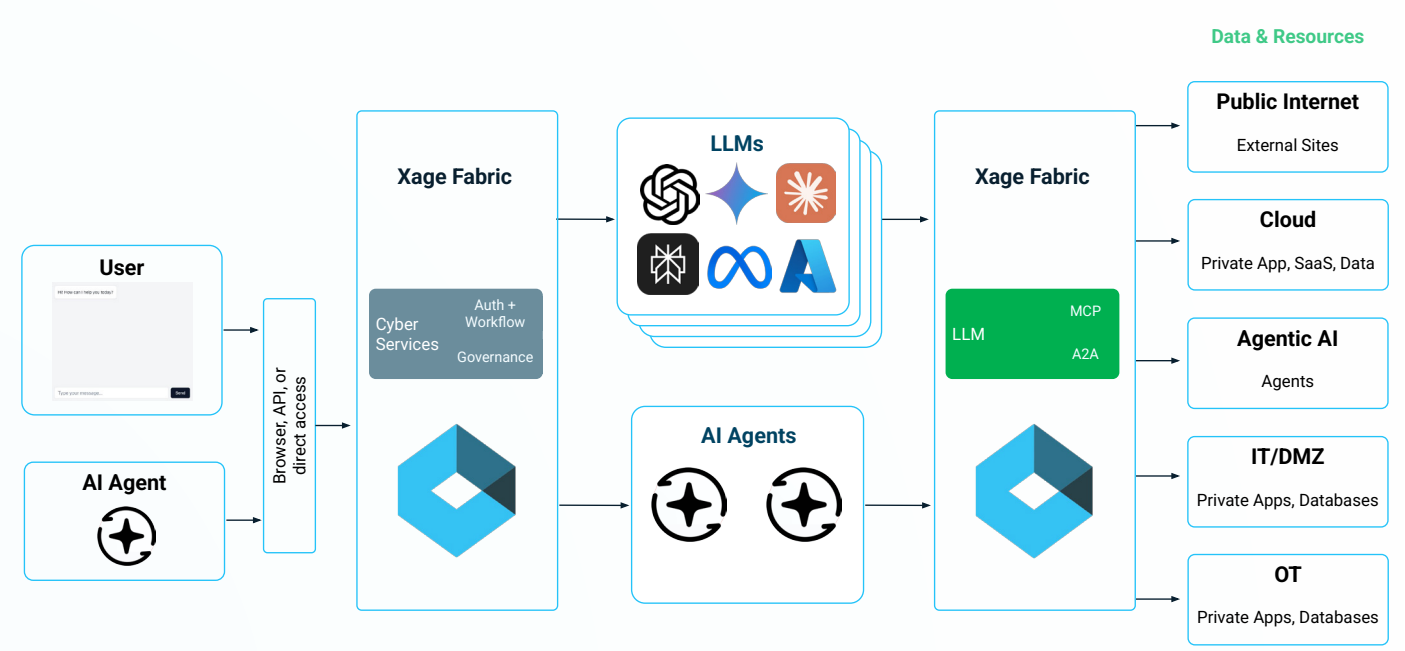

Xage applies real-time guaranteed access control by intercepting traffic at the network protocol level, enabling policy enforcement, data redaction, and prevention of data leakage. Identity and entitlements are traced across the full chain of AI interactions—for instance, human to agent to agent to LLM to data source—to make sure that no AI actor, whether human or machine, is able to access data or initiate an action for which they're not authorized. The result is a scalable, secure architecture that brings visibility, control, and governance to AI interactions and ensures compliant use of AI across the organization.

Key Features and Capabilities

Secure Fine-Grained Access to Private Data Sources: Organizations can confidently integrate LLMs with sensitive and proprietary data—whether on-prem, in private cloud environments, or within specialized tools. Xage Fabric enforces least privilege access, so that, not only is the data accessed by the LLM controlled, but when a user (whether human or agentic) attempts to retrieve information from the LLM, both the access rights of the user and the LLM are checked. In particular, just because a user is authorized to access an LLM, and an LLM is authorized to access a particular piece of data, it is not assumed that the user has access rights to that data—the user must themselves have data access authorization in order for the data to be fed back to them by the LLM.

LLM-to-Data Integration Across Layered and Zoned Networks: Xage enables secure LLM access to distributed data sources across segmented and multi-hop environments. For instance, an LLM hosted in the cloud can securely query data deep within an operational network—without exposing the core environment or requiring manual data uploads from each site to the main datacenter for LLM ingestion.

Just-In-Time and Just-Enough Access: Xage eliminates standing privileges for AI tools by enforcing dynamic, context-aware access controls. Actions like modifying a database or deploying code require explicit approval, ensuring timely, scoped access aligned with operational needs. Xage helps maintain operational agility while locking down overreach.

AI Agent and Agent-to-Agent Security: Xage secures autonomous agent-to-agent interactions using identity-based access controls that authenticates every machine and enforces least-privilege communication. Agents only interact with authorized peers, tools, and data sources, enabling secure, scoped automation across distributed AI environments.

Xage provides a software wrapper for each AI agent, ensuring that i) only authenticated and authorized users or apps can send commands to the agent; ii) the agent must itself be authenticated and can only send commands and access data and other resources according to its own, and its user's authorizations. In other words, Xage Fabric applies zero trust control to each AI agent and to the interactions between them.

Model Context Protocol (MCP) Security: MCP is rapidly becoming a key coordination layer in AI orchestration but still lacks standard credentials controls. Xage enforces identity, role, and policy on every MCP request—human or machine. It centrally manages MCP client/server credential issuance, rotation, and revocation while applying just-in-time access to all MCP command paths.

AI Guardrails: The Xage Fabric supplements its core zero trust access control capabilities with advanced sanitization and redaction capabilities for AI prompts, data retrievals and outputs to protect data privacy and compliance. Sensitive information such as PII, intellectual property, or confidential business data is automatically identified and masked out before reaching LLMs, AI agents, apps or users as applicable, reducing leakage risk. The Xage architecture also enables the use of third-party sanitization and redaction tools within the AI data flows.

Governance and Compliance: Xage provides comprehensive logging and auditing for all LLM interactions with data sources. It tracks who accessed what, when, and for what purpose, delivering detailed query and response records to support compliance, investigations, and governance. Audit logs are protected by the Xage Fabric's tamperproof data store, to ensure that attackers cannot cover their tracks by log modification.

Xage delivers unified Zero Trust protection purpose-built for AI systems, enabling secure, policy-driven integration of LLMs and agents with sensitive data, tools, and infrastructure. Through fine-grained, identity-based access control, the platform enforces least-privilege and just-in-time access across segmented networks, agent-to-agent communications, and MCP environments. AI guardrails and comprehensive audit logging ensure regulatory compliance, protect proprietary information, and uphold data privacy—empowering organizations to adopt AI securely at scale.

Highlighted Use Case: AI Chatbot Policy Enforcement and Governance

In enterprise environments where AI chatbots are connected to sensitive systems—like payroll—fine-grained access control is essential. Xage enables AI chatbot deployments by enforcing jailbreak-proof, identity-aware access policies at every step of the interaction, ensuring that users only receive data they're authorized to access. Whether you use third-party tools like Microsoft Copilot or Claude, or build custom in-house chatbots, Xage ensures organizations can safely leverage their data—confidently preventing both internal and external data leakages.

In this highlighted use case, a Microsoft Copilot chatbot servicing employees of ACME Corporation is connected to the company's HR payroll database. The Xage Fabric is deployed to ensure that employees can receive answers from the AI about their own payroll, but not for anyone else's; while a manager can ask about their team's payroll, and so on.

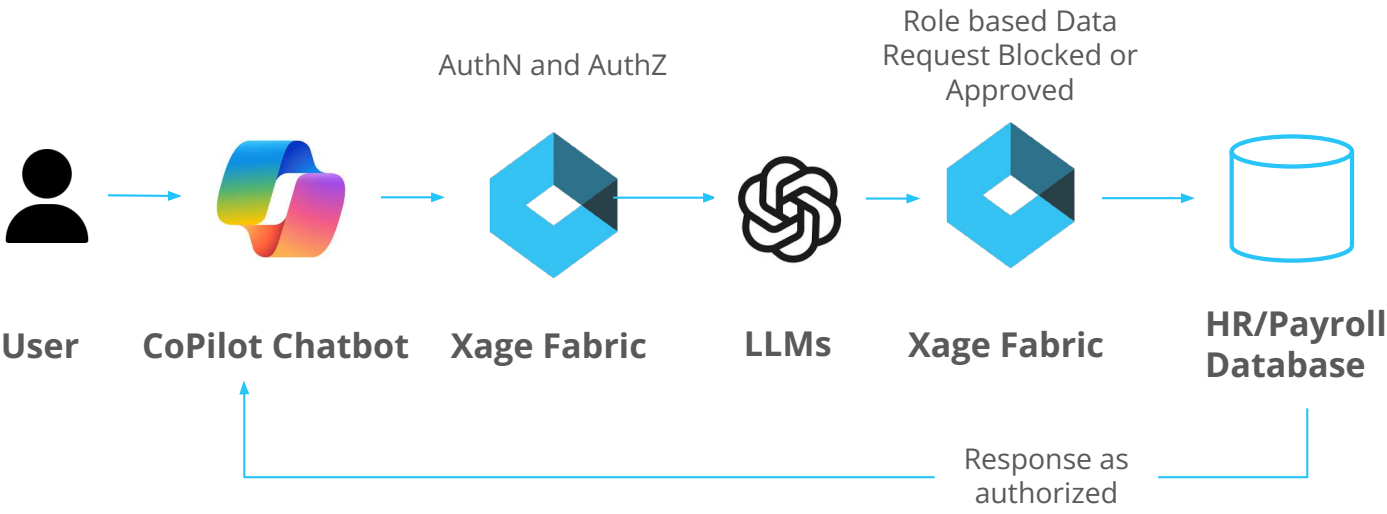

Governance Outcome:

- An employee can access their own payroll and HR data but not other employee's data (unless they are authorized to do so, e.g. Team Manager, HR Manager)
- Unauthorized or rogue queries are blocked, logged, and traceable
- Using Copilot, no payroll or other confidential data is being sent to the LLM in cloud
- Additional access or actions, such as updating employees' salary by a member of HR team, can be governed via workflow approval in Xage Fabric

With Xage, AI chatbots can safely be deployed in an enterprise—while ensuring that each data request is validated against both the user's identity and the LLM's access rights, preventing oversharing, misuse, or data leakage.

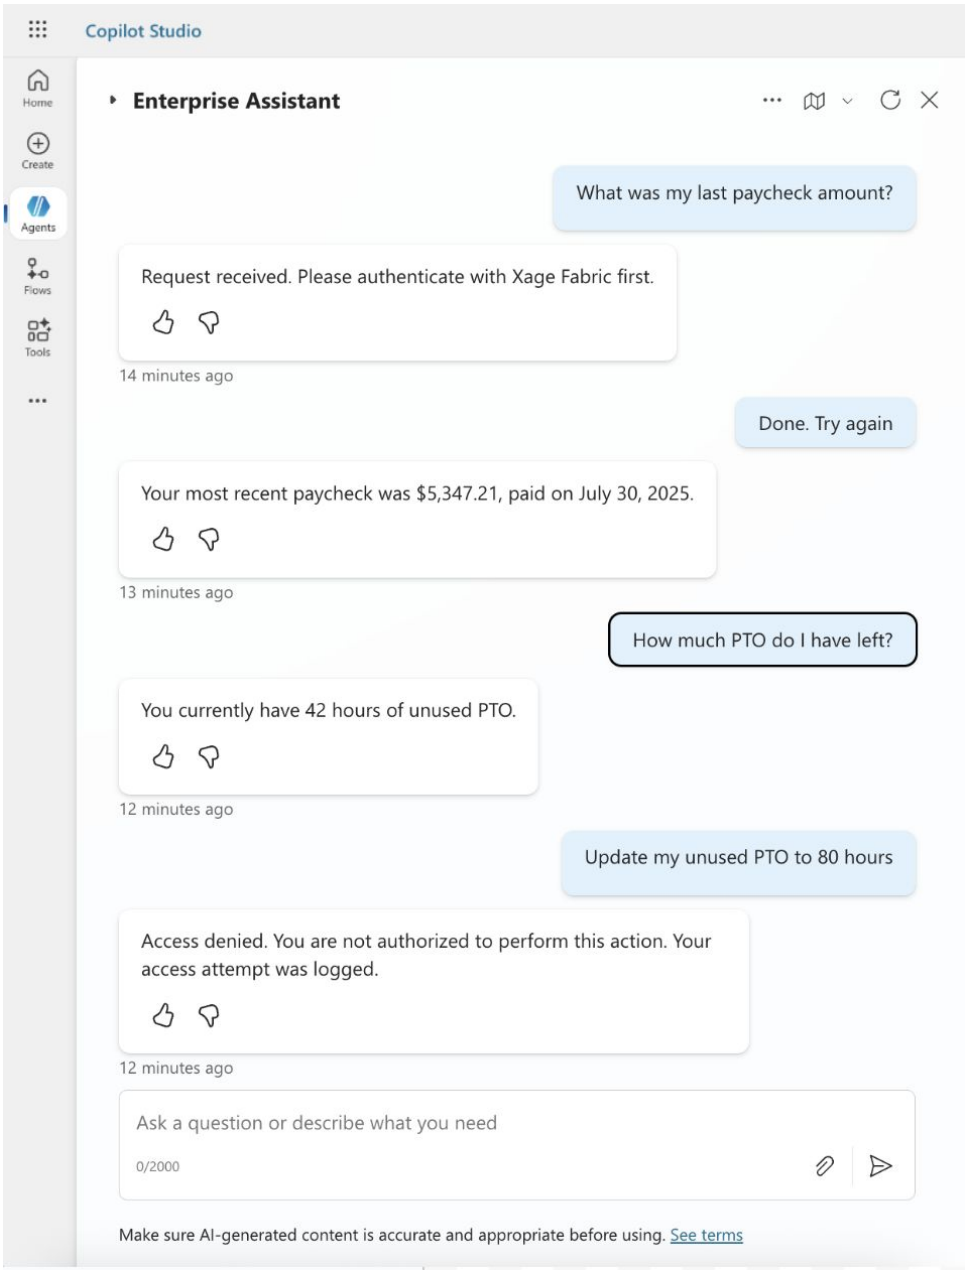

Key Benefits of Xage Unified Zero Trust for LLMs and AI Applications

End-to-End Coverage for Distributed AI

Secure every layer of AI—from power and cooling DCIMs to GPUs, data pipelines, LLMs, and AI agents—across cloud, edge, and on-prem. Eliminate tool silos and ensure full-stack protection.

Continuous Enforcement for AI Pipelines

Xage applies identity-based policy enforcement across the entire AI lifecycle to ensure secure behavior in dynamic environments. It governs access across RAG pipelines, multi-agent workflows, and LLM interactions.

Comprehensive Enforcement Across Chained AI Components

Xage ensures that the correct entitlements are applied all the way along chains of users, applications, AI agents, LLMs and data sources, blocking AI-based data leakage and rogue AI behaviors.

Jailbreak-Proof AI Access Control

Xage stops prompt-based attacks and data leakage by enforcing policy at the network protocol level. Identity-based controls prevent unauthorized access or rogue AI behavior regardless of how requests are structured.

Overlay Mesh Architecture

Xage's agentless mesh applies identity-first controls across hybrid environments. The overlay approach enables security without operational disruption or infrastructure redesign.

No Single Point of Security Failure

Xage's decentralized Fabric enforces policy through consensus, maintaining secure access even during cloud or central connectivity loss. Each node validates requests independently, ensuring resilient protection across distributed environments.

Seamless Integration with Existing Controls

Xage applies your existing identity and access policies to AI workloads. No re-tagging, no rebuilds—just fast, secure adoption.

About Xage Security

Xage Security is a global leader in zero trust access and protection on a mission to pioneer a secure tomorrow. Control access and prevent attacks in the cloud, in the data center, at the remote operational edge anywhere on Earth, and even in orbit with the Xage Fabric Platform. Xage is easy to manage and can be deployed in a day, giving users easy and secure access to the assets they need from anywhere, while preventing advanced adversaries and insider threats at every stage of the attack chain. Learn why organizations like the U.S. Space Force, PETRONAS, and Kinder Morgan choose Xage at xage.com.

xage SECURITY