How Zscaler Secures the Agentic AI Era with Zero Trust

June 18, 2026

AI just crossed a threshold that changes everything for security teams.

For two years, the enterprise AI story was about productivity. Faster research, smarter writing, better decisions. That was the warm-up. What's here now is categorically different: AI agents that don't just generate answers; they take action.

They query your databases, call your APIs, trigger workflows, move data across systems, spawn sub-agents and much more They do all of this at machine speed, with identities that are ephemeral, permissions that are often over-broad, and behavior that most security tools were simply never built to see.

At Zenith Live 2026, we announced exactly what enterprises need to govern this new reality: the industry's first complete Zero Trust platform for Agentic AI.

Not a proof of concept. A deployable architecture built on the Zero Trust Exchange™ that already processes 750 billion transactions a day.

Why Traditional Security Models Are Not Enough Against Agentic Threats

Legacy security was designed around humans: known identities, predictable access patterns, static directories. AI agents break every one of those assumptions.

An agent may carry valid credentials, act on a legitimate user's behalf, and interact with approved systems. This can pose a serious risk if it's over-permissioned, loosely governed, or invisible to your security stack. The challenge isn't just what an agent can access—it's what it's allowed to do once access is granted.

Anthropic recently made this point directly in their Zero Trust for AI Agents framework: perimeter-based defenses cannot keep pace with AI-accelerated threats. Their conclusion aligns with ours: Zero Trust isn't just relevant for the agentic era, it's the only model built for it.

Zscaler has successfully demonstrated for years how Zero Trust works at scale for users, branches, and cloud workloads. We're now extending that same architecture, with new purpose-built capabilities, to AI agents.

Here's what we launched at Zenith Live.

Zscaler AI Broker

AI agents communicate with each other and with enterprise data through emerging protocols like MCP (Model Context Protocol) and A2A (Agent-to-Agent). Most security tools can't see these channels at all.

AI Broker sits inline on these communications, enforcing fine-grained access controls across every agent interaction. The integrated Agent Registry gives your team a clear, governed view of what each agent is permitted to access and enforces it in real time. No more black-box agent activity.

Zscaler AI Access Graph

This is the visibility layer that makes everything else possible. Powered by our acquisition of Symmetry Systems, AI Access Graph maps how identities, AI applications, and data sources connect across your enterprise in real time. It surfaces over-privileged access before it becomes a breach, tracks data lineage across every channel, and integrates directly with the Zero Trust Exchange so you can move from insight to enforcement in the same platform. When an agent touches your data, you'll know exactly who authorized it, what it accessed, and where that data went.

Zscaler Endpoint AI Security

Your endpoints are already running AI whether IT knows about it or not. AI-powered IDEs, local models, browser plugins, developer extensions are the layers that legacy endpoint tools were never designed to inspect.

Endpoint AI Security reaches into exactly those layers to detect AI-related threats, enforce policies, and stop risks that traditional EDR solutions miss entirely. It's Zero Trust enforcement at the device level, for the AI era.

Major Enhancements to Zscaler AI Protect

Building on AI Protect, launched in January 2026, we're also shipping significant new capabilities across all three pillars:

- AI Asset Management: Now discovers embedded AI in SaaS and internet traffic, identifies AI agents and MCP servers in public cloud environments, scans agentic codebases for risk, and extends visibility to AI activity on endpoints.
- Secure Access to AI: Prompt extraction controls now cover 2,900+ GenAI apps, with full conversational views, Anthropic and OpenAI Compliance API support, and intent-based guardrails for multi-turn agent conversations.
- Secure AI Infrastructure and Apps: New AI red teaming for MCP servers, a standalone prompt hardening service, and compliance heat maps to strengthen AI governance across your environment.

The Bottom Line

Enterprises don't need to slow down their AI adoption. They need security infrastructure that can keep pace with it.

AI agents are a new class of digital actor: autonomous, fast, and capable of operating at a scope and scale that humans can't match. Governing them requires the same Zero Trust discipline that transformed how we secure users and cloud workloads. It just needs to be applied with more precision, coverage, and urgency.

This is what Zscaler has built, and it's available now.

Ready to see it in action? Learn more and schedule a demo.