Zero Trust for a Workforce That Never Sleeps:
Zscaler Stakes Its Claim on Agentic AI Security

At Zenith Live in Las Vegas, Zscaler laid out its bid to become the security backbone for the coming wave of enterprise AI agents — and argued that governing them is now an architectural decision for the board, not a feature to bolt on later.

Zscaler built its name on a simple, disruptive idea: replace the traditional network perimeter with a zero-trust platform where nothing is trusted by default. At its Zenith Live user conference in Las Vegas, the company unveiled what it sees as the next chapter of that story — positioning itself as the foundational zero-trust layer for agentic AI. The message from CEO Jay Chaudhry's opening keynote was blunt: securing AI agents, along with their connections, data paths, and device footprint, has become a board-level architecture question, and answering it will force enterprises to rethink security from the ground up.

AI Agents Are the New Risk Surface

Chaudhry framed agentic AI as the next great technology wave after cloud and mobile — except this one is arriving faster and carries a fundamentally different risk profile. His warning to the audience: enterprises should expect dozens of AI agents for every single employee, each one running around the clock, capable of spawning additional agents, and autonomously reaching into corporate systems and data. Unlike people, agents never clock out, never sleep, and can multiply themselves — a shift that moves the threat model from human-centric to machine-centric.

That reframing has a sobering implication: human users become just one slice of a much larger digital workforce, and in many cases the agents will hold more privileges than the people. The governance mechanisms built for humans — periodic access certifications, security training, manual approval chains — simply cannot keep up with software making decisions in milliseconds. Chaudhry's argument is that the industry needs a new control plane anchored in identity, data, and application context rather than in networks and IP addresses, because policies written for people break down when machines are the ones acting.

The Zero Trust Exchange Becomes an Agent Fabric

The second big storyline from the keynote is the evolution of Zscaler's flagship Zero Trust Exchange. What began as a fabric connecting users to applications is being extended into an agent fabric — one that brokers interactions among users, workloads, and AI agents alike. Zscaler's long-held thesis is that the internet itself can serve as the corporate network, with applications hidden behind a policy-driven exchange. The company is now extending that same logic to AI agents, treating each one as an untrusted outsider until identity and policy prove otherwise — exactly the posture it has always applied to human users.

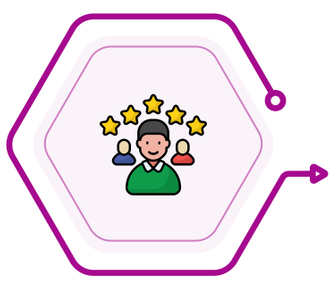

For existing customers, the continuity matters. Organizations that have already standardized on Zscaler for zero trust and Security Service Edge will not need to stand up a separate, parallel AI security stack. Instead, AI agents onboard into the same fabric that already connects users and apps, where the platform can enforce least-privilege access, keep internal applications hidden from direct exposure, and watch every interaction for anomalous behavior. It positions agentic security as an extension of the architecture customers already run — not a rip-and-replace project.

AI Broker: A Policy Engine, Not Another Gateway

On the product front, the headline announcement was the Zscaler AI Broker, which sits between AI agents and the systems they touch — including agents built on the Model Context Protocol and direct agent-to-agent interactions. Paired with an integrated Agent Registry, the Broker tracks each agent's identity, its stated purpose, and the data and actions it is permitted to perform. That enables fine-grained policies: a finance agent can be confined to specific systems, while a customer-support agent can be walled off from personally identifiable information.

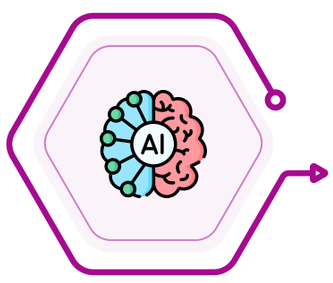

This represents a step beyond the first generation of AI gateways, which concentrated mostly on filtering prompts and routing traffic between models. Chaudhry positioned the Broker as the control plane for an emerging agent ecosystem — the point being not merely to observe what agents do, but to define what they are allowed to do from the outset. For enterprises experimenting with internal orchestrators and AI frameworks, that centralized governance offers a way to contain the blast radius of any single agent and to demonstrate compliance by treating agents the way mature organizations treat highly privileged service accounts: with continuous authorization.

Endpoint AI Security: Bringing Shadow AI Into the Light

Zscaler also acknowledged an uncomfortable truth: a large share of AI experimentation happens at the endpoint, through browsers, extensions, local tools, and plugins — far from the sanctioned cloud stack. Its answer is Endpoint AI Security, which extends the company's visibility into AI activity on devices. The capability detects and blocks behaviors like malicious browser extensions operating as agents, unmanaged AI tools reaching into sensitive files, and data quietly leaving the building through AI assistants.

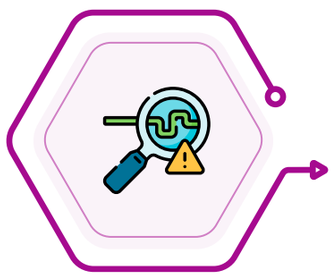

Notably, Zscaler is not trying to morph into a traditional endpoint detection and response vendor. Instead, it is correlating what it already sees — encrypted traffic and SaaS usage at massive scale — with endpoint-level AI behavior. The aim is to rein in shadow AI without smothering the experimentation that drives value. Chaudhry's framing captured the dilemma well: employees are going to use AI whether or not a policy exists, so the real question is whether the organization has visibility and control when they do.

AI Access Graph and AIGuardian: From Telemetry to Governance

The final piece is the AI Access Graph, built partly on technology from Zscaler's Symmetry Systems acquisition. The graph maps how identities, data, and applications connect across the enterprise, making it possible to answer the questions regulators and boards increasingly ask: which users and agents can reach a given sensitive dataset, and what chain of access led to a specific AI action. That kind of lineage is fast becoming table stakes for AI governance.

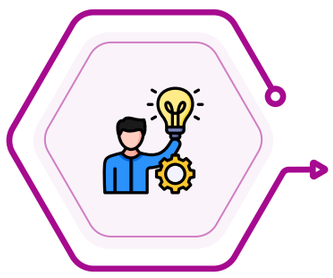

The Access Graph rolls up into a broader initiative called AIGuardian, which bundles Zscaler's Zero Trust Everywhere framework, the AI Broker, Endpoint AI Security, and the Access Graph together with consulting and integration support from global system integrators. The packaging reflects a candid recognition: securing agentic AI is as much an operating-model problem as a technology problem, and many large enterprises will need a guided path — not just products — to move AI from pilot purgatory into production.

Why Mythos Strengthens Zscaler's Hand

One topic that surfaced in the analyst Q&A rather than the keynote was Mythos — the frontier-model development that has rattled parts of the SaaS security market. Chaudhry's take was that Mythos actually creates a long-term tailwind for Zscaler, because it validates the company's core thesis: eliminating the attack surface matters more than racing to patch every new vulnerability, particularly in an era when frontier AI can discover and weaponize bugs at machine speed.

The doomsday narrative around Mythos assumes AI-accelerated vulnerability discovery is an existential threat to SaaS security vendors. But Zscaler's architecture differs structurally from a typical exposed SaaS application: the Zero Trust Exchange is designed to hide applications from the public internet, eliminate public IPs and open ports, and make users and workloads reachable only through identity- and policy-driven connections. Early findings from Anthropic's Project Glasswing — and the lessons of the Claude Mythos leak — point in the same direction: the most catastrophic exposures overwhelmingly trace back to misconfigured internet-facing services, not exotic exploits. That reality reinforces Zscaler's long-running mantra that anything reachable is breachable, and that shrinking what is reachable is the only sustainable defense against AI-driven reconnaissance. As an early Glasswing partner feeding telemetry from hundreds of billions of daily transactions, Zscaler is betting it can turn the same frontier AI that frightens the market into a hardening engine for its own stack and its customers' attack surfaces.

The Takeaway for CISOs and CIOs

The clearest message from Zenith Live is that AI security can no longer wait for AI projects to mature. Chaudhry acknowledged what many security leaders already know: most organizations are stuck in pilot mode not for lack of AI ideas, but because they do not trust their own ability to govern AI access to sensitive systems and data. By extending zero-trust principles to agents and anchoring everything in a single platform, Zscaler is offering a credible bridge from experimentation to production — with guardrails.

01

Zooming out, the rising tide of enterprise AI is likely to lift the whole Secure Access Service Edge and SSE category, which makes it critical for Zscaler and its peers to articulate clearly how their agentic strategies integrate with — or compete against — the emerging crop of AI security fabrics. For enterprises, the blueprint coming out of the event is straightforward: converge user, application, and agent connectivity onto one zero-trust fabric; treat every AI agent as untrusted but governable; and build AI governance on data-centric visibility instead of network topology.

02

None of this is surprising coming from Zscaler's CEO — the pitch fits the company's DNA. But the underlying logic is hard to dismiss. IT complexity has only compounded, and AI multiplies it by an order of magnitude. Zscaler's founding principles — shrink the attack surface, limit lateral movement, minimize the blast radius of any breach — were built for exactly this kind of moment. As agents flood into the enterprise at machine speed, those fundamentals may decide which security teams keep pace with the business and which fall behind.