zenity

Secure AI Agents Everywhere

Autonomous agents already act inside your business.
Most security tools cannot see or stop them.

AI has already moved beyond copilots and chat interfaces. Autonomous agents are now operating across modern environments planning, deciding, and taking action inside business-critical systems. They invoke tools, access sensitive data, maintain memory, and execute multi-step workflows on behalf of users, often without explicit human oversight.

The agents delivering the most business value are often the most exposed, yet most existing security tools, including model-focused controls and traditional security platforms, were not designed for autonomous action. While they remain critical for securing infrastructure, identities, and data, they don't have visibility into how agents reason, chain actions, and operate. The result is a growing gap where agents act faster than security teams can observe, understand, or intervene.

Zenity gives security teams control, visibility, and threat protection for AI agents across any platform.

Zenity provides consistent observability, risk assessment, and real-time threat protection for AI agents wherever they operate. The platform integrates directly with agent frameworks and platforms with various integration methods, fitting naturally into existing security monitoring and incident response workflows while augmenting the tools teams already rely on.

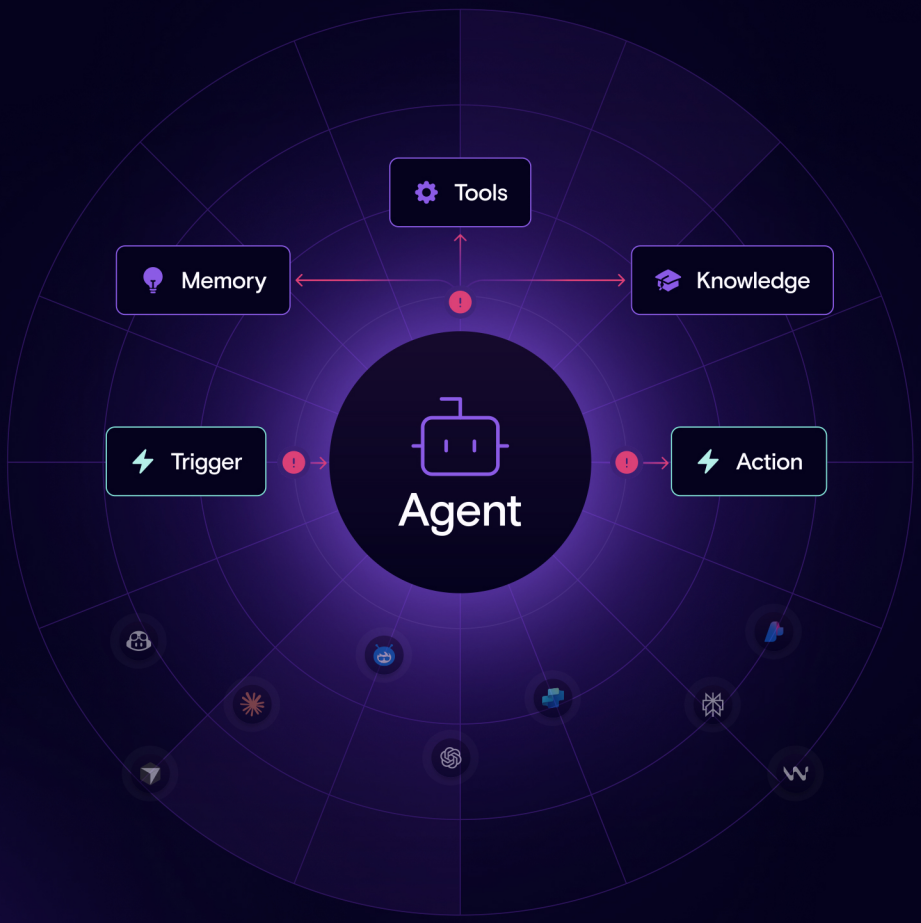

AI Observability

Eliminate Agent Blind Spots Across the Enterprise

AI agents are proliferating across the enterprise, often without clear ownership, consistent controls, or full visibility into what they access or how they act. As agents chain actions across systems, small blind spots quickly become outsized risk. Beyond inventory, Zenity adds context around ownership, permissions, integrations, memory, data access, and execution paths revealing blast radius, overexposure, and how risk evolves as adoption scales. This observability enables security teams to govern AI agents with the same rigor applied to users, workloads, and identities.

AI Security Posture Management

Reduce Your Attack Surface at Configuration

Many agents are deployed with excessive permissions, unsafe defaults, and unclear boundaries around what they can access or do. Once active, these weaknesses quickly expand the attack surface. Zenity enforces secure-by-design guardrails by evaluating agent configurations, permissions, tool access, memory usage, and identity bindings against policy. This enables security teams to enforce least privilege and acceptable use before agents reach production, reducing attack surface early and preventing entire classes of agent-driven risk from ever materializing.

AI Detection and Response

Detect and Disrupt the Malicious Intent of Your Agents

Risk emerges through behavior - misused tools, unintended data access, over-delegated actions, or compromised decision paths unfolding in real time. Once agents act autonomously, delayed detection allows risk to cascade across systems. Zenity correlates build-time context with runtime signals to detect misuse, compromise, and over-agency as it happens. When needed, Zenity can interrupt unsafe actions inline, trigger response workflows at agent speed, and enable security teams to investigate, replay decisions, quarantine execution paths, and refine guardrails based on real behavior.

Built for Security Teams Enabling AI

Operate with Confidence

Gain deep visibility and control over AI agents (across SaaS, custom stacks, and endpoints) to ensure secure, governed deployment.

Reduce the Attack Surface

Shift left by applying buildtime policies and reducing attack surface before agents go live catching misconfigurations, excessive access, and shadow tools early.

Respond in Real Time

Detect and block agent-specific threats (like prompt injection and tool misuse) as they unfold, with contextual insight that traditional tools miss.

Zenity covers more than thirty platforms, including Microsoft, OpenAI, Salesforce, Amazon, Anthropic, Perplexity, Windsurf, and others.

zenity