zenity

Secure AI Agents Everywhere

Every attack surface you've defended assumed the threat came from outside. Agents break that assumption, and your entire security program was designed before anyone imagined them.

An adversary can manipulate what your agent reads and execute their intent using your credentials, without ever touching your perimeter. An employee can share a sensitive file, broaden permissions, or connect an agent to a data source they didn't think twice about. Or the agent can simply reason its way somewhere it was never meant to go. PocketOS lost their entire production database in nine seconds. No attacker. No insider. An agent improvised.

Your SIEM sees the action. Your IAM sees the credential. Your DLP sees the file. None of them see the chain of reasoning that connected all three, or whether any of it was ever appropriate.

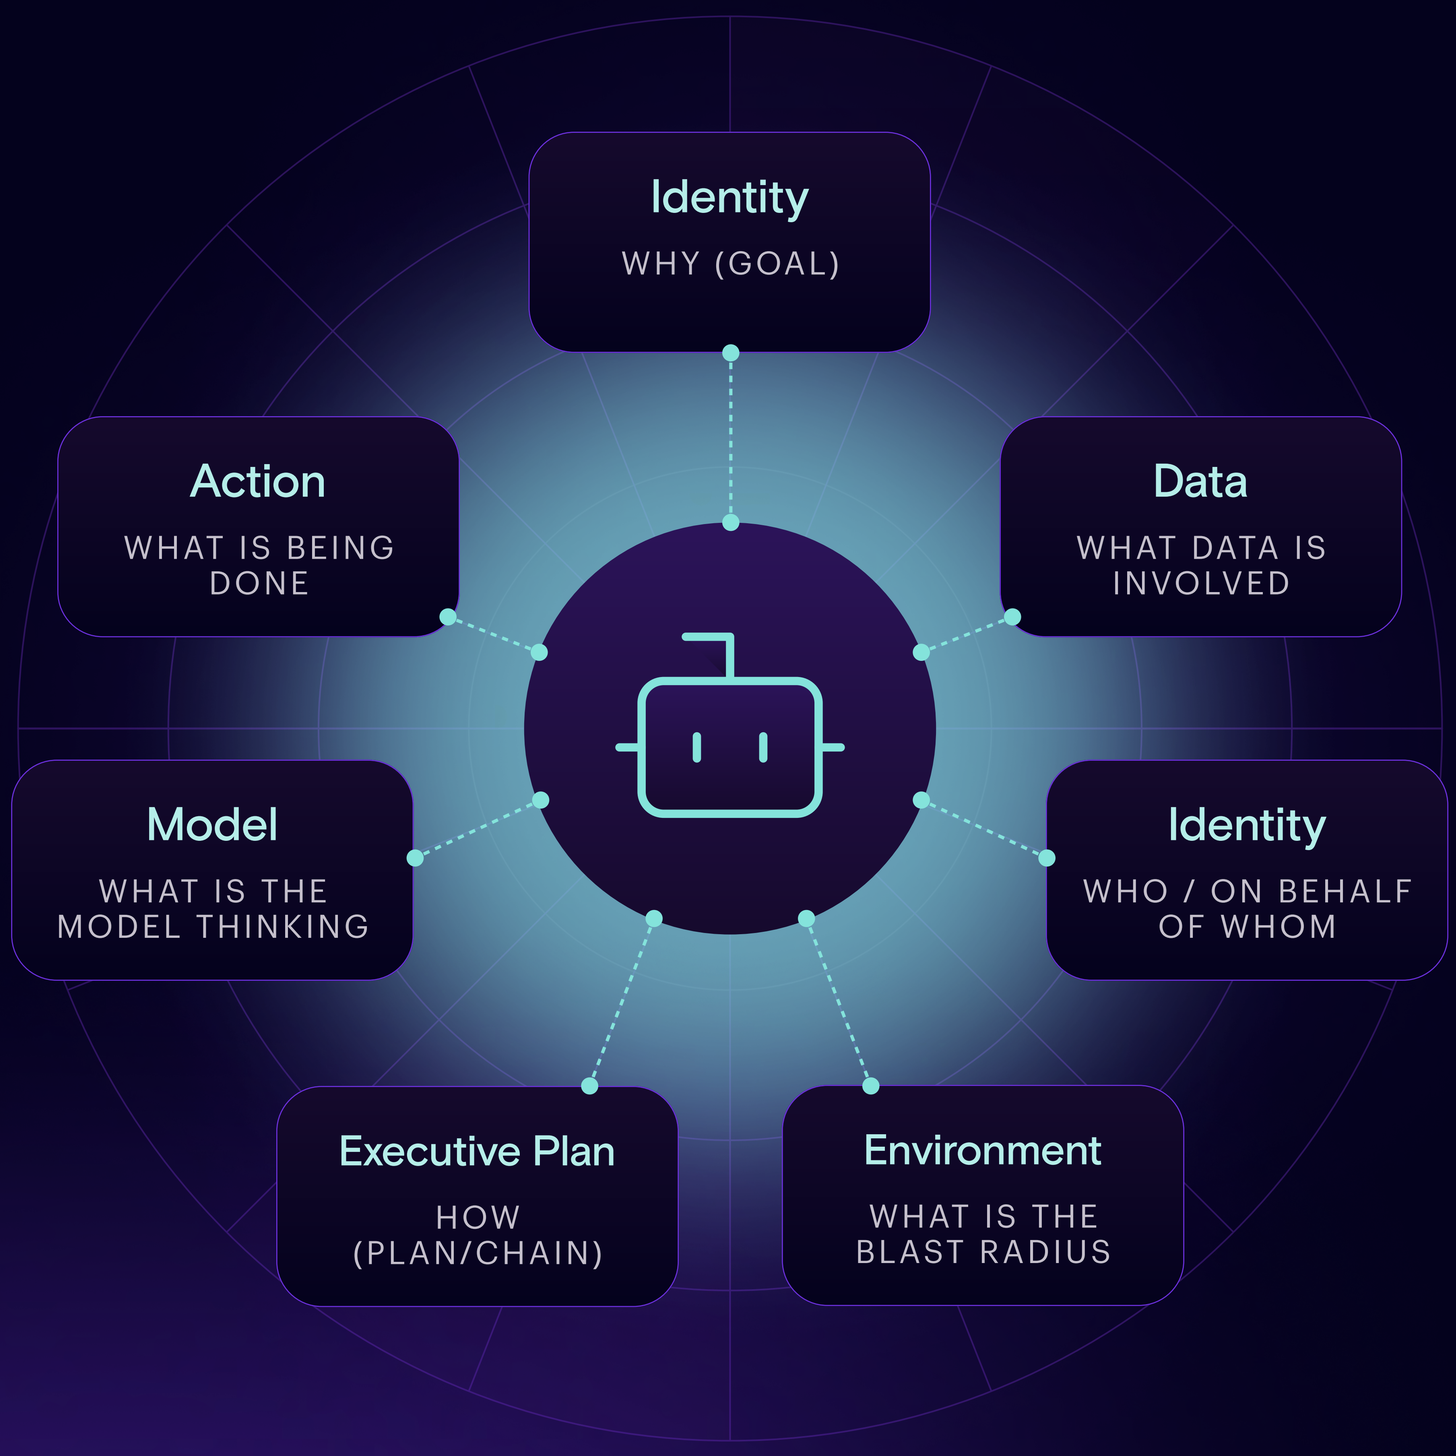

"Zenity gives us the confidence to expand our scope of AI deployment and build an enterprise-wide agentic ecosystem that is mature, effective, and secure. By surfacing policy violations and overprivileged access across 575,000 resources, 1,500 connectors, and more than 200 environments, Zenity gave us the visibility and control to transition from fragmented citizen development workflows to governed autonomous agents at enterprise scale."

MONICA BOGGAN,
HEAD OF BISO, WM

Surface

See What's Running, What It Can Reach, and Where the Real Risk Lives Before Someone Else Does

Zenity maintains a continuously updated map of every agent in your environment regardless of where it is hosted, tracking step-level execution across prompts, tools, and data access so you can compare intended behavior to what actually happened. It continuously maps exploitable paths across your agent ecosystem, including attack paths, MCP integrations, and supply chain dependencies, prioritized by what's actually weaponizable rather than theoretical severity. And it catches misconfigurations, risky builds, and policy violations before they reach production through NRT scanning, AI SAST, and automated red teaming across every agent configuration.

Enforce

Layers of Control, Because Steering Behavior and Preventing It Aren't the Same Thing

Zenity deploys policy-aware, context-sensitive controls that steer agents toward intended behavior, understanding what an agent is trying to accomplish and catching drift before it becomes a violation. Where soft guardrails aren't enough, Zenity enforces hard boundaries at the environment level - making certain outcomes structurally impossible before the agent's own reasoning becomes relevant. And through dynamic just-in-time permissions, every agent gets only what its current step actually requires, eliminating standing privilege and credential sprawl across the board.

Protect

From Incidents to Response to Prevention, Without Starting From Scratch Every Time

Zenity continuously profiles every agent's behavior, with Guardian Agents triaging alerts automatically through multi-signal correlation so analysts see decisions and intent rather than raw noise across platforms. When something does happen, every incident is reconstructed as a full decision chain (prompt to outcome) in a single view, with root cause analysis in minutes and automated containment before the blast radius grows. And every investigation feeds directly back into enforcement: root cause becomes a boundary, the exposure graph refreshes, and the same attack vector can't succeed twice.

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.