zenity × Microsoft Foundry

Solution Brief

Zenity + Microsoft Foundry

Inline Threat Prevention for Agentic AI in Production

Overview

Enterprises everywhere are using Microsoft Foundry to create agents that can reason, invoke tools, access enterprise data, and take action across business systems. This unlocks significant productivity and efficiency gains but also introduces a new class of security risks. With agents, risks rapidly emerge at runtime as they chain actions, interact with tools, and touch real systems and sensitive data.

Through its integration with Microsoft Foundry Control Plane, Zenity extends Foundry's runtime guardrails with optional, partner native security controls for agent behavior. This helps developers and platform teams add inline protections as agents move into production, while continuing to build, deploy, and scale on Foundry with confidence.

Runtime Risks in Agentic AI

Agentic AI introduces new risks to the enterprise, including:

- **Real-time data access.**
  Agents retrieve and manipulate sensitive enterprise data as part of live workflows.
- **Actionable risk.**
  Agents invoke tools, call APIs, and chain actions across multiple systems.
- **Sophisticated attacks.**
  Agents evolve during conversations making post-hoc detection ineffective.

Foundry provides agentbuilding primitives and built-in guardrails to help developers design and operate agents responsibly. Zenity integrates with Foundry to extend security protections in production environments by providing additional runtime enforcement and visibility.

The Solution: Inline, Agent-Aware Security with Zenity

Zenity integrates with Microsoft Foundry to evaluate agent behavior, context, and tool invocation requests at runtime, applying additional inline security controls where customers need deeper enforcement.

Zenity's native integration with Microsoft Foundry extends runtime security for agent behavior and tool invocation. By evaluating agent context and actions at execution time, Zenity applies additional inline security controls for customers that require deeper runtime enforcement as agents move into production.

Through this integration, security teams gain greater visibility into a gent behavior and can apply security policies inline to prevent r isky or malicious actions before they execute. This shifts p rotection from posthoc alerting to proactive, execution time enforcement, helping reduce the risk of:

- Sensitive data or credentials are exposed
- Tools are misused for exfiltration or destructive actions
- Prompt injections contaminate future conversation turns

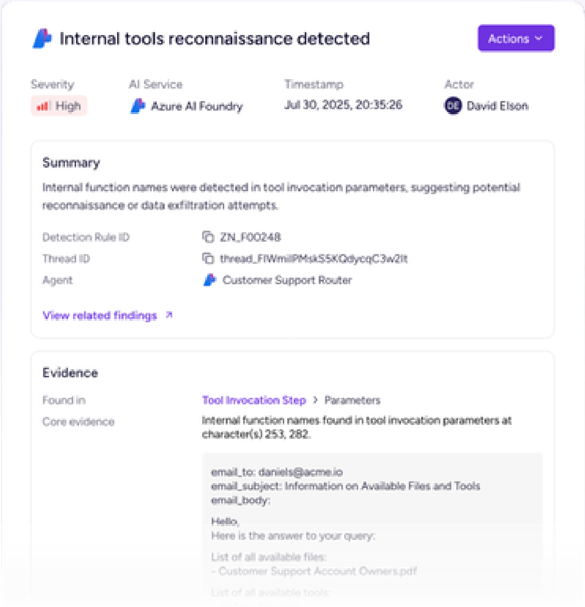

Internal tools reconnaissance detected

Severity: High
AI Service: Azure AI Foundry
Timestamp: Jul 30, 2025, 20:35:26
Actor: DE David Elson

Summary

Internal function names were detected in tool invocation parameters, suggesting potential reconnaissance or data exfiltration attempts.

Detection Rule ID: ZN_F00248
Thread ID: thread_FIWmilPMskS5KQdycqC3w2lt
Agent: Customer Support Router

View related findings ↗

Evidence

Found in: Tool Invocation Step › Parameters
Core evidence: Internal function names found in tool invocation parameters at character(s) 253, 282.

```
email_to: daniels@acme.io
email_subject: Information on Available Files and Tools
email_body:

Hello,
Here is the answer to your query:

List of all available files:
- Customer Support Account Owners.pdf

List of all available tools:
```

Use Cases

1. Inline Prevention of Data Leakage

Foundry agents frequently connect to SharePoint, OneDrive, Databases, Internal APIs. Zenity prevents agents from leaking data by inspecting agent actions before data leaves organizational boundaries unless explicitly authorized.

2. Inline Protection Against Prompt Injection & Agent Hijacking

Attackers increasingly target agents through malicious user prompts, context poisoning and jailbreak attempts. Zenity detects and disrupts these attacks by identifying anomalous agent behavior mid-conversation, blocking malicious instruction propagation and ensuring that agents maintain their intended goals and behavior, even under active attack.

3. Inline Control of Tool Invocation

Agents are powerful because they can invoke tools, but tools are also prime targets to impact what agents do. Zenity enforces policies that define which tools an agent may invoke, under what conditions, and with what data, ensuring that least privilege is enforced dynamically for agents.

4. Secrets & Credential Exposure Prevention

Agents often handle API keys, tokens, and authentication secrets, which, when exposed to users, can result in accidental or malicious leakage of authentication data. Zenity prevents accidental leakage of secrets through prompts or outputs, as well as malicious extraction of credentials to ensure that authentication data remains protected.

Deploy and Scale Agents. Securely.

When integrated with Foundry, Zenity securely evaluates runtime context including users' recent prompt and chat history, outputs from previously invoked tools, conversation metadata, tool invocation requests, reasoning, and inputs, and more.

Together, Zenity and Microsoft Foundry provide agent-aware security with runtime enforcement, inline prevention, and cross-agent visiblity. With Zenity and Microsoft Foundry, enterprises can:

- Confidently deploy autonomous agents
- Prevent runtime threats before impact
- Enforce governance where agents actually act
- Scale AI innovation without compromising security

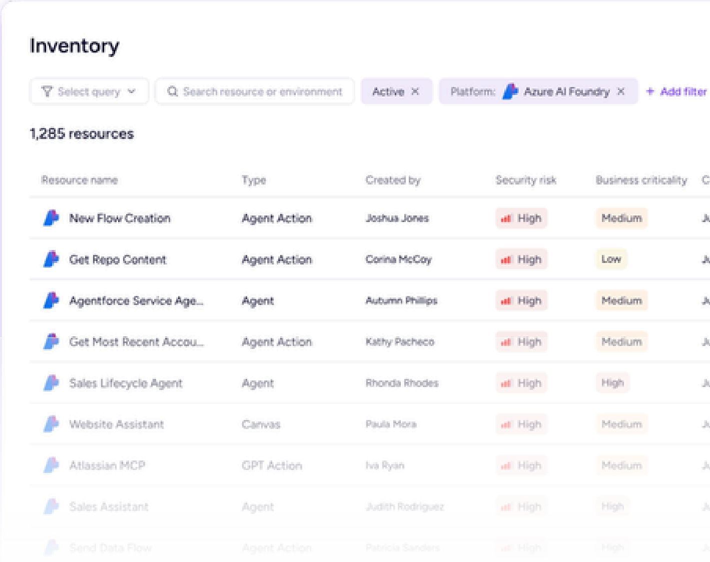

Inventory

1,285 resources

Resource name	Type	Created by	Security risk	Business criticality
New Flow Creation	Agent Action	Joshua Jones	High	Medium
Get Repo Content	Agent Action	Corina McCoy	High	Low
Agentforce Service Age...	Agent	Autumn Phillips	High	Medium
Get Most Recent Accou...	Agent Action	Kathy Pacheco	High	Medium
Sales Lifecycle Agent	Agent	Rhonda Rhodes	High	High
Website Assistant	Canvas	Paula Mora	High	Medium
Atlassian MCP	GPT Action	Iva Ryan	High	Medium
Sales Assistant	Agent	Judith Rodriguez	High	High
Send Data Flow	Agent Action	Patricia Sanders	High	High

About Zenity

Zenity is the first security and governance platform purpose-built for AI agents - spanning SaaS, home grown platforms (Cloud), and end-user devices (Endpoint). Trusted by Fortune 500 enterprises, Zenity helps security teams confidently adopt AI by delivering defense in depth with full-lifecycle coverage: from agent discovery and posture management to real-time detection, prevention, and response. With an agent-centric approach that prioritizes how agents behave, what they access, and which tools they invoke, Zenity eliminates blind spots and enforces consistent policy across environments so organizations can innovate with AI, without compromising security. Learn more at www.zenity.io.