SOLUTION BRIEF

SECURING AI WITH ORCA SECURITY

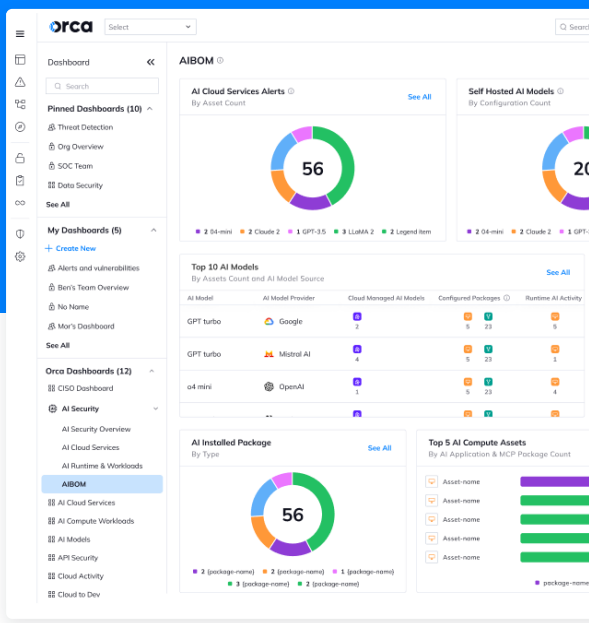

THE CHALLENGE:

AI Is Expanding Faster Than Teams Can Keep Up

AI adoption is rapidly transforming the enterprise risk landscape. Today, 84% of organizations use AI in the cloud, and 62% have deployed AI software packages with known vulnerabilities, while 90% of organizations still lack standards to defend against emerging AI-driven threats. As organizations deploy AI models, agents, and services across cloud environments, the attack surface expands dramatically. Security teams must now account for shadow AI services and agents deployed outside of governance, AI systems connected to sensitive data sources, AI identities and permissions with broad access, and autonomous AI actions occurring across cloud infrastructure.

This shift requires organizations to rethink how they understand and secure their environments. Security teams must now answer critical questions: Where is AI being used across the organization? What data and access do AI systems have? And what actions can AI perform across cloud infrastructure? and how can risky usage be detected and governed? Without clear visibility and control, AI can quickly introduce new pathways for data exposure, privilege misuse, and unintended automation across the cloud.

THE SOLUTION:

AI Security Across the Lifecycle

By combining AI discovery, security posture management, runtime monitoring, and AI-driven automation, Orca enables organizations to safely scale AI adoption while maintaining control over data, identities, and infrastructure.

- ✓ **Build:** Orca helps teams prevent AI risk early by identifying AI models, frameworks, and integrations in code and pipelines. Security teams gain visibility into AI components, dependencies, datasets, and tool integrations, allowing them to detect insecure configurations, excessive permissions, and risky model connections before they reach production.
- ✓ **Visibility:** Once deployed, Orca provides continuous visibility into AI exposure across cloud environments. Security teams can automatically discover AI services, models, and agents across multi-cloud infrastructure, understand how they connect to identities and sensitive data sources, and prioritize risks based on real exposure and blast radius.
- ✓ **Runtime:** Orca helps organizations detect and stop AI abuse by monitoring model activity, prompts, and AI-driven actions across workloads during execution. This visibility allows teams to identify suspicious AI usage, potential prompt injection attempts, unauthorized access to AI services, and risky automated behavior.

AI security requires visibility and protection across the full lifecycle from development to deployment and runtime operation. Orca Security helps organizations secure AI by providing early risk prevention, continuous posture visibility, and runtime monitoring across cloud environments.

Build: Prevent AI Risk Early

AI Bill of Materials (AI-BOM)

AI-BOM provides an inventory of AI components used across development environments. Orca identifies AI models, datasets, frameworks, libraries, agents, and integrations within repositories and CI/CD pipelines so teams can track dependencies and detect risky components before deployment.

AI Code & MCP Visibility

Orca analyzes source code and pipelines to identify how AI systems are integrated into applications, including Model Context Protocol (MCP) usage, tool integrations, model endpoints, plugins, and more, providing visibility into how AI systems interact with infrastructure, data sources, and services.

AI CI/CD Guardrails & Configuration Hardening

Orca monitors CI/CD pipelines for AI-related changes and analyzes configurations for insecure defaults or excessive permissions. Security teams can enforce policies, detect risky AI integrations, and prevent insecure models, prompts, or agents from reaching production.

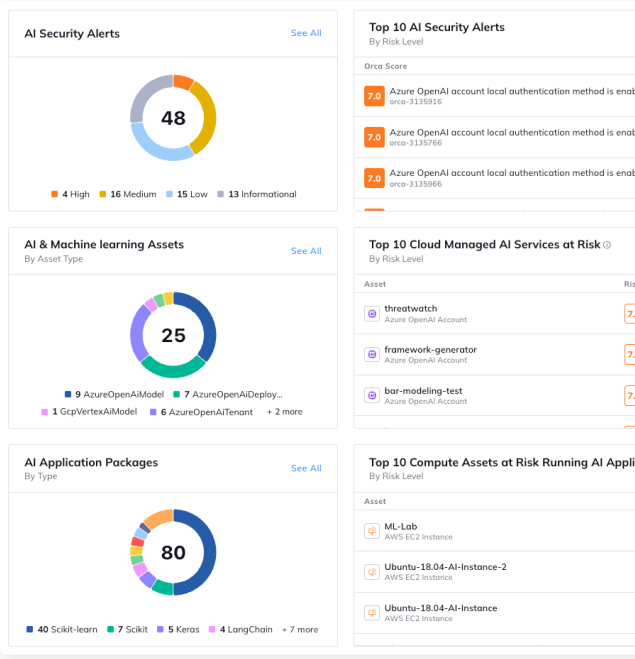

> *We use Orca to scan for any GenAI related vulnerabilities identified in packages and libraries in our cloud assets. Then we use this data to reach out to product teams for remediation."*

SRI TIRUVEEDHULA
Principal Engineer
Autodesk

Visibility: Understand AI Exposure

AI Inventory & Security Posture Management

Orca automatically discovers AI services, models, frameworks, and agents across multi-cloud environments. By analyzing configurations, exposed endpoints, identities, and permissions, Orca provides a centralized view of AI risk and helps teams understand blast radius and prioritize remediation.

AI Data Exposure Protection

Orca identifies sensitive data that may be exposed through AI workflows by detecting credentials, secrets, proprietary data, and regulated information in AI-related files, prompts, and storage locations. Teams gain visibility into where sensitive data resides and receive prioritized remediation guidance.

AI Vulnerability Management

Orca continuously scans AI workloads for vulnerabilities affecting AI frameworks, packages, models, and supporting infrastructure. Findings are enriched with contextual risk scoring so teams can focus on vulnerabilities that are exploitable or connected to sensitive systems.

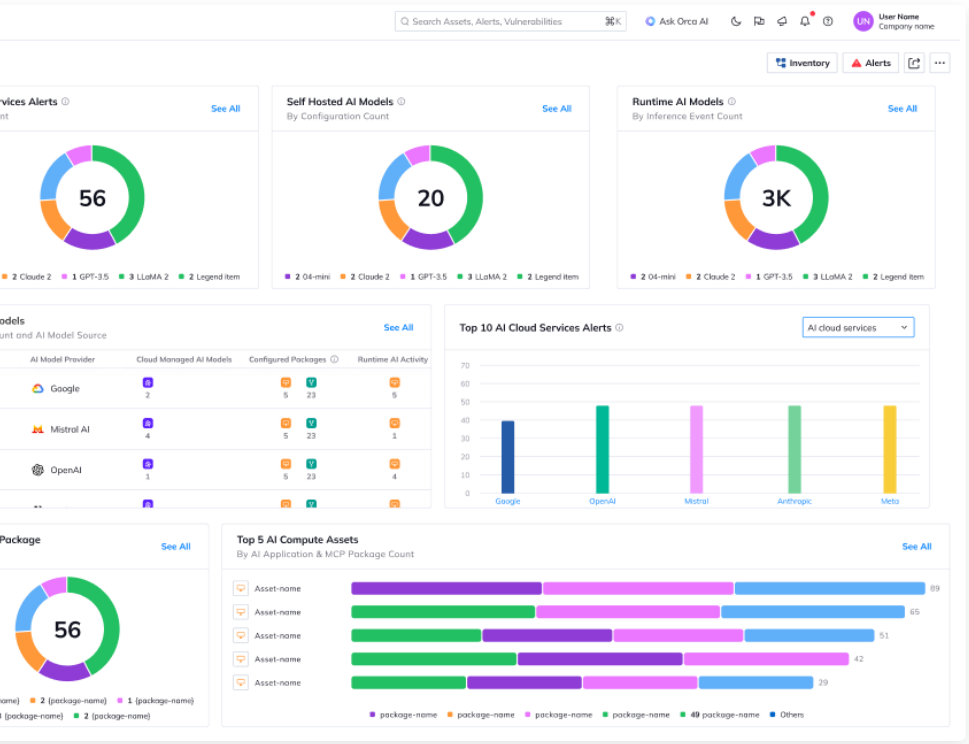

Runtime: Detect and Stop AI Abuse

AI Runtime Visibility

Orca provides real-time visibility into AI activity running across cloud workloads. Security teams can monitor AI providers, models, APIs, and endpoints in use, along with the identities and workloads interacting with them.

AI Prompt Monitoring & Behavior Detection

Orca analyzes prompts, sessions, and AI activity to detect malicious behavior, including prompt injection attempts, policy violations, sensitive data exposure, and abnormal AI access patterns. This helps identify misuse of AI services and suspicious automated actions.

AI-Driven Security Automation

AI Security Agents

Orca provides AI-powered security agents that automate investigations and security workflows. Organizations can deploy built-in deep research agents or create custom agents that integrate with security tools to triage alerts, investigate incidents, and automate repetitive security tasks.

About the Orca Platform

The Orca Cloud Security Platform is trusted by hundreds of organizations and identifies, prioritizes, and remediates risks and compliance issues across your AWS, Azure, GCP, Oracle Cloud, and Alibaba Cloud estates—without requiring a single agent. Orca deploys in minutes, and detects vulnerabilities, malware, misconfigurations, lateral movement, API risks, sensitive data at risk, anomalous events and behaviors, overly permissive identities, and more.

Learn more at https://orca.security.

Ready to try it out?
Sign up for a demo. Visit orca.security/demo