ORCA AI

The Reasoning Engine for AI Agents, Workflows, and Assistants

Orca AI is the reasoning layer of the Orca Platform. It powers Orca's own agents, workflows, and assistant. Through the Orca MCP Server, it extends that same reasoning into the AI tools your teams already use, from Claude and Cursor to your own custom LLM applications. One reasoning engine, grounded in one unified view of your cloud, code, and AI systems, wherever the work happens.

Defend at Machine Speed. Scale Beyond Human Limits.

AI-driven attacks are running faster than human responders can keep up with. AI Agents match that speed, running investigation, triage, and response at machine scale. Every response is grounded in the Unified Data Model, with reasoning shown at every step so your team stays in control.

01

Investigate at machine speed

AI Agents run correlation, evidence gathering, and triage in parallel across your entire queue, not one analyst at a time. Time-to-verdict shrinks and no signal waits.

02

Turn evidence into a verdict

AI Agents apply domain expertise to the evidence and reach a conclusion on every alert: real risk or false positive. What's proven gets escalated, what isn't gets dropped.

03

Free senior talent for the next threat

With AI Agents handling triage and remediation prep, senior staff shift from clearing queues to threat modeling, architecture review, and hunting the next attack pattern.

Visibility. Context. Security for the Way You Work.

Orca AI thrives off the first two pillars that power the Orca Platform, delivering its reasoning in the tools your team already uses instead of another console they have to visit.

01 / VISIBILITY

Patented SideScanning, the Orca Sensor, and native connectors give Orca AI a complete view of workloads, identities, code, data, and AI systems. An agent can only reason about a risk it can see.

02 / CONTEXT

Every asset, alert, identity, and dependency is already correlated in the Unified Data Model before the AI is asked a question, so answers factor in reachability, blast radius, and business impact.

03 / YOUR WAY OF WORKING

Agents investigate at machine speed, the MCP Server extends Orca into Claude, Codex, and Cursor, and findings land in the ticketing tools your team already uses.

AI Agents

Orca **Core Agents** are specialized AI teammates, grouped into pods that work together within a security discipline. Expertise alone isn't the differentiator: every agent runs on the same Unified Data Model behind Orca's risk prioritization and attack path analysis.

RED POD

Detects and Validates What Attackers Could Exploit

Attack Surface. On-demand AI-driven DAST and AI penetration testing that probes your public-facing assets. Each test starts with reconnaissance context already assembled from Orca's live graph, so findings arrive with real risk context.

Code Security Auditor. Traces cross-file data flows across your repository to reconstruct the attack chains an attacker would take, catching the classes pattern-matching SAST cannot see: broken authorization, privilege escalation, and business-logic flaws.

BLUE POD

Analyzes and Investigates Threats in Your Environment

Threat Investigation. Runs an alert's full investigation lifecycle, correlating signals, validating facts, and producing a transparent report with recommended containment actions.

AppSec Triage. Analyzes SAST and secrets findings using code context, data flow, and sanitization patterns to separate true positives from likely false positives, then adjusts risk scores accordingly.

GREEN POD

Turns Findings Into Fixes and Tracks Compliance

Green Pod agents close the loop between finding and fix, so validated findings become owned, tracked, verified work.

CUSTOM AGENTS

Shape One Around Your Environment

Start from templates and frameworks, or build from scratch, so an agent fits your environment, processes, and use cases.

Orca MCP Server and AI Skills Hub

Orca was the first cloud security platform to support the Model Context Protocol. The MCP Server extends Orca AI's reasoning into Claude, Codex, Cursor, and any MCP-compatible client, so security data goes where the work already happens: the IDE, the chat window, the terminal.

Ask Claude to triage an alert, and it calls Orca's MCP tools to pull the asset, the attack path, the blast radius, and the remediation steps in a single response. Teams use it for CVE remediation plans, executive reporting, and shift-left workflows.

The **AI Skills Hub** sits on top of the MCP Server: pre-built workflows with the logic already written, so analysts type "triage alert orca-9012345" and the skill handles retrieval, analysis, and formatting.

> **"The backlog dropped to 0, and now we're able to manage it and really hit meaningful SLAs as a business.**

Andrew Orr Ewing
Senior DevSecOps Engineer, Wayflyer
Read the Wayflyer case study →

ALSO WORTH A LOOK

Orca doesn't just extend into Claude. It also secures how your organization uses Claude and other AI apps.
Explore AI AppGen Security →

AI Assistant

The Orca AI Assistant is a natural-language interface built on Orca's full environment context, from cloud and identities to code and AI systems. Any employee, whether security engineer, developer, SOC analyst, or compliance lead, can ask precise questions in plain language and get answers in seconds.

It reasons over the same Unified Data Model that powers Orca's agents, so answers arrive with reachability, identity permissions, sensitive data proximity, and business impact already factored in.

Supported in 50+ languages, it closes the skills gap that forces teams to route every question through senior analysts.

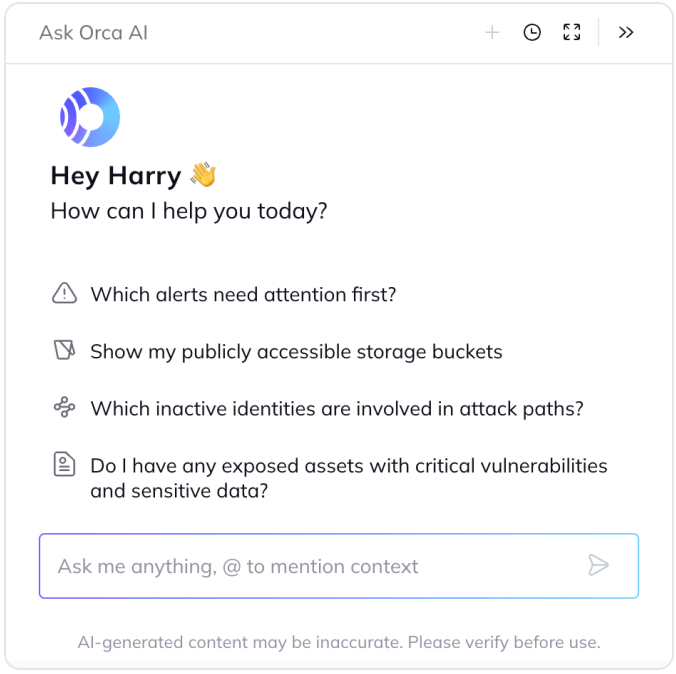

AI Discovery and AI Code Fixes

AI Discovery. Describe what you are looking for in plain language and Orca returns the query that finds it, ready to run, refine, or save. Ask for internet-facing databases holding sensitive data, and the query comes back written against the Unified Data Model. You find what you want without learning yet another query language.

AI Code Fixes. Orca AI generates remediation code for AppSec and IaC findings, with the option to open a pull request directly. Developers can copy it into a command line, paste it into their IaC tool, or ask follow-up questions to fine-tune the fix.

A finding without a fix becomes engineering's problem to decode. AI Code Fixes arrive as working code in the developer's own workflow, so the fix ships in the next commit instead of waiting in a queue.

Why the Foundation Matters

AI-powered security is the new normal. Every vendor will ship agents, so the question is no longer whether the AI exists. It is what the AI reasons against, who it works alongside, and where it can go.

Context is the price of entry. Orca AI reasons over the Unified Data Model that existed before AI arrived, so a question about an exposed endpoint returns with reachability, permissions, sensitive data proximity, and business impact factored in. Tools that stitch sources together at query time return fragmented conclusions.

The advantage compounds from there. A family of specialized agents works that same graph across cloud, AI, code, and runtime, each with its own memory and reasoning on display. The MCP Server carries that reasoning into Claude, Codex, Cursor, and whatever your teams adopt next, so you are never locked into one vendor's chat window.

ABOUT ORCA SECURITY

Orca Security delivers security for the companies that build. As cloud, AI, and app generation expand the attack surface, the Orca Platform provides complete visibility, correlates risk, and prioritizes what matters most. Trusted by SAP, Autodesk, Gannett, Lemonade, and Digital Turbine. Backed by Temasek, CapitalG, ICONIQ Capital, and Redpoint Ventures.

See Orca AI in Action

Orca AI is available today across the Orca Platform. Schedule a personalized 1:1 demo to see how it fits your environment.

Request a demo →