DATASHEET

AI SECURITY IN CONTEXT: FROM CODE TO CLOUD

Your pipelines now ship AI packages, model endpoints, provider keys and agents with their own permissions. Most AI security tools see one layer: the code, the model, or the agent. Orca connects all of them to the cloud resources and data they reach, so your team fixes the AI risk that's actually exploitable. Agentless, deployed in 15 minutes, and in the same queue your team already works.

Key challenges

AI moved into production faster than the controls around it. 90% of developers now regularly use at least one AI tool at work (JetBrains, 2026), and every one of them can add a model, a package or a key to the pipeline in a single commit.

AI runs with permissions nobody reviewed. An agent inherits the IAM role of its workload and the credentials of its MCP servers, so its risk is everything it can reach.

AI dependencies arrive with hidden payloads. Public model hubs and package registries now carry malicious models and poisoned packages that run code as soon as a pipeline loads them. In March 2026, backdoored LiteLLM releases on PyPI stole credentials from the AI toolchains that installed them (Cloud Security Alliance, 2026).

AI keys leak into code and configs. Leaked secrets tied to AI services rose 81% in 2025 to more than 1.27 million, and 24,008 secrets turned up in MCP configuration files on public GitHub (GitGuardian, 2026). One leaked OpenAI or Hugging Face token can expose models, training data and spend.

Coding agents act on untrusted input. In April 2026, researchers showed that a crafted pull request title or issue comment could make AI coding agents running in CI leak their API keys and GitHub tokens (Cloud Security Alliance, 2026).

Shadow AI escapes governance. Teams adopt models, coding assistants and AI services faster than security can inventory them, so AI risk builds up in assets nobody registered.

Point tools see one layer. One tool sees the model, another the code, a third the agent. None of them knows the exposed model endpoint sits two hops from a production database.

Benefits at a glance

- **Cloud, AI and code in one platform:** cloud, AI and code security share one Unified Data Model, replacing point tools and fitting the workflows you already run
- **One AI-BOM, code to cloud:** 50+ AI packages and frameworks in your repositories, plus Amazon Bedrock, SageMaker, Azure OpenAI, Google Vertex AI, self-hosted models, MCP servers and AI agents in your cloud
- **Context, not just findings:** every AI model, key and agent mapped to the IAM role, workload and data it can reach
- **Prioritized by exploitability:** AI risk scored on the same attack paths as the rest of your cloud
- **No new infrastructure:** read-only and agentless, with no gateway, proxy or endpoint agent to deploy
- **Deploy in 15 minutes:** connect your cloud and code once, and Orca starts mapping your AI risk

> ***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

©2026 ORCA SECURITY. ALL RIGHTS RESERVED.

Make AI security actionable with Orca

Find shadow AI without chasing engineering. See the models, AI packages and keys teams added on their own.

See every AI agent and what it can reach. Agents mapped to their models, MCP servers, IAM roles and data.

Fix the exposed package first, not the noisy one. Each instance is scored on its own attack path.

Catch risky code before it ships. Orca flags vulnerable code, secrets and IaC risks in the IDE, coding agents, PRs and CI.

Find flaws before frontier AI models do. AI analysis of your code surfaces broken access controls and zero-day flaws.

Trace every finding to its root cause. Production findings link back to the repository, owner and line of code.

One platform, one queue. AI findings land in the same alerts and workflows as the rest of your cloud.

Prove AI governance to auditors. Map AI findings to OWASP Top 10 for LLM Applications, ISO/IEC 42001 and the EU AI Act.

Orca capabilities

Orca covers AI across three layers: the code where it enters, the cloud where it's deployed, and the runtime where it acts. All three land in one Unified Data Model.

AI PLATFORMS

AI platform posture and DSPM

Secures access to Claude Enterprise and ChatGPT Enterprise, with DSPM for sensitive data in prompts and files

CODE

AI-BOM from repositories

Inventories AI packages and frameworks (PyTorch, TensorFlow, LangChain, Hugging Face and 50+ more) in GitHub, GitLab and Azure DevOps, with versions and known CVEs

CODE

AI secrets detection

Finds OpenAI, Anthropic, Hugging Face and other AI provider keys across repositories and full Git history

CODE

CI/CD guardrails

Orca CLI scans IaC, dependencies, secrets and container images in any CI tool and fails builds that break policy

CODE

Coding agent security

Finds Claude Code, Cursor and Codex across your repositories, ties findings to the agent that wrote the code, and scans agent edits as they happen with Orca hooks and skills

CLOUD

Agentless AI discovery

SideScanning™ finds cloud-managed AI services like Amazon Bedrock, SageMaker, Azure OpenAI and Google Vertex AI, plus self-hosted models, including shadow AI

CLOUD

AI posture management

Alerts on public exposure, over-privileged roles, missing encryption and insecure defaults on AI services

CLOUD

AI agent security posture

Maps each agent to its models, MCP servers, IAM role and reachable data; flags MCP configs holding live credentials and agents set to skip human approval

CLOUD

Data security (DSPM) for AI

Classifies PII and secrets in training data, AI projects and data stores your models reach

CLOUD

AI attack path analysis

Traces how an exposed model, key or agent leads to crown-jewel data, scored with the same dynamic risk model as every other finding

RUNTIME

Runtime AI Security (Orca Sensor, optional)

Observes LLM and MCP calls from each workload and detects prompt injection, jailbreak attempts and sensitive data in prompts

ALL LAYERS

AI compliance

OWASP Top 10 for LLM Applications, ISO/IEC 42001, the EU AI Act, Orca AI Best Practices and custom frameworks

Specifications

Area	Supported
Cloud providers	AWS, Azure, Google Cloud, Oracle Cloud, Alibaba Cloud and Tencent Cloud
Kubernetes	Managed and self-managed clusters, including OpenShift, in any cloud or on-premises
Managed AI services	Amazon Bedrock, Amazon Bedrock AgentCore and Amazon SageMaker; Azure OpenAI, Azure AI Foundry, Azure Machine Learning and Azure AI Content Safety; Google Vertex AI
Self-hosted AI models	Models served through Ollama, Hugging Face, vLLM and other popular frameworks on compute workloads in every supported cloud
Enterprise AI assistants	Claude Enterprise, ChatGPT Enterprise
AI app builders and hosting	Lovable, v0, Bolt, Base44, Replit, Vercel, Netlify, Supabase
Source code management	GitHub, GitLab, Azure DevOps and Bitbucket
CI/CD and developer tools	Orca CLI for any CI pipeline, with native integrations for GitHub Actions, GitLab CI, Azure Pipelines, Jenkins, CircleCI, AWS CodeBuild and more; pull request scanning with inline comments; VS Code and JetBrains IDE plugins; pre-commit hooks; Claude Code hooks and skills
Deployment	Agentless SideScanning™ with read-only cloud access, in SaaS, in-account or private mode; optional Orca Sensor (eBPF) for Runtime AI Security
Security and data handling	Certifications, data handling and privacy practices are published in the Orca Trust Center

About Orca

Orca Security delivers security for the companies that build. As cloud, code, AI and app generation expand the attack surface, Orca transforms security risk into the context teams need to act. The Orca Platform provides complete visibility across cloud, AI, and application environments, correlates risk across every layer, and prioritizes the exposures that matter most so organizations can remediate faster and innovate with confidence. Trusted by hundreds of organizations, including SAP, Autodesk, Gannett, Lemonade, and Digital Turbine, Orca is backed by leading investors including Temasek, CapitalG, ICONIQ Capital, and Redpoint Ventures. Learn more at orca.security and book a personalized demo.

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

©2026 ORCA SECURITY. ALL RIGHTS RESERVED.