SOLUTION BRIEF

Securing AI Workloads with Aqua Secure AI

Full Lifecycle Protection from Code to Cloud to Prompt

Secure Your Runtime. Reduce Risk. Stay Ahead.

AI workloads break the rules of traditional security. These applications are dynamic, autonomous, and unpredictable, often operating beyond the reach of legacy tools. Risk does not start and end in development; it escalates in production, where visibility and control are weakest. Security teams often lack awareness of where LLMs are deployed, have no insight into real-time behavior like prompt injection or unsafe outputs, and face friction from tools that require code changes, SDKs, or proxies.

At the same time, according to Gartner, more than 70% of AI applications are built and deployed in containers running on Kubernetes and cloud native infrastructure. Aqua Secure AI delivers full lifecycle security tailored for this modern stack. It combines deep development visibility with intelligent runtime defense, all without slowing down innovation. Built on Aqua's decade of container security expertise, Secure AI empowers organizations to innovate with GenAI while staying in control.

Key Benefits

- ✔ Gain full visibility into which AI models, platforms, and versions are in use.
- ✔ Govern AI behavior at runtime without adding agents or changing code.
- ✔ Detect and stop threats like prompt injection and jailbreak attempts as they happen.
- ✔ Enforce AI usage policies based on OWASP Top 10 for LLMs.
- ✔ Shift left to detect unsafe AI patterns and logic flaws before deployment.
- ✔ Make faster, more informed decisions with complete AI risk context across code, cloud and prompt.

Prevent AI Attacks in Real Time

Stop advanced AI threats the moment they appear using Aqua's advanced behavioral detection. Stay ahead of runtime risk without needing to change your application architecture.

Stop Prompt Based Threats at Runtime

Monitor prompt injection, jailbreak attempts, and other malicious inputs using eBPF-powered runtime detection, with no SDKs or code changes required.

Identify Unsafe or Rogue Model Behavior

Observe runtime activity to catch unauthorized model access, data leakage attempts, or anomalous inference behavior, enabling quick incident response.

See and Govern Everything

Gain a single view of AI risk across your environments based on OWASP Top 10 for LLM.
Aqua maps LLM use, model types, and platform activity, giving you clear insights to drive governance and control.

Protect AI Applications Across Environments

Secure your AI workloads wherever they operate in containers, Kubernetes, cloud platforms, or even mainframes using the same consistent runtime protection.

Integrate AI Security Early in the SDLC

Catch AI risks early by scanning code for insecure inputs and outputs, and misconfigured LLM use. Prevent vulnerabilities before they reach production, improving security while maintaining developer speed.

Detect AI in Code and Pipelines

Identify embedded LLMs and AI logic across codebases and CI or CD workflows.
Surface undocumented use and risky configurations before they create downstream exposure.

Validate Input and Output Logic

Analyze how prompts are used and how model outputs are handled to prevent injection attacks and response manipulation.

Secure Third-Party AI Services

Evaluate cloud AI services like OpenAI and AWS Bedrock with Aqua AI SPM. Ensure configurations are secure and monitored to reduce supply chain exposure.

Apply GenAI Assurance Policies

Define and enforce policies aligned to OWASP Top 10 for LLMs.
Apply them throughout development to catch risky AI usage patterns before code is committed.

How It Works

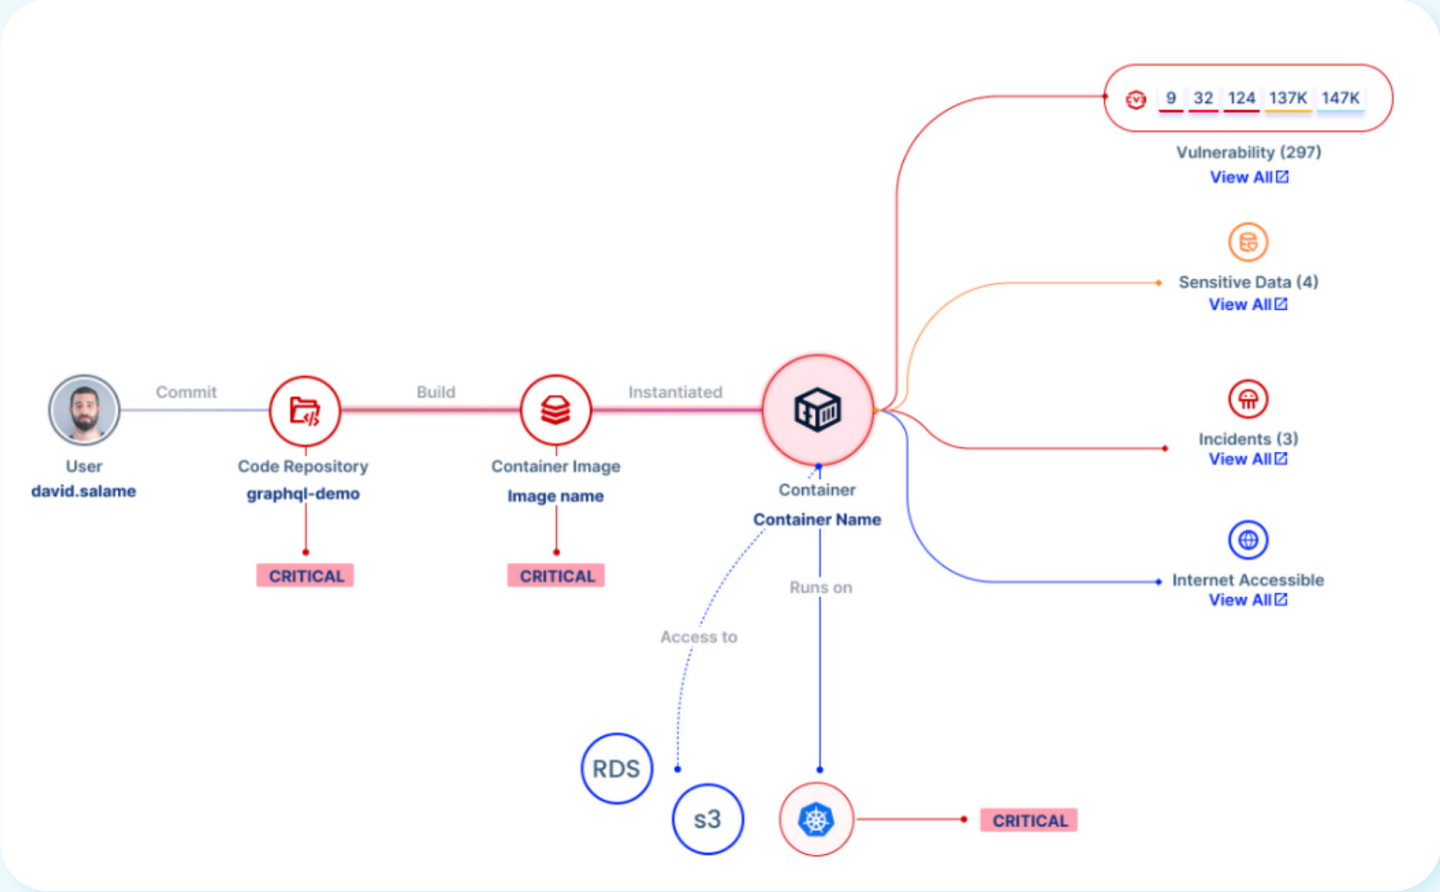

1

Comprehensively scan source code for unsafe AI usage

2

Implement strict guardrails that enforce AI security best practices

3

Stop attacks with real-time threat detection and response

The Aqua Platform

Securing Every Application, Including AI

Aqua's Platform now includes Secure AI, a purpose-built solution for protecting AI applications throughout their lifecycle. As AI becomes containerized and integrated into cloud native stacks, Aqua extends its proven runtime, code scanning, and policy enforcement capabilities to address new GenAI threats. From detecting unsafe prompt logic in development to blocking live attacks in production, Aqua helps teams innovate securely with confidence that every AI workload is protected from code to cloud to prompt.

Gain visibility, enforce policy, and stop AI threats in real time.

Get a personalized walkthrough of Aqua Secure AI and see how it fits into your AI strategy.

Schedule a demo >

aqua

Aqua Security is the pioneer in securing containerized cloud native applications. The Aqua Platform, a Cloud Native Application Protection Platform (CNAPP), enables organizations to secure every cloud native application everywhere, from code commit to runtime. With enterprise scale that doesn't slow development pipelines, Aqua secures your future in the cloud. Founded in 2015, Aqua is headquartered in Boston, MA and Ramat Gan, IL protecting over 500 of the world's largest enterprises. For more information, visit https://www.aquasec.com