SOLUTION BRIEF

Aqua Platform

Protect What Runs in the Cloud

Key Benefits

- ✔ Reduce real production risk by controlling workload behavior at execution time.
- ✔ Contain threats immediately without waiting for remediation cycles.
- ✔ Maintain consistent protection across hybrid and multi cloud environments.
- ✔ Enforce workload behavior at execution time.
- ✔ Strengthen investigations with preserved runtime evidence.
- ✔ Provide clear, defensible insight into security effectiveness.

The Shift from Prevention to Protection

Security teams can no longer rely on visibility and remediation workflows to control risk. Vulnerabilities will reach production and exploitation now happens faster than teams can respond, especially as AI can analyze large codebases, surface deep logic flaws and generate exploit paths with minimal human involvement. The time required to move from discovery to exploitation becomes extremely short, and if action requires a ticket, the window to stop the attack has already passed. The backlog of vulnerabilities does not go away, but it becomes less relevant because attackers can act on it faster than teams can resolve it.

Aqua brings security directly into runtime where applications execute and attacks occur. It enables organizations to detect, understand and contain threats immediately without waiting for remediation. By combining deep runtime telemetry, real time enforcement and agent driven response, Aqua allows teams to move from identifying risk to actively controlling it in production environments. Rather than adding more findings or earlier alerts, Aqua focuses on controlling behavior in production and reducing the impact of exploitation in real time.

Protect Workloads Where Risk Becomes Real

Enforce runtime policy inside production workloads so unsafe behavior is contained before it disrupts critical applications.

Preserve workload immutability with enforced runtime policy

Prevent unauthorized changes by allowing only expected processes and executables to run, keeping workloads safe and harder to exploit in production.

Maintain consistent enforcement across environments

Apply the same controls across containers, Kubernetes, virtual machines, and serverless environments, regardless of where workloads run across hybrid, multi-cloud, and air-gapped environments.

Reduce the runtime attack surface

Limit lateral movement and privilege escalation by restricting behavior inside and between workloads at execution time, containing attacks and preventing spread across the environment.

Stop malware without disrupting applications

Enhance runtime security by automatically alerting, blocking or deleting advanced malware upon download or execution based on the threat severity and business risk.

Act at Machine Speed with Autonomous Enforcement

When an attack occurs, there is no time for tickets or escalation paths, with Aqua decisions are made directly at the point of execution.

Stop threats instantly, not after analysis

Replace sensor-based visibility and delayed enrichment with in-workload enforcement and compensating controls that block unsafe behavior the moment it happens.

Maintain protection as environments change in real time

Continue enforcing decisions as workloads scale, restart, and move across clusters without requiring revalidation or reconfiguration.

Extend security with custom agents

Allow teams to build AI agents that use Aqua's runtime intelligence and enforcement to investigate, decide, and act within their own workflows and environments.

Enable agent driven response

Move from alert to action in seconds with Compass, an embedded MCP server that allows teams to analyze incidents and recommend actions with human oversight.

Focus Remediation on Exposure That Matters

Prioritize vulnerabilities based on their impact on running workloads so security teams invest effort where it lowers real risk.

Evaluate vulnerabilities in live environments

Identify which packages are active, which code paths are reachable and which services are exposed, allowing teams to distinguish exploitable weaknesses from false positive noise.

Prevent high-risk artifacts from reaching production

Set up assurance policies to define risk level for accepting artifacts for deployment, proactively preventing vulnerabilities from reaching production.

Connect findings to accountable owners

Trace vulnerabilities from running workloads back to the specific image and service owner responsible so developers can act quickly with clear, production-level context.

Contain risk when remediation must wait

Reduce exploitability when immediate patching is not feasible by applying compensating controls, such as a virtual patch for immediate protection.

Investigate Incidents with Reliable Evidence

Provide security teams with factual runtime insight so investigations are grounded in real attack patterns and behavioral context.

Detect known and novel threats through real workload behavior

Identify attacks by observing process, file, network, memory and AI execution behavior inside running workloads, not just configuration data or pre-deployment signals.

Preserve memory level evidence at the moment of attack

Capture critical data from container memory before workloads terminate, so investigations reflect the actual state of the system at the time of compromise.

Gain deep runtime visibility at the kernel level

See exactly what is happening in production with granular insight into every action executed inside the workload, enabling precise and reliable threat detection.

Reduce alert noise with high-fidelity detection

Surface precise runtime signals that indicate real compromise, eliminating noise from theoretical, unreachable or low impact risk.

Demonstrate Security Effectiveness with Production Data

Translate runtime enforcement into measurable insight that supports governance, compliance, and executive reporting.

Surface risk exposure in business context

Quantify exposure in financial terms by estimating potential loss based on exploitability, likelihood of attack, and workload impact.

Support audit and compliance efforts

Deliver audit-ready reports with verified evidence of runtime policy enforcement, reducing manual data collection and preparation.

Measure reduction in exploitable risk

Provide reports that show how many meaningful exposures have been resolved over time and where material production risk still exists.

Integrate data into existing systems

Extend reporting beyond the dashboard by exporting structured runtime data into SIEM, GRC and internal reporting platforms through APIs and direct query access.

Secure AI Applications at Runtime

Apply the same production-level discipline to AI-driven workloads so innovation does not outpace control.

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.

See and Govern Everything

Gain a single view of AI risk across your environments, mapping LLM use, model types and platform activity, to OWASP Top 10 for LLMs to drive AI governance and control.

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.

Protect AI Applications Across the Lifecycle

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

Autonomous Runtime Security for Modern Cloud Environments

The Aqua Platform is built to control what happens in production. By operating inside workloads and enforcing policy at execution time, it allows organizations to move beyond visibility and take direct control of risk. Security teams no longer have to rely on perfect remediation or perfect prevention. They can detect, understand, and contain threats as they occur, reducing real business risk at the moment it matters most.

Schedule a demo ›