Why and How to Add Runtime Defense to Your Cloud Security Strategy

Today, cloud native technologies like containers and serverless functions are not exactly new ideas. They have been around in their modern forms for most of the past decade. Docker dates back to 2013, for instance, and AWS Lambda serverless functions appeared in 2014.

What is changing rapidly now is the rate at which cloud native technologies are entering into production use. A March 2020 report found, for instance, that 29 percent of organizations were running more than half of their production applications inside containers—an increase of about 31 percent from just six months prior. Gartner predicts that 70 percent of organizations will be running at least three applications in containers by 2023. And a Forrester report found that more than four-fifths of CIOs currently include expanded use of containers and other cloud native technologies on their lists of priorities.

This rapid adoption of cloud native technologies means that DevSecOps tools and processes are now becoming an integral part of the development process. In turn, teams are implementing additional layers of security to help protect cloud native workloads at all stages of the CI/CD delivery chain.

Chief among those additional approaches is runtime security. By automating security for fast-moving, dynamic applications like those that run in containers, runtime security addresses the unique security and compliance needs of cloud native environments and allows teams to apply DevSecOps processes fully to their cloud native workloads.

What Is Runtime Security?

Runtime security refers to end-to-end monitoring and validation of all activity within containers, hosts, and serverless functions.

Runtime security works by leveraging application control and allow listing to establish a baseline of normal behavior for each host, container, serverless function, or other objects within a cloud native environment. Then, through real-time observation of file systems, processes, and network activity, runtime security tools detect suspicious or anomalous activity, and alerts teams as needed.

Real-time security monitoring is not a new idea, of course. Security Information and Event Management (SIEM) platforms have been capable of monitoring application environments for anomalies for years.

What's different about cloud native runtime security, however, is that it applies to fast-moving environments where baselines don't exist in the traditional sense. When container and serverless function instances constantly spin up and down, and load balancers continuously redirect traffic between different instances, conventional data sources like logs and network traffic are not enough on their own to detect anomalies that could signal a security breach. Runtime security goes deeper by interpreting how behavioral trends vary over time to establish a dynamic baseline. From there, runtime security tools can detect changes in internal container processes, file system activity, and so on, that deviate from the norm—even within environments that rapidly scale up and down.

Put another way, runtime defense is the set of features that provide both predictive and threat-based active protection for rapidly changing environments.

Why Do You Need Runtime Security?

By enabling anomaly detection in dynamic environments, runtime security offers a variety of advantages that traditional monitoring tools can't provide in cloud native environments.

Detect threats in dynamic environments

The first and most important advantage is that runtime security is the only way to secure cloud native applications at scale. As noted above, establishing a baseline of "normal" behavior, then identifying deviations from it, is difficult to do in fast-moving, distributed environments—especially when those environments involve hundreds of container and serverless functions, and dozens of microservices.

Runtime security uses AI and machine learning to automate the process of modeling healthy activity and detecting deviations from it, even in complex environments where it would be very difficult for human engineers to identify normal activity manually.

Mitigate the impact of a breach, or prevent it in the first place

In addition, runtime security helps harden environments against vulnerabilities. By delivering control over file systems, processes, and network activity for each container and serverless function, runtime defenses mitigate the damage that can result if a security breach does occur within your environment.

On top of this, runtime defense tools can automatically model application-safe behavior and enforce rules that prevent dangerous activity on the container or host. They protect against situations such as a compromised container executing processes that spread the breach to other containers or the host, for example.

Enabling incident response

The data collected by runtime security tools plays a role in incident response as well. By capturing and storing audit data for cloud native applications, runtime security tools provide the information teams need to understand what went wrong following an incident, even if the cloud native environment no longer exists in its earlier form when the investigation occurs.

Challenges in Implementing Runtime Security

Runtime security is a complex concept. It can be difficult to implement it efficiently and effectively for several reasons. The challenges that follow can be overcome, but they require runtime security tools that are designed for modern, highly complex environments.

Constantly changing environments

Again, in cloud native environments where configurations and scale change constantly, establishing a meaningful baseline for normal activity is challenging. It requires factoring in a large number of variables, such as how traffic patterns change between times of day and days of the week, and how container orchestration tools react to shifts in application demand and pod failure.

Diverse technologies

Cloud native environments consist of a range of different types of technologies: containers, serverless functions, virtual machines, cloud services, and more. Runtime security tools must be able to interpret the architectural traits and behavioral patterns of each of these components to model and enforce safe behavior.

Risk of inaccurate alerts

It's easy to send alerts. It's harder to ensure that each alert is meaningful, that alerting rules don't "rot" as environments change, and manual alert configurations that previously delivered valuable alarms don't result in false positives or false negatives.

What to Expect from Runtime Security

When properly implemented, a runtime security solution should deliver several core features:

Automatic application behavior modeling

Runtime security tools must be able to automatically understand what secure application behavior looks like. When you have dozens of services to manage and hundreds of containers, serverless functions, and VMs hosting them, you don't have time to configure behavior models manually, or even to collect behavioral data by hand. You need tools that can do it all automatically.

Control application behavior

In addition to modeling safe behavior, runtime defenses should be capable of automatically defining allowed and disallowed behavior for each container, serverless function, or another object in the environment. For example, they should be able to define another container as a container they can communicate with, or which data storage volumes it can access. Enforcing these rules is critical to ensuring that if a breach does occur, its fallout will be limited.

Send meaningful alerts

Although runtime security tools can automatically implement some defenses, they must also be capable of alerting your team when interventions are necessary. To do this, they must monitor and alert for suspicious changes to resources like processes, network connections, or file system read/writes within cloud native infrastructure.

They must also be able to decide whether or not to send an alert based on dynamic alert rules. Activity that might signal a security threat at one moment could be benign the next, which means static alerting rules are not capable of keeping up with cloud native threats.

Integrate with other security solutions

Runtime security represents only one of the layers of defense that should exist within your organization's security stack. Automated data security protections, access control and auditing tools, container image scanners, and so on, are equally important.

To deliver the greatest value, runtime security solutions must be able to integrate with other security tools to provide full depth and context for incidents, as well as understand how the impact of a threat at one layer of your technology stack (like the runtime environment) impacts another (like data at rest).

Getting the Most Out of Runtime Security

Runtime security tools offer a broad set of features. The value you leverage from them depends in part on how you approach runtime defense. To get the most out of runtime security solutions, strive for the outcomes that follow.

End-to-end coverage

Monitoring only part of your environment, or focusing on only key services or infrastructure, is not enough to guarantee detection of all security threats. For best results, apply runtime security to all layers of your environment, and use it to protect both dev/test and production workloads.

Detect incidents in real time

Although runtime security is capable of mitigating the impact of a breach after it occurs, your runtime solutions will ideally allow you to find and remediate threats in real time, before they have an opportunity to escalate.

Treat each resource uniquely

Every host, container instance, serverless function, etc., in your cloud native environment, has a unique configuration and behavior. You should therefore model each one separately. Don't assume that all containers will behave the same even if they are based on a common container image, for example. Doing so leads to a sampled approach that limits visibility into security incidents.

Runtime Security with Prisma® Cloud

When you move to cloud native technology, runtime security becomes a critical part of your security strategy. You can't make full use of solutions like containers, serverless functions, and cloud services without runtime defenses to protect your investment.

Prisma Cloud, a comprehensive cloud native security platform, offers runtime security solutions, as well as other defenses—such as data protection, access control management, and cloud security posture management—that organizations need to keep their cloud native environments safe, no matter how complex those environments are or how rapidly they scale. Learn more by requesting a demo.

paloalto® NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com