

Cloud Workload Security: What You Need to Know

Part 1

Cloud proliferation is on the rise, and more than ever before, security teams are on the lookout for solutions that align with evolving cloud security paradigms. Given the evolving threat landscapes and more sophisticated cyber attacks being reported daily, it's clear that your security teams need to have a well-defined strategy with the right controls in place to protect the mission-critical workloads you've deployed in the cloud.

Since the controls and categories you take into account when developing your cloud security strategy can make or break your cloud deployments, as well as the security of your workloads, we've created a blog series to help you understand and identify these controls and provide guidelines for implementing them on the cloud platform of your choice.

In subsequent parts of this series, we'll get you up to speed on the security tools and services delivered by the leading cloud platforms—**Azure**, **AWS**, and **GCP**—as well as which controls they support. We'll also explore the maturity of each platform's security capabilities and conclude the series with a full comparison of their services.

But here in Part 1, we take a deep-dive into the following areas: What you need to focus on when developing your cloud security strategy; what security controls you should consider, along with the relevance of each; and finally, what's the best approach for implementing these controls.

Security in the Cloud: What Should You Focus On?

Organizations usually find it overwhelming to distill down cloud security controls while deploying workloads in the cloud. It's equally important to maximize security while being on the lookout for possible attacks, plus be ready with mitigation plans in the event of an attack. It has to be a well-balanced effort—focusing on one area while ignoring the other leaves your cloud workloads vulnerable and exposed.

Let's explore the main focus areas you need to take into account in order to develop a mature cloud security strategy with maximum coverage.

Reduce the Attack Surface

Understanding the attack surface and making efforts to decrease its size will ensure cloud security, meaning you will effectively decrease the likelihood of an attack. Controls you can implement to achieve this include segmentation of connected networks, patch management, runtime vulnerability management, and container image scanning. You should also build in the necessary security measures early in your application's lifecycle and align them with DevSecOps practices.

Note that continuous monitoring and optimization is crucial to aligning your security strategy for ongoing protection. This is not just the responsibility of your security team—it's on everyone who accesses and uses the environment, i.e., developers, DevOps engineers, the monitoring team, etc. Even with ongoing efforts, it may not be practically possible to eliminate the attack surface completely. So, a more pragmatic approach is to minimize it as much as you can by focusing on breach detection and response.

Detect Attacks/Breaches

The longer it takes to detect a breach, the greater the damage will be; so a successful cloud security strategy needs to minimize the time from an initial attack to detection.

It should also be noted that isolated events that don't look suspicious could potentially point to undetected attacks when correlated with other events in your environment. For example, increased east-west traffic when correlated with a new process running in a compute environment could be related to an attacker trying to gain lateral access to components in a network from a compromised VM/container.

For comprehensive detection, you should make use of services that work at the network and compute-resource levels. Also, make sure to update your threat/attack baseline databases with the [latest threat information](#), as threats are evolving on a daily basis.

Respond to Attacks

The last piece of the puzzle is to minimize the time from detection to full recovery through an effective threat response strategy. You also need to investigate and differentiate between real threats and

false alarms. This involves an in-depth analysis of the attack vector information captured by your monitoring tools and logs. Then, there are auto-remediation features, like webhooks and automation scripts, integrated with alerts to help you create a first line of defense.

Remember, what you learn from an attack and the subsequent mitigation must be fed back in to reinforce your security strategy through adaptive tuning.

Cloud Security Controls and Categories

There are other elements, such as identity and access management, data protection, and application security, that are ingrained in the three strategies that we discussed in the previous section. However, these controls are fairly standard or too application-specific, so we will not be focusing on them in this blog.

Instead, let's move on to some practical implementation methods for the cloud security strategies discussed so far.

This process can be mapped to different security controls and categories, some of the most prominent being: network perimeter of the workloads, managing the configurations per best practices, runtime protection mechanisms, choosing the right CWPP solution, and integrating reliable SIEM solutions.

We will elaborate on these categories and controls here below.

Network

Segmentation of Your Network

Compared to traditional on-premises networks, cloud networks are based on SDN and offer a great level of flexibility to enforce microsegmentation. Steps such as basic isolation using different virtual networks, policy-based segmentation of container workloads, and segmentation of workload tiers help in restricting east-west traffic, they also create clear boundaries to apply focused security policies.

For example, boundaries and policies at the network layer could help protect against attacks from unknown IPs or those targeting exposed known ports. However, more sophisticated attacks happening at the application layer require the segmentation of application components and policies targeting those boundaries.

Web Application Firewalls

Applications hosted in the cloud are more susceptible to attacks and so need advanced protection layers configured. Web application firewalls offered by cloud service providers can be considered here, as they provide comprehensive protection at the application layer against known vulnerabilities and exploits. WAFs based on attack detection rules from [OWASP](#) can help detect and protect against common application security risks like Injection (SQL, NoSQL, OS, LDAP), sensitive data exposure, security misconfiguration, XSS flaws, and broken authentication and access control.

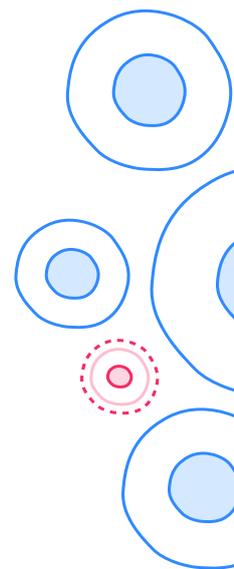
DDoS Protection

Organized DDoS attacks can bring down your cloud applications in a matter of seconds. These attacks could target the cloud platform or specific workloads hosted on the platform. So, a multi-layer DDoS protection approach is inevitable and should start from the platform and provide protection all the way up to the application layer. All leading cloud service providers have native capabilities built-in to protect against platform attacks as well as take any necessary mitigation measures.

In addition to this, you should also consider configurable DDoS protection services focused on specific IaaS/PaaS services. Desirable characteristics of a service include configurable mitigation policies and automated mitigation services, as well as rich analytics services that offer insights into the attack.

Cloud Security Posture Management (CSPM)

Service misconfigurations, manual errors, or mismanagement are often the root cause of successful attacks on cloud workloads. CSPM solutions continuously monitor cloud deployments for misconfigurations and report any deviations so that the necessary remedial measures can be implemented. This control-plane view, often quantified using security scores for cloud environments, provides a high-level overview of the cloud environment's security posture.



In addition to default security policies aligned with leading compliance standards, an effective CSPM solution should also enable organizations to define custom security policies specific to their industry vector or business requirements. CSPM solutions can report deviations from an organization's security policies as well as recommend remediation steps or even implement them. The scope of a CSPM solution should cover the entire cloud service landscape, including VMs, network, storage, and PaaS services, and also extend to container and serverless environments.

The capability of continuous risk and compliance assessment and reporting plays an important role in reducing the attack surface, making CSPM an integral part of your cloud security strategy.

Vulnerability Management

A real-time vulnerability scanning and remediation mechanism is as important in the cloud as it is on-premises. Considering the growth of microservices in the cloud, vulnerability detection and management tools should also be able to protect containerized environments through options like container runtime scanning, integration with CI/CD pipelines, etc. Ideally, they should be able to continuously analyze workloads for vulnerabilities, generate reports, and display results in dashboards, as well as auto-remediate vulnerabilities whenever possible.

A systematic patch management process for both Windows and Linux machines helps to reduce an operating system's attack surface area. Organizations can also leverage the built-in patch management solutions available on cloud platforms or third-party integrations to achieve this. Overarching reports that cover reported vulnerabilities and the mitigation steps completed, as well as patch cycles executed, will help in maintaining a historic view of security compliance, which is necessary to meet the compliance standards of most audits.

Cloud Workload Protection Platform (CWPP)

With threat actors shifting to the cloud and adopting [sophisticated](#) attack methods, the focus of security must shift from rigid parameters used by on-premises systems to a more evolved workload-centric approach. The approach for detecting and reporting breaches needs to be holistic, and that is where Cloud Workload Protection Platforms (CWPP) help with their workload-centric security approach. CWPP solutions are designed to monitor the compute resources in your deployments, i.e., VMs, containers, and functions, with the help of agents to provide comprehensive insights into your security posture.

Applications can span across hybrid or multi-cloud environments, and CWPP solutions offer single-pane visibility and protection for them. CWPP capabilities include, but are not limited to, breach/threat detection, system integrity assurance, hardening of services, Application Control, and in-memory protection. Signature-based anti-malware scanning programs could be an outdated approach for cloud workloads especially those hosted in Linux; thus, CWPP should additionally provide advanced threat detection methods and built-in application-control features. A CWPP threat detection strategy goes one step beyond the traditional approach and ensures that all the code running in your compute resources are from a trusted source. Any unauthorized or malicious code execution will be identified by the threat detection capabilities of the tool.

Traditionally, Linux was considered a "secure by default" operating system. However, threats like [Doki](#) and [IPStorm](#) targeting Linux and cloud environments, along with undetected security loopholes (including vulnerable plugins and unpatched software), have made Linux systems equally vulnerable in the cloud. While traditional security solutions are mostly focused on detecting threats in Windows, CWPP solutions should specialize in security management and threat detection for Linux to bring definitive added value.

Container Security

Container security includes the protection of containers as well as orchestration platforms, the most popular one in the cloud being [Kubernetes](#). While most of the cloud service providers offer a managed Kubernetes service, customers also choose to deploy their own Kubernetes clusters if they need access to the control plane.

Industry-standard security baselines should be defined for the Kubernetes clusters as well as the containers, with continuous monitoring against those standards and reporting of any deviations. Any malicious activities at the container or host level, like privileged container access, API access from suspicious sources, and web shell detection, should be reported in real time as well as analyzed for underlying security flaws. This capability is often included as part of a CWPP solution.

The container-image scanning feature also helps reduce the attack surface by identifying and flagging vulnerabilities in the image before it is pushed to a container registry. To achieve this, customers can leverage native container tools provided by cloud service providers for security baseline monitoring and image scanning. When native tools are not available, third-party tools can be leveraged to do the same.

Security Information Event Management (SIEM)

Cloud-native SIEM helps to collect relevant logs and signals at cloud scale from multiple sources, both in the cloud and on-premises. They can then rapidly analyze and detect threats. SIEM tools help to correlate data from different components of the architecture, irrespective of where they are deployed; this helps with enhanced threat detection and response.

These tools should ideally offer native integration with multiple data sources and provide API-based automation for remediation activities. Most SIEM tools also offer rich visualization capabilities that help with easier reporting and visibility into the frequency of malicious events, data anomalies, network infiltrations, etc.

Additional Threat Detection Capabilities

In addition to the security coverage offered by CWPP solutions, organizations should also consider services from their cloud service provider that offer advanced threat detection and protection capabilities. The scope of service here should cover the different layers of the application, compute resources, and data resources, as well as the supporting service layer, including the cloud control plane, network traffic, and key management solutions. Some examples of scenarios to be covered include identifying suspicious logins and activities, monitoring the administrator activity log, and reviewing usage patterns to detect rogue users.

Available Security Logs and Monitoring

All leading cloud providers offer comprehensive logging options for relevant service signals. The logs of both the control plane and the data plane should be analyzed and monitored to ensure end-to-end security. These logs, which would include service activity logs, network flows, IAM logs, data ingress/egress logs, etc., will be a critical factor when conducting investigations.

Summary

Assuring cloud security is not a destination but rather an ongoing journey that requires continuous optimization as your workloads mature. The controls discussed in this blog will provide the baselines you need to define your initial cloud security strategy and accelerate the process of securing cloud workloads. Some additional references and recommended reading materials are listed below:

[Cloud Security Posture Management: Why You Need It Now](#)

[Posture Management: Cloud Security Tools Rise in Wake of Breaches](#)

[Is the Cloud Secure?](#)

AWS, Azure, and GCP are the leading cloud platforms on the market and offer a number of native solutions and services that can help you to implement security controls. Advanced SIEM capabilities or threat detection features, such as in-memory detection or runtime scans, might still require that you use third-party solutions alongside native features. We will explore such tools in detail in the upcoming posts of this series, discussing these cloud service providers and giving a comprehensive comparison of their features.

