Guide to Cloud Security Posture Management Tools

paloalto® NETWORKS | PRISMA® BY PALO ALTO NETWORKS

Table of Contents

Good cloud security hygiene starts with complete visibility into the security and compliance posture of every resource you deploy into your cloud. It's one thing to achieve this visibility in a single cloud environment—you can lean heavily on the native monitoring and auditing tools of your cloud provider, using third-party solutions to fill in gaps (e.g., threat detection)—but in a multi-cloud architecture, maintaining robust cloud security posture becomes exponentially more complex.

It is much more difficult to achieve centralized visibility as well as consistently enforce policies and compliance rules within a multi-cloud environment. It's also more complicated to detect threats and fix vulnerabilities quickly due to the sheer complexity of threats across distributed, multilayered architectures.

You can address these challenges, though—and you need to, if you want to take advantage of multi-cloud architecture without compromising on security. This guide will walk you through the specific challenges of cloud security posture management (CSPM) within a multi-cloud architecture. Then, we'll discuss how to build a CSPM tool set and strategy that effectively address those challenges by providing centralized visibility, compliance management, threat detection, data protection, and automation purpose-built for multi-cloud environments.

The Unique Challenges of Multi-Cloud CSPM

CSPM that works in a single-cloud environment can't simply be scaled up to meet the needs of a multi-cloud architecture. From a security perspective, multi-cloud environments fundamentally differ from single clouds in a variety of ways.

Disparate, Distributed Data

Data in a multi-cloud environment is spread across multiple clouds. You may store high-volume data in one cloud to take advantage of lower pricing, for example, while keeping other data in a different cloud that charges more but offers faster access to the data. In another case, you might distribute data between different cloud regions to place it closest to end users who will access it.

When data is distributed, ensuring it's secure and malware-free is more challenging. You need to be certain that proper identity and access management (IAM) rules are in place for each data store on each cloud. Moreover, you also need to ensure each bucket is properly configured, and what is "proper" differs across cloud service providers (CSPs).

Distributed Applications

Often, applications and the tools that deliver them are also distributed across multi-cloud environments. For instance, you may run duplicate instances of the same application in different clouds so that it will remain available if one cloud fails, or you may host a development toolchain in one cloud but deploy from it into another.

Effective CSPM in this context requires understanding the security posture of all the individual services and resources that compose the application. It’s not enough to monitor each instance of a multi-cloud application separately. You must know how the security state of one instance could impact others. Can instances interact in ways that could allow a breach in one cloud environment to escalate into another, for example? Continuously monitoring application configurations to ensure they don’t deviate from the policy guardrails in place helps protect against such risks.

At the same time, because deploying applications across multiple clouds is more complex, it becomes even more important to integrate security into the application development pipeline, rather than tacking it on as an afterthought. It takes much more time to address a vulnerability once the application reaches production—even more if the vulnerable code is deployed across multiple environments. By integrating security controls into the application development pipeline, you minimize the risk of having to perform these types of reactive, corrective actions once an application is already in production.

Multiple Threat Types and Vulnerabilities

Modern threats come in many forms, ranging from malicious insiders misusing APIs to cryptojacking, data exfiltration, malware inside a container image, SQL injection vulnerabilities within applications, and more. No single type of threat detection can guard against them all. Malware scanning or configuration auditing alone won’t protect your environments. Instead, you must gather threat intelligence from a variety of sources, analyze it, and map the results to known threats. You must also be able to augment rule-based policies with machine learning-based policies to detect unknown threats.

Multiple Users and Multiple Permissions

Single cloud environments already challenge security teams to enforce good IAM hygiene. With numerous different policies (e.g., CSP- or user-managed policies; policies attached to other groups, roles, resources, or access control lists) attached to users, enforcing least-privileged access for a single cloud is hard enough.

Multi-cloud environments further complicate this as user permissions and entitlements are inconsistently defined across CSPs. You must not only monitor a different set of IAM configurations for each cloud, but also be able to correlate user roles and permissions defined for each cloud with each user's requirements. You can't simply audit for the same set of credentials for the same users on every cloud.

Broader Attack Surface

More clouds mean more accounts, access control policies, services, and so on. All of this adds up to a broader attack surface, creating more potential opportunities for attackers to take advantage of a misconfigured resource, lax permissions, or a code vulnerability to break into your environment.

At the same time, the lack of centralized visibility and control into a multi-cloud environment makes vulnerabilities and threats more difficult to detect. When you can't monitor and audit all your configurations across all your services using a single CSP's tools, it's harder to prevent the types of misconfigurations that can lead to a breach.

Mastering Multi-Cloud CSPM: Four Key Features

Addressing the aforementioned multi-cloud security challenges requires a CSPM strategy and tool set that provide four key features.

Complete Visibility, Compliance, and Governance

The foundation of successful CSPM for multiple clouds is the ability to continuously monitor and audit all resources across all CSPs. Whenever a new service or workload is deployed or a configuration is changed, your tools must be able to detect and scan the update to ensure it complies with security requirements and best practices.

If it doesn't, the tools should both alert your team that something is wrong and recommend ways to fix it. Your tools should be able to apply simple fixes (e.g., updating a mistyped IP address, adding a missing statement to an IAM policy) automatically, remediating such issues without waiting on security teams to make the change.

Preventing insecure configurations from reaching production in the first place is also important because it can reduce the number of runtime alerts generated. CSPM tools should be able to scan infrastructure-as-code (IaC) templates for misconfigurations and enforce policies not just at runtime, but also across the software development lifecycle.

Comprehensive Threat Detection

The complex nature of threats in multi-cloud environments means CSPM tool sets need to collect threat intelligence from a variety of sources to gain accurate risk clarity. Those sources include IaC configurations, container images, and cloud virtual machine (VM) images, which many teams already scan for vulnerabilities.

Simply scanning these components, however, is insufficient to deliver full threat intelligence and detection. For that, your organization must also maintain high-fidelity threat intelligence so you can identify the latest threats and assess their severity level. The ability to detect anomalies on the network and correlate them with other types of threat data is important, too, for gaining full context on the potential risk impact of any threat. You must also do the same with user and entity behavior analytics (UEBA) data.

In other words, modern threat detection requires analysis of multiple data sources, combined with the ability to correlate and contextualize that data. This is important not just for identifying threats within complex, multilayered environments, but also for helping teams understand how to quickly prioritize risk and remediate threats. Only through comprehensive threat detection can you associate network anomalies with an insecure container image, for instance, or determine which account is the source of a breach. When your team can understand threats more quickly, you can also fix them more quickly, minimizing your mean time to remediate.

Integrated Data Security

Whatever kinds of data you store in the cloud, and whether or not it includes personally identifiable information (PII), keeping it safe requires a multipronged defense that provides deep visibility into the state and status of your data. That starts with the ability to monitor the configuration of each storage bucket across your various storage services to ensure data is not accidentally exposed to unauthorized users or applications. At the same time, you must audit the contents of buckets to determine whether they contain PII that is subject to special compliance rules or other requirements.

Detecting malware within data at rest is another important yet often overlooked part of cloud data protection. Malware identification requires scanning not just storage buckets, but also databases, VM file systems, container storage volumes, and even short-lived container file systems for signs of malware.

Finally, because cloud data security often requires striking the right balance between protection and availability, your CSPM tools should allow you to calculate the exposure risk of data and make recommendations to help limit the potential impact of a breach. Which level of access control is appropriate for your cloud data based on the sensitivity of the data? Should you use less granular access policies to simplify management? Tools that can calculate exposure risk will help you answer questions like these.

Automation for Alert Remediation

By definition, multi-cloud environments are complex, large-scale environments. Enforcing the aforementioned security processes and oversight within them is impossible without the help of tools that can automatically monitor for and help remediate security risks.

This isn't to say multi-cloud CSPM should be a totally hands-off, automated affair. Manual intervention will always be necessary to respond to complex security incidents or assess risks that are too complicated for your CSPM tools to handle alone. However, routine security monitoring, audits, and remediations should be automated so your team can focus on the big-ticket items.

The Limitations of Cloud Vendor Tools

CSPs provide a variety of tools that can address some of the risks described in this guide. Data protection services like Amazon Macie® and Google Cloud DLP can assess data vulnerabilities in storage buckets and databases, for instance. Monitoring tools such as Amazon CloudWatch and Microsoft's Azure® Monitor can generate alerts for certain types of events that may indicate security risks. These tools are useful, and you may want to take advantage of them to help build out your CSPM tool set. However, CSPs' tools are subject to two main limitations:

1. **They are not designed to be comprehensive, end-to-end CSPM solutions**. They may be able to audit some configuration files or find some PII within certain data stores, but they won't continuously scan container images, detect network anomalies, or automatically remediate threats.
2. **Each CSP's tools work only with that CSP's clouds**. In other words, Google's tools only work with Google Cloud, Amazon's only with Amazon Web Services, Microsoft's only with Azure, and so on.

Relying on CSPs' tools alone in a multi-cloud environment requires juggling a long list of disparate tools—a difficult and inefficient task. It's impossible to correlate data efficiently between all these tools, and when you can't correlate data, you can't perform comprehensive threat detection based on all sources of data available in your multi-cloud environment.

Conclusion

In short, multi-cloud CSPM requires a comprehensive security platform that can continuously monitor for and alert on misconfigurations, vulnerabilities, and threats with high accuracy.

Prisma® Cloud by Palo Alto Networks provides the automation and centralized visibility necessary to effectively address multi-cloud security challenges. By ingesting data from flow logs, configuration logs, and audit logs stored on each of the clouds you run, Prisma Cloud provides a centralized view for security teams to monitor the security and compliance posture of your entire environment.

In addition, because Prisma Cloud provides threat detection based on anomalous activities and impact, it helps your team understand the severity of each threat and take action accordingly. Prisma Cloud means no more time wasted guessing which threats to respond to first, and no more trying to identify the root cause of complex security incidents.

By taking advantage of the automated remediation features of Prisma Cloud, your team can resolve many types of security-related misconfigurations quickly and automatically while monitoring the status of remediation through a central dashboard.

To learn more about how Prisma Cloud can help your team manage the security posture of complex multi-cloud environments, you can watch a demo of the platform in action or book a meeting with a Palo Alto Networks expert.

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