


Cloud Security Blind Spots: Detecting and Fixing Cloud Misconfigurations

At first glance, identifying and fixing security gaps in a cloud architecture may not appear very different from doing the same for on-premises environments. Most of the deployment technologies that run in the cloud, such as virtual machines (VMs) and containers, also run on-premises. In addition, modern management tools typically support cloud-based and on-premises environments equally well.

Yet even if the deployment processes and tooling for cloud environments look basically the same as those of on-premises infrastructure, there are a variety of nuanced differences that can be easy to overlook. If you fail to appreciate and account for them, these misconfigurations can cause security blind spots in your cloud architecture.

Those blind spots can arise in a number of different ways. Some of them result from the ephemeral nature of cloud workloads, which is one of the major distinguishing factors between clouds and on-premises environments. Others are the product of services unique to the cloud, like identity and access management (IAM). Still others have to do with tooling, such as infrastructure-as-code (IaC), which is especially prevalent in large-scale cloud environments, although it may be used for on-premises configurations.

We've designed this guide to help organizations identify the cloud security blind spots they may be overlooking within their cloud environments. It walks through a number of common configuration oversights that affect cloud infrastructures, explains why they can lead to serious vulnerabilities, and offers tips for designing cloud architectures and technology stacks that protect against the risk of misconfiguration.

Because every setup is unique, we can't guarantee that this guide will alert you to every blind spot that may be lurking in your cloud environment. It can, however, offer actionable guidance for designing architectures that are resilient against easily overlooked security threats as well as for detecting misconfigurations within workloads you have already deployed.

What Is a Security Blind Spot?

A security blind spot is any type of configuration, tool, or process that appears to be secure but is subject to potential vulnerabilities that lurk beneath the surface and are easy to miss. In IT, a security blind spot is analogous to locking your front door and feeling safe, forgetting that intruders could easily use the spare key under your doormat to get inside. It's like configuring a home security system but overlooking the fact that it only monitors the first floor of your house, leaving you defenseless against an intruder who climbs in through a second-story window.

Security blind spots are the opposite of well-known, easily recognized configurations. For example, failing to configure IAM policies for your cloud services isn't a security blind spot; it's blatant security neglect. So is storing login credentials in a plaintext file hosted on a public file server. Every administrator worth his or her salt knows not to make mistakes like these.

In contrast, security blind spots are, by definition, misconfigurations that even skilled administrators often fail to see, such as IAM configuration drift across a multi-cloud environment. Avoiding this type of risk requires understanding the oversights that are commonly made when setting up cloud environments or deploying workloads to them.

In addition, you can use auditing to help detect security blind spots after services are already deployed. However, an ounce of prevention equals a pound of cure. Knowing which types of mistakes to avoid in the first place will go far in mitigating your risk of security blind spots.

Cloud Security Blind Spots: Five Common Examples

To illustrate what security blind spots look like in practice, consider the following five examples.

Failure to Understand Shared Responsibilities

Public cloud providers operate according to a so-called shared responsibility model. While the details of such models vary from one cloud service provider to the next, they boil down to the notion that the cloud provider takes responsibility for securing resources that only it can control—such as the underlying servers that host its infrastructure as a service (IaaS)—while cloud end users are responsible for securing workloads that they deploy on top of cloud services.

Shared responsibility sounds straightforward in theory, but it can lead to security blind spots when users make incorrect assumptions about where their cloud service provider's security responsibility stops and theirs begins. As a result, users may fail to secure some facets of their workloads because they mistakenly believe that the cloud provider will handle them.

This is an especially easy mistake to make when using new types of cloud services. When you're working with a classic public cloud service, like VMs or object storage, it's simple enough to delineate shared responsibility: your cloud provider secures the underlying host servers, and you secure the VMs or data that you deploy on top of them.

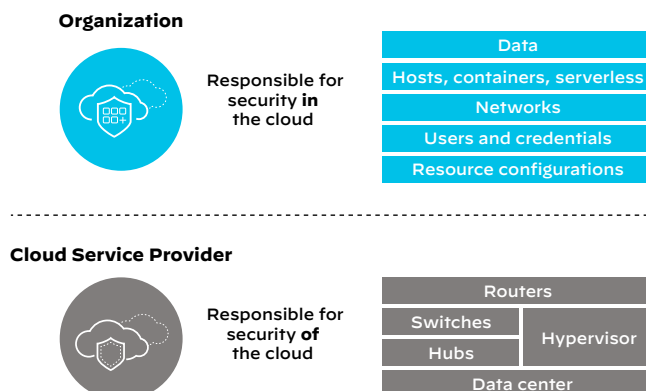


Figure 1: The basic shared responsibility model

Things become trickier, however, when you work with a new, more complex type of cloud service. Consider, for instance, a managed Kubernetes® service that runs in the public cloud. On this type of service, the cloud provider automates the provisioning and, to a certain extent, monitoring of your Kubernetes cluster. It may therefore be easy to assume that the cloud provider manages security, too. The reality is more complex. Although the cloud provider secures the host infrastructure and provisioning tools, it does not address other major security risks that could arise in this context, such as malware inside untrusted container images that you deploy on your cluster or misconfigured access control policies that could invite a breach. If you overlook these sorts of risks and fail to implement your own solutions to address them, you end up with security blind spots in your cloud-based Kubernetes strategy.

Another place where confusion might arise in the context of shared security responsibility is in a cloud-based productivity suite, like Microsoft 365™ or Google Workspace™ (formerly Office 365 and G Suite, respectively). In these software-as-a-service (SaaS) platforms, the cloud service provider secures both the host infrastructure and the applications within the productivity suite. However, responsibility for securing data downloaded from those applications, and managing access to the applications, lies with the end user. Microsoft can't guarantee that weak access control won't allow an unauthorized user to access sensitive documents stored in Microsoft 365, nor can Google ensure that email attachments downloaded from Gmail are free of malware. Responsibility for plugging these security gaps lies with the organizations that use these SaaS platforms.

Sprawl, Shadow IT, and Technical Security Debt

Today, IT teams are under unprecedented pressure to deliver software early and often while embracing a fast-moving “fail forward” culture. This strategy can drive faster innovation and greater agility, but it also increases the risk that the IT team, in its quest for rapid advancement, will make mistakes or oversights that lead to cloud security blind spots.

Specifically, a fast-moving team that relies heavily on the cloud is prone to three types of blind spots:

- **Sprawl:** Cloud services are easy to launch but harder to manage and consolidate in a logical way. Fast-moving teams may find themselves spinning up VMs, databases, and the like haphazardly, with no centralized management strategy in place. By extension, systematically discovering and securing all of those workloads becomes challenging.
- **Shadow IT:** Along similar lines, teams or individuals may launch so-called shadow IT, or workloads only they know about and that are not integrated with central IT management systems. These workloads are also difficult to detect and secure in a centralized fashion.
- **Technical security debt:** To innovate quickly, teams may overlook weak spots in their security postures in ways that increase the time and effort security teams must spend finding and fixing vulnerabilities. For example, a team might set up a container registry without bothering to integrate automated image scanning. This scanning must therefore be done manually, which is inefficient. Until the problem is corrected, this constitutes a form of technical security debt.

Risks such as these are inherent to the fast-moving culture that dominates IT departments today, especially those that leverage the cloud as an easy and highly scalable way to launch workloads quickly. That doesn't mean, however, that you have to accept these risks as a necessary trade-off for fast innovation. By enforcing strong IT governance policies via automated audits, you can find workloads that deviate from your organization's best practices or that are not connected to central management systems.

Infrastructure-as-Code Misconfigurations

To automate the process of provisioning cloud environments, many organizations turn to IaC tools. These solutions allow administrators to describe how an environment should be configured. The tools then apply the configuration automatically, eliminating the need to set up each server, cloud service, and so on by hand.

IaC is great from an efficiency point of view. From a security perspective, however, it can be devastating, because it means that an insecure configuration may be applied automatically across a large number of servers or hosts. Indeed, research detailed in a [2020 Unit 42 Cloud Threat Report](#) found that nearly half of configuration files created for use with CloudFormation, a popular IaC tool, contained misconfigurations (e.g., lack of proper access control) that could lead to security breaches.

This is not to say that you should avoid IaC tools; on the contrary, they are critical for managing large-scale environments. You should audit your IaC configurations carefully, though, making certain to adhere to policies such as [least privilege](#)—a security best practice wherein each user or service account gets the minimum amount of entitlements needed to perform a given task; or [Zero Trust](#)—another best practice that demands each device, application, and microservice be responsible for its own security.

Ephemeral, Fast-Changing Cloud Environments

Modern cloud workloads often rely on deployment architectures that are inherently ephemeral, such as containers. A container instance can be spun up in seconds and then destroyed a few minutes later in response to fluctuations in demand.

This means not only that cloud environments are constantly changing, but also that it is difficult to identify the specific location of a given workload. If you deploy applications on Kubernetes, the orchestrator will automatically spin containers up and down on different nodes depending on which configuration it deems most effective at a given moment. This makes it very difficult to know which individual servers are hosting which applications at any particular point in time. Similarly, network endpoints in a cloud environment constantly change as hosts go up and down, making traffic routes extremely complex to track. From a security perspective, the ever-changing nature of cloud environments, combined with the difficulty of obtaining concrete visibility into their state, means that many traditional security strategies break down. You can't rely on a simple firewall to block malicious traffic if your network configurations are constantly changing. You can't deploy security tools that monitor operating system logs for signs of a breach if your workloads are constantly shifting between different servers.

Instead, you need to deploy cloud native security tools that provide visibility into the dynamic, ephemeral nature of the modern cloud. Relying on conventional security strategies leaves you prone to blind spots because you may believe that you have taken proper steps to secure your cloud workloads when, in fact, those tools are not effective in the cloud.

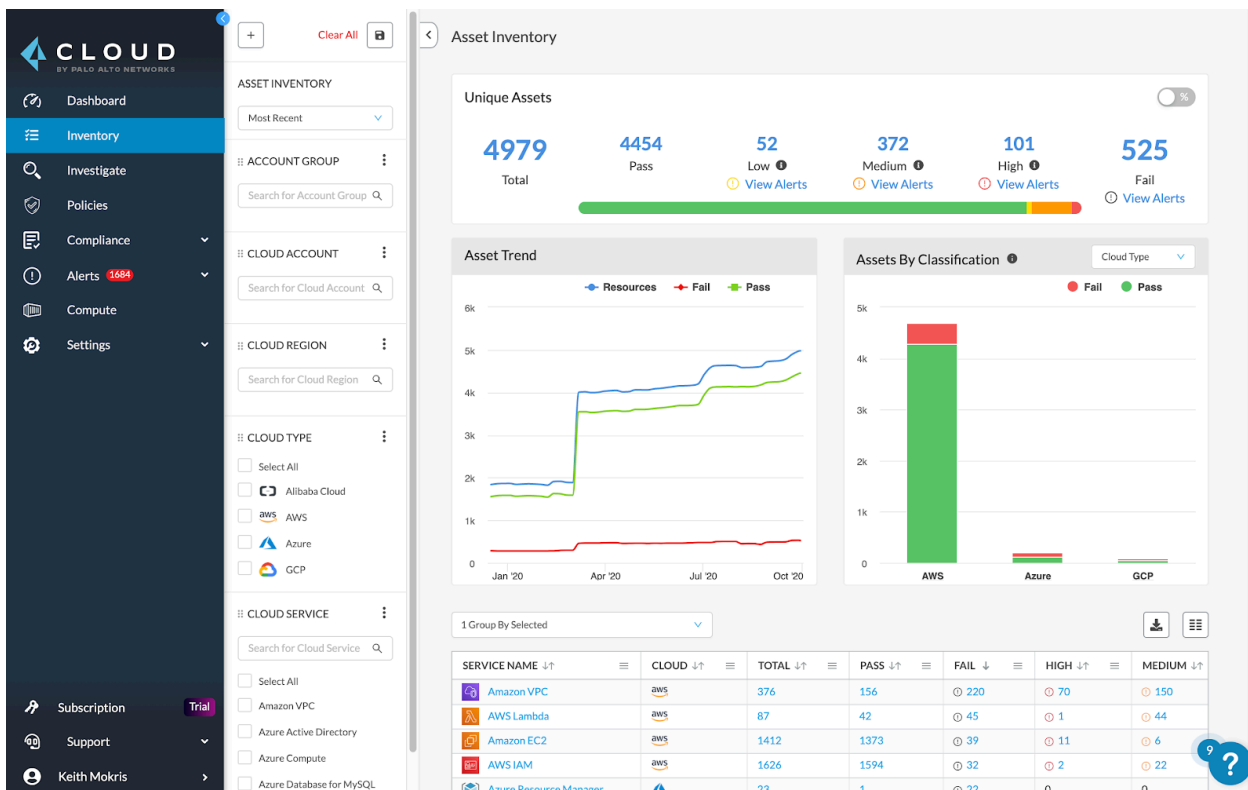


Figure 2: Multi-cloud asset inventory in Prisma Cloud

IAM Mistakes

Most organizations know that using cloud IAM systems to manage access control is a requirement for ensuring basic cloud security. However, IAM policies can be easy to misconfigure, especially in large-scale environments.

In the interest of efficiency, teams may end up over-permissioning by granting users higher levels of access than they need. This may reduce the need to update IAM configurations later if a user's access needs change, but it creates security risks because it ignores the principle of least privilege.

Likewise, an organization might establish IAM policies that apply the same privileges across entire categories of users, giving all employees or devices of the same type the same access rights. This strategy lacks granularity and is a recipe for over-permissioning.

The point here is that, just as setting up a password for a PC is only a first step in securing it, simply having an IAM policy in place is no guarantee that the cloud resources it protects are actually secure. You must audit your IAM policies, ensuring that they are sufficiently restrictive to prevent abuse.

Preventing, Finding, and Addressing Cloud Security Blind Spots

How do you prevent cloud security oversights like those we've described? The answer depends in part on where you are on your cloud journey.

If you have the luxury of building all your cloud workloads from scratch using cloud native technologies, security blind spots are relatively easy to avoid. Unfortunately, few organizations enjoy that freedom. Instead, most find themselves adapting existing workloads to run on the cloud. In this case, you must take extra steps to avoid security blind spots. Consider a few best practices.

Avoid a Blind Lift and Shift

This approach, which refers to taking a workload that runs on-premises and moving it to the cloud with minimal changes, leaves you at a high risk for blind spots. Configurations that are secure in an on-premises environment may be less secure in the cloud. For this reason, careful evaluation of the security implications of a lift-and-shift strategy is critical before you move workloads to the cloud.

Automate Configuration Auditing

Whenever you write a new configuration file of any kind—whether it's an IaC template, an IAM policy, or something else—it should be audited automatically for potential security risks before it is applied. In addition, audits of running workloads should be performed on an ongoing basis to detect risks that may be introduced by changes to a configuration after the workloads have been deployed.

Assume Maximum Responsibility

Just as it's a best practice to enforce least-privileged access for each cloud resource, organizations should strive to assume the maximum level of responsibility for securing cloud workloads. Doing so minimizes the risk of oversights related to misinterpreting your cloud provider's shared responsibility policy. If there is any tool or process you can use to help secure any facet of your cloud services, apply it. It's much better to overstep the shared responsibility line than to fall short of it.

Adopt a Cloud Native Security Platform

At the end of the day, cloud security is a fundamentally different game from on-premises security due to the ephemeral, highly scalable nature of the cloud. Managing security responsibilities in the cloud therefore requires a cloud native security platform (CNSP)—one that is specifically built to address the challenges of cloud native development and that integrates with existing tools and processes.

No matter which cloud architecture or workloads you deploy, [Prisma® Cloud](#) from Palo Alto Networks can provide deep visibility across complex stacks. Prisma Cloud is the industry's most comprehensive CNSP, addressing security needs across all stages of the cloud software development lifecycle, no matter which cloud services you use or how they are managed.

With Prisma Cloud, teams can automatically detect vulnerabilities that are difficult for even the best-trained security engineers to notice manually. Prisma Cloud can also enforce best practices to help ensure that all members of your team operate in a secure fashion when deploying workloads to the cloud.

To learn more, you can [request a personalized demo](#) or [watch a recorded demo at your convenience](#).



Cybersecurity
Partner of Choice

3000 Tannery Way
Santa Clara, CA 95054
www.paloaltonetworks.com

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

© 2021 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at <https://www.paloaltonetworks.com/company/trademarks.html>. All other marks mentioned herein may be trademarks of their respective companies. [prisma_wp_cloud-security-blind_073021](#)