panoptica

Demystifying Cloud Security: Dispelling Common Misconceptions for Robust Protection - Panoptica

Table of Contents

Introduction

According to a Gartner cloud adoption survey, it was discovered that **76% of organizations are leveraging multi-cloud architectures**. As organizations continue to grow and expand their multi-cloud presence, it likewise means that their attack surface is growing. With digital-first experiences quickly becoming the primary engagement customers have with brands across industries and markets, it's essential that organizations ramp up their security efforts.

Growing security at scale is often easier said than done, and orgs may fall into the path of least resistance, instead of truly scoping the projects to the depths they require to harden all their cloud layers. Security, especially in the cloud, where workloads are constantly dynamically shifting and expanding, requires expertise and a longer, more strategic view to understand the many pitfalls and common errors that can occur.

In this eBook, we aim to kill some of the common cloud security misconceptions in the market.

We will cover:

01. Common Vulnerabilities and Exposures (CVEs)

CVEs are the most pervasive risks your environment may have, but does that mean cloud security should start from there? Moving beyond CVE patching and instead taking an “inside-out” perspective of cloud security will help your organization better batten down its hatches.

02. Cloud Workload Protection Platform (CWPP)

Many organizations add a CWPP tool to their arsenal and believe all will be quiet on the Western front. However, you wouldn’t wait for a burglar to enter your home before buying and setting an alarm system, would you? CWPP solutions are not enough to keep your network safe from penetration – you require proactive positioning and solution to ensure your network doesn’t fall victim to opportunistic attackers.

03. The Principle of Least Privilege

While least privilege claims to protect against lateral movement and privilege escalation, in the cloud an attacker can operate within your environment using just a couple of permissions. Should the privilege of least principle then truly apply to the cloud?

04. There is No Private in Your Public Cloud

Do private networks really ensure that your assets are safe from external access? Data safeguarded in private networks can be breached – private doesn’t mean impenetrable.

Join us as we debunk the most common cloud security misconceptions and help highlight the avenues of risk that still exist alongside some of these cloud security defaults. With a better understanding of these misconceptions, you can more effectively secure your cloud environments and apply more proactive security practices to your business.

Part 1: Common Vulnerabilities and Exposures (CVEs)

Going beyond the common vulns to better secure your environment

Background

Common Vulnerabilities and Exposures (CVEs) – a short descriptor of exactly what we're getting – the exposures and vulnerabilities that are most common. But what happens when uncommon issues are discovered and exploited by attackers? What if attackers just want us to think they'd only exploit common issues and vulnerabilities?

Securing CVEs seems like it should be the right place to start when attempting to secure your broader cloud environments. However, exploring and exploiting common vulnerabilities and exposure are where script kiddies start from... a common entry point for bots to exploit. We don't want to land in the security hall of shame typically set aside for organizations that were exploited and impacted by ransomware thanks in part to an unpatched CVE from months ago.

The importance of remediating vulnerabilities in our environment is obvious. But are we misleading ourselves by thinking that remediating vulnerabilities on public assets will eliminate all threats? Starting from the perimeter often leaves a false impression that we are safe. It's easier to communicate such efforts to managers than it is to harden a Linux base image with SELinux.

But alas, since we all want to check that simple task off on our endless and growing list of tasks, we end up putting our efforts into this low-hanging fruit. But the more important question is, "does the fact that we don't have any open CVEs present in our environment right now, mean that we're secure?" The simple answer is no.

Let's have a look at some reasons why this is the case.

Application Security

Application security starts with the opposite of CVEs, focusing on the unknown issues. These issues are often discovered by various security testing methods (Static Application Security Testing (SAST), Dynamic Application Security Testing (DAST), Runtime Application Self-Protection (RASP), etc.) and are prioritized by the Open Web Application Security Project (OWASP) Top 10.

The business use case of having a web application or an API exposed to the internet applies to most of us. Web applications are what many may consider their “crown jewels” and most organization’s defense strategy focuses on preventing them from being accessed and exploited.

Although customer databases are the so-called “crown jewels,” penetration testers typically start their cyber security testing scenarios from the internet. We want their path to the crown jewel to be as long and circuitous as possible to minimize their chances of obtaining this treasure.

If the database is the crown jewel, that’s where the penetration tests should start. We put most of our security testing effort on web applications and then put extra effort around the crown jewel.

However, that’s exactly how you end up with a web application without CVEs and that’s free of XSS (good for you!) but with a weak password that is hardcoded in the config file, communicating in an unencrypted channel in the backend database.

CVE vs. Bug Bounty

Crowdsourcing manual security testing is a great thing. It’s a win-win for all sides. Whitehat hackers end up being rewarded for their hard work, we eliminate unknown threats and vulnerabilities in our code, security is improved exactly where we want, and awareness regarding security is enhanced.

Compared to common vulnerabilities, most of which do not have a public exploit, vulnerabilities found in bug bounties are always exploitable – so why are we so afraid of going out publicly for bug bounties?

That’s another misconception that remains a mystery. For some reason, organizations seem to think that remediating common vulnerabilities should be first, and bug bounties are a security concept to be saved for mature security enterprises.

Open Sources

While there are huge benefits that come with open-source code, we can't overlook the fact that any developer can write a library which can be adopted by developers, without anyone ever noticing a security bug.

Luckily, the wild jungle of open-source security has experienced a significant improvement over the last few years, wherein companies like Snyk have developed automated solutions to alert developers of common vulnerabilities in open-source libraries in use. That's a tremendous achievement for security owners, where security has shifted left and the responsibility for remediation is on the developers.

Bypassing Your Perimeter

The career path that many information security experts have taken usually does not include being on the offensive side. Panoptica's founders had the opportunity to spend multiple years as offensive researchers and white-hat hackers, specifically after having held security ops and engineering roles. This experience shifted the way our founders looked at things in cloud security.

The most important concept that shifted is the idea of bypassing the perimeter. Like many, our team thought that exploiting a common vulnerability or a zero day would be required to bypass the perimeter (but that's saved for nation-states or top APTs — or maybe that's just another misconception). But in the first two years on the offensive side taught our founders otherwise - there are just so many ways of bypassing that same perimeter that teams often invest a great deal of time into. Whether it's done by using default passwords for exposed management interfaces or dashboards, using leaked credentials, taking advantage of servers left open without authentication, or a variety of other methods, there are so many ways to bypass the perimeter.

Secure From the Inside-Out

Common vulnerabilities and exposures are simply the first step an attacker would take to penetrate your organization. Making this step harder won't stop them and won't prevent data breaches.

Instead of making the attacker's life harder with every step they take, common approaches are making it easier for attackers, because teams are securing from the outside-in and not the inside-out. The closer one is to the crown jewels, often the less effort is put into security.

This shouldn't be the default or status quo. Organizations need to take an inside-out approach to ensure that from the deepest and more valuable layer of the environment to the more exterior layers are impenetrable.

Moving Beyond CVEs

Let's imagine a world without common vulnerabilities, a world where the first asset, our public asset, is accessible to everyone, not just hackers, but to everyone.

Randomly choose an instance in your cloud environment, not a public one, and review its security posture. You'll probably find one of the following:

- A role with risky permissions attached (FullAccess ones or Admin)
- Security Group with 0.0.0.0/0 enforced
- Unnecessary sudoers (usually in Linux)
- Sensitive information in old files in the operating system

Adopting a state of mind in which it's understood that the first asset will always get breached would dramatically improve cloud security posture. It's time to start approaching security from the inside-out, securing the crown jewels and making every element of configuration super-hardened. In this construct, the attacker would be the one who must put a great deal of effort into moving forward or performing any lateral movement across the system.

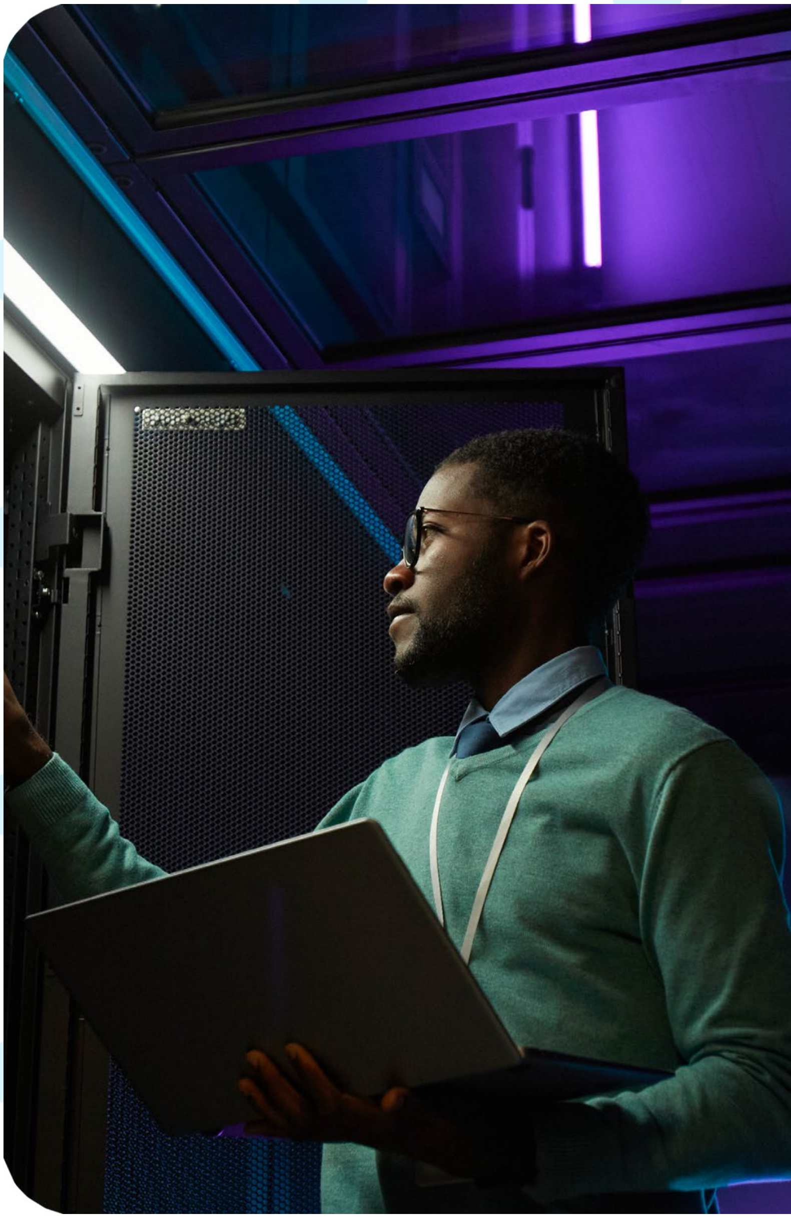

Part 2: Cloud Workload Protection Platform (CWPP)

From "right now" protection to proactive prevention in your cloud

As organizations accelerate their move to the cloud, it's no surprise that many CISOs want to stick with what they know when it comes to security. On-premises, traditional best practices dictate that every endpoint or server needs to have its own endpoint protection. This is why anti-virus software on machines is usually a top priority. However, lifting and shifting this approach to the cloud just doesn't work.

Let's take a deeper dive.

New Cloud Environments Have Changing Security Needs

Think about the legacy firewall approach on-premises, where a perimeter-based firewall around the network was seen as sufficient to keep attackers at bay. Today, it's well established that the idea of the traditional perimeter is dead, and that hybrid and complex IT environments can no longer be protected using a North-South perimeter approach. After all, most of the traffic is already inside the network, East-West. The next logical step is therefore endpoint protection, where each machine or server has its own protection in place. Even if an attacker enters the network, (and it's now best-practice of a zero-trust model to assume access) critical assets are kept under lock and key.

However, in our hybrid and multi-cloud environment, security solutions need to evolve past a reliance on endpoint protection, too. On the cloud, endpoint protection technologies like Endpoint Detection Response (EDR) are comparable to workload protection, where Cloud Workload Protection Platforms (CWPP) use signature-based detection and anomalous behaviors to identify suspicious activity. *By design, this is only effective as a security measure once attackers are inside the network.*

The Cat and Mouse Game of CWPP

Seeing as so many organizations are putting their security in the hands of this kind of workload protection, let's put CWPP under the microscope.

As everything that we do as security professionals is intended to keep the attacker at bay, it makes sense to do this while thinking about the psychology of the people behind today's cyber – attacks – the attackers themselves.

Any hacker goes through something called the kill chain, which starts from reconnaissance, and goes through various stages, from weaponization and delivery, to exploitation to deliver on their objective. From the outside looking in, we see a single path that an attacker takes to succeed in their campaign, without focusing on the one thing that can make all the difference – the time and effort that goes into the attack.

Today, even when cyber warfare is an almost daily occurrence, there is still a limit on the time and resources that attackers can put into each campaign. That's why we see malware being reused in different campaigns over time, and it's why even sophisticated attack patterns are often based on long-known signatures and threats. There is one obvious insight we can take from this. We clearly want to make it more time consuming and difficult for an attacker to make it to their final objective, namely – the data or critical assets they are looking to reach. Now let's consider whether the CWPP approach works towards that goal.

We can compare it to a security guard and a guard dog walking around the base of a mountain with a flashlight. The mountain holds everything they want to protect. The guard is listening for sounds of a trespasser, looking for any signs of movement, while the dog uses his sense of smell to detect anyone or anything out of the ordinary. However, the mountain itself is not that high or strong, and the guard and the dog can only be at one point along its landscape at any one time. On top of that, they may have a long list of what they would consider "unusual behavior," but they can only protect against their existing schema, which is limited to what they have already seen or experienced in the past.

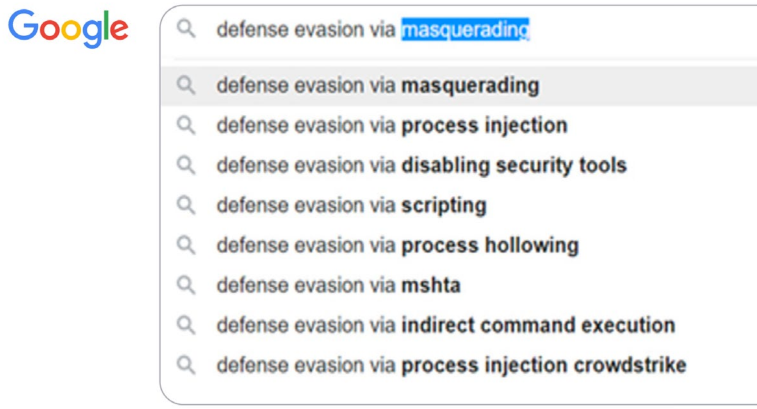

defense evasion via masquerading

defense evasion via masquerading
defense evasion via process injection
defense evasion via disabling security tools
defense evasion via scripting
defense evasion via process hollowing
defense evasion via mshta
defense evasion via indirect command execution
defense evasion via process injection crowdstrike

CWPP solutions effectively do these same actions – They scan the network for known attack signatures, review audit and event logs, and process traffic and activity. In return, attackers know exactly what security professionals are looking for, and can evade these processes by changing even a minute part of a known threat, down to a single hash or checksum, or by simply finding a vulnerability in the workload protection.

Below you can see a short list of the ways that attackers evade detection from CWPP or EDR tools. While these are just search terms, you can even find 326,000 results for readymade tools on GitHub that allow attackers to bypass signature-based workload or endpoint protection.

The New Cloud Security Posture Management: Contextual Cloud Security to Break the Cycle

In contrast, **Cloud Security Posture Management (CSPM) tools take a proactive approach to cloud security**. Instead of looking for known attack patterns and protecting against them, CSPMs use techniques such as configuration hardening, that makes the mountain you're protecting harder to climb in the first place. Rather than wait for an attacker to enter the network, CSPM also shows you where an attacker is most likely to make an entry attempt, so that you can shore up your defenses ahead of time. Your mountain becomes far too much effort to climb, and the attacker takes his limited resources elsewhere.

IBM puts configuration hardening among the top three items for CISOs to invest in when securing the cloud. Configuration hardening isn't based on signatures, anomalies, or profiles. It's a process where your security teams will be able to look at the entire configuration of your environment, in a single view. From there, they can uncover any threat, from dangerous defaults, to cleartext credentials, risky network configurations or vulnerabilities, unsafe permissions, and more.

When implemented intelligently and with visualization at its core, you can then detect the exact path that an attacker might take to access your crown jewel applications or data, allowing you to harden exactly where and how it's needed, ahead of time. In Panoptica's experience, this approach can help you to **close 80% of infrastructure issues with minimal effort**.

Shifting to a Proactive Approach to Cloud Security

Looking back at the IBM study, researchers also found that the average time it takes to identify and contain a breach is 279 days. This is the promise of CWPP tools that use anomaly detection and signatures. When we make CWPP our priority, we're sending a clear message to the hackers. We're saying, "You have plenty of time to play inside our network" and "we're willing to continue engaging in this endless cat and mouse game even though we always seem to be 2 steps behind."

In today's complex cloud landscape, the concept of "wait and see" is not an effective strategy.

CISOs need to be more efficient about how they protect their organizations, focusing on pre-emptive and proactive measures that harden security ahead of time, rather than relying on those that simply react to known attack patterns once an intruder is already inside their system.

Starting with CSPM tools that provide a full view of your cloud infrastructure will enable you to uncover the routes that attackers could take to your critical assets and data and shore these pathways up via configuration hardening.

Now, even while assuming access to the network, you're making it exponentially harder for hackers to take that next step, or to reach anything of value. This sends a clear message to the attackers: "Move on, you won't find opponents for cat and mouse under this roof."

Leveraging a cloud security platform that provides the tools to continuously visualize, detect and block any dangerous attack path in your cloud environment creates a shift if your perspective – instead of operating as a security team protecting single workloads, you can view everything as an attacker would, by looking at the whole picture. Only with this view can you truly move from reactive to proactive mode, allowing you to harden each and every element of your configuration against today's level of threat.

Part 3: The Principle of Least Privilege

Should the privilege of least principle really apply to the cloud?

With the principle of least privilege, a concept largely used for on-prem, is a huge amount of work to apply to the cloud, for little reward or risk reduction.

What is the Principle of Least Privilege?

The principle of least privilege (PoLP) is a security concept for computer systems where you give users exactly the permissions they need to do their job, and nothing more. It was created for and applied to on-premises security environments, and on-premises at least, it can be extremely effective at reducing risk.

Here's how it works.

All users have access permissions, which will govern what they can interact with in any network, whether that's access to read-only versions of files, read write access, root privileges, or any others. Excessive privileged access means that a user has more access to an environment than they need to given their designated role.

If a freelance designer is given all-access to their client's network, with virtually unlimited privileges, when all they really need is permissions to a single folder where the logos are kept – it's quite clear that there's a problem with their level of privilege access policies. In contrast, if a super user account can only view read-only files, they obviously don't have enough access, and least privilege has gone awry somewhere, probably with too coarse policies dictating their access.

In a smart and secure on-premises environment, tools like micro-segmentation and zero trust policies have used the principle of least privilege to narrow down risk in several ways:

Ring-fencing critical assets

This is the process of isolating specific "crown jewel" applications so that even if an attacker could make it into your environment, they would be unable to reach that data or application. As few people as possible would be given credentials that allow this kind of access, therefore following least privilege access rules. Crown jewel applications could be anything from where sensitive customer data is stored, to business-critical systems and processes.

Establishing role-based access control

Based on the role that they hold at the company, RBAC or role-based access control allows specific access to certain data or applications, or parts of the network. This goes hand in hand with the principle of least privilege, and means that if credentials are stolen, the attackers are limited by what access the employee whose credentials were obtained holds. As this is based on users, you could isolate privileged user sessions specifically to keep them with an extra layer of protection. Only if an administrator account or one with wide access privilege is stolen, would the business be in real trouble.

Isolate applications, tiers, users or data

This task is usually done with micro segmentation, where specific apps, users, data, or any other element of the business is protected from an attack with internal, next-gen firewalls. Risk is reduced in a similar way to the examples above, where the requisite access needed is provided using the principle of least privilege to allow access to only those who need it, and no one else. In some situations, you might need to allow elevated privileges for a short period of time, for example during an emergency. Watch out for privilege creep, where users gain more access over time. This can be hard to track and often can go uncorrected.

More recently, with the rise in cloud-native deployments and data centers, many security professionals have attempted to lift and shift the concept of least privilege to the cloud.

When it comes to securing the cloud, here is why we believe that least privilege is dead.

One or Two Permissions is Enough to Achieve Superuser Privileges

Least privilege is predicated on the idea of escalating privileges and credentials. An attacker breaches the network, and they can make lateral moves from one account to another, moving across the network and gaining greater and greater access until they reach the ultimate payload. This is a very "on-prem" concept.

To take full account ownership on the cloud, a user would really only need to have a couple of permissions, in an environment where there are often tens of thousands. For example, on AWS there are more than 10,000 different IAM actions, distributed across the various services in the cloud. These permissions are segmented between read, write, or management actions. To create a data breach, all an attacker needs is to gain access to a single IAM role or set of user credentials that has S3:GetObject only with a wildcard in the resource. In short, any S3 bucket that doesn't have explicit restrictions to the account is open for business to an attacker. We talk a lot about the risks of S3 buckets in this article on the risks of misconfigured cloud buckets and what you can do about them.

An example in practice: The CapitalOne breach

This is exactly what happened in the CapitalOne case. Security experts continue to call out the famous data breach as an issue that could have been prevented with least privilege in mind, but this really isn't the case. With Capital One, just two IAM actions, '**List S3 buckets**' and '**Get objects**' (a sync command) allowed the attackers access, the ability to prompt a $190M settlement, and a data breach of more than 100M customers.

With this attack path available from just two permissions, the principle of least privilege just isn't possible. It's not about restricting privilege access permissions or using role-based user access on cloud platforms like AWS, Azure, or GCP. The truth is you need a completely different approach.

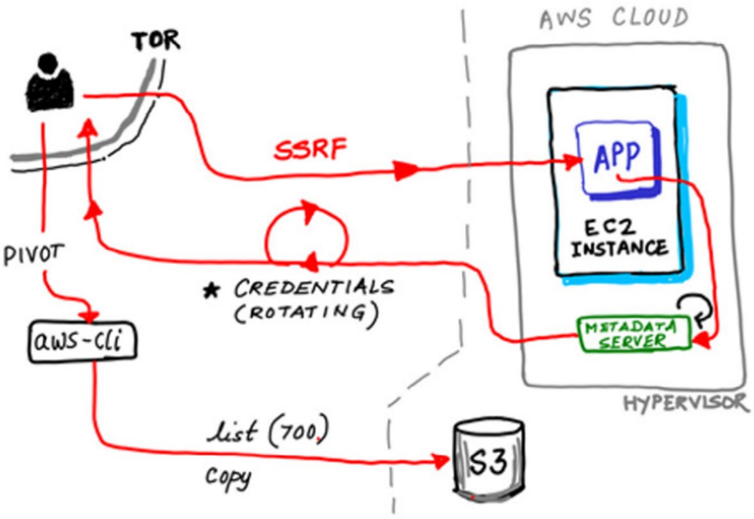

Image Source: ShiftLeft

The same is true when you think about threats due to Privilege Escalation in the cloud. While least privilege claims to protect against lateral movement and privilege escalation, in the cloud you can achieve this with just a couple of permissions.

Think about AWS, for example. As RhinoSecurity points out in this great blog, an attacker could use **ec2:RunInstances** and **iam:PassRole** as just two permissions that allow for full account ownership. It's not about gaining specific privileged credentials at all, as this could be done with minimal privilege or superuser privileges alike.

Built-in Managed Policies and Roles for User Access

The principle of least privilege also relies on the idea that you can manipulate and manage permissions and access rules, fully controlling your environment.

That's also not really an idea that can be widened to fit a hybrid or cloud-native environment.

Examples of these policies

On both AWS and Azure, there are managed policies and roles that are set out by the Cloud Service Provider that can't be manually changed. In most cases, users and developers are unaware of the risks of these policies and roles.

Two examples are the **AWSLambda_ FullAccess** and **AWSElasticBeanstalkManaged** Updates-CustomerRolePolicy.

Although these are used by-design to achieve rapid agile development cycles out of the box, making DevOps lives easier, both also allow for privilege escalation when exploited.

One good example is the **AWSElasticBeanstalkService** managed policy that AWS announced for deprecation, and is no longer available for attachment to users, groups or roles, as of April 15th, 2021.

AWS recommends using the new managed policy, **AWSElasticBeanstalkManagedUpdatesCustomerRolePolicy**.

In this policy, some actions have had certain resources and conditions defined, but privilege escalation is still possible by exploit, a huge security gap.
Here's a gist Azazar created for more information and a proof of concept.

The Endless Work of Working Towards a Least Privilege Model

The Cloud Infrastructure Entitlement Management (CIEM) domain is built around tackling these issues and supporting businesses in achieving the principle of least privilege in the cloud.

However, this causes a huge headache for DevOps teams and cloud security owners when attempting to maintain lean policies under strict resource requirements. In some cases, CIEM tools will claim to offer ways to achieve the principle of least privilege creation at the build stage, baked-into CI/CD pipelines. This ROI is in fact, impossible, and the attempt causes an endless amount of maintenance and frustration.

Most developers in cloud-native companies are unaware of the risks of managed policies or wildcard policies in the cloud, and that’s fine! It’s not part of their role to gain such an in-depth knowledge of cloud risks. In fact, attempting to get up to speed with all the changing risks of a dynamic environment would slow down the very DevOps pace that they are aiming to achieve by leveraging the cloud.

Consider the time that would be spent every time we modify the policies used by our users or applications, for example when we add a new database or new IAM actions. Trying to achieve privilege enforcement here would be beyond tedious. This process involves continuously checking each identity and specific permissions, asking yourself – is it necessary for this identity to be using these specific permissions? If not, and we decide to remove one, will it break production or cause problems for the user? It’s simply not maintainable or scalable as businesses grow.

Moving Past the Principle of Least Privilege in a Cloud Security Environment

Cloud-native companies have just **two options**.

01. They can take the direction of policies shared between many applications, using default AWS or Azure permissions, which is not the principle of least privilege

OR

02. They can attempt to create policies for each application and user

The second option generates so much maintenance and overhead for security and operations teams, that it negates the benefits of the cloud in the first place. You might be working towards a least privilege model, but you won’t have much else!

For system stability, and to restrict access in a smart way, what kind of privilege enforcement does make sense to apply to the cloud?

An Alternative Option for Privilege Access Management

At Panoptica, our approach is to offer Guardrails to solve this use case. By creating your own managed policy unique to your environment that blocks specific attack paths, there's no need to fight with the least privilege model. You can apply access controls using Guardrails, and your teams can work unimpeded without any impact to the pace of DevOps.

Instead of getting fine-grained about specific one-to-one connections between the account and a user or working endlessly on permissions into any given resource (all of which need to be maintained over time), you just **set and forget Guardrails that block the potential attack.**

As part of our cloud security platform, we offer automated Guardrails that are dynamically built inside your DevOps tools, for example via Terraform or JSON, but can be customized to fit your specific environments' requirements. A misconfiguration can be found and alleviated automatically, eliminating the vast majority of the manual work that DevOps will need to do. If an asset has an overly-permissive role where it should have just read-only access, the required Guardrail will automatically appear with all the necessary denies, shoring up this risk. No impact on production, and no need to change the read-only permissions themselves.

Much smarter than relying on the principle of least privilege for an environment it was never intended for, right?

Do private networks really ensure that your assets are safe from external access?

The idea of keeping your assets secure using private networks is a pervasive myth.

Many people that use the public cloud rely on segmenting assets in private networks inside their cloud environments to keep their assets secure. Networking-wise their assets are then private, so it's easy to see where this misconception comes from. While it's true that if there is no public IP, or if the private network isn't exposed to the internet, then these assets can't be accessed externally. However, that doesn't mean there is no access at all and that their data is safe from manipulation, extraction, or risk.

Your Internal Servers Aren't as Private as You Think They Are

The truth is that every single asset in AWS, or on any public cloud is accessible through a cloud service provider's API. It doesn't matter whether they are private or public networks, they can all be accessible, and allow attackers to shut down services or make changes to your environment, if the right permissions are granted.

All that needs to happen is a single user credential or public access key exposure – something as simple as a compromised developer workstation that has sufficient permissions, and the attacker will be able to access any private asset in the cloud.

Let's look at three real-world examples of how an attack on "private" networks can occur, and how the misconception of data safeguarding in the cloud occurs.

01. Information Stolen from a Public GitHub Repository

First, let's take the example of a developer who generates an access key for his AWS account. The user's access key includes the permissions to deploy and modify existing EC2 instances, as well as their network access rules. Within the environment, the company keeps private workloads that run the customer database. Inside this database there is sensitive financial information, from credit card details to credentials such as usernames and passwords.

Our developer is only human, and as part of his code, he forgets his access key ID and the secret access key in a public GitHub repository. As many attackers have automation working around the clock to constantly scan public GitHub repositories, it's only seconds or minutes later that the access key ID and secret access key are in the hands of a nefarious actor. Quickly, they open the existing server to the public internet, which means the customer database (and all the sensitive information held within) is now accessible through a publicly exposed server.

02. A Malicious Open-Source Library

Another common exposure point, your admin workstations – likely to have some of the highest permissions and therefore a regular target. In this scenario, a malicious bot is looking to steal secret credentials to cloud environments from admin workstations, and it accomplishes this using an open-source library from inside an asset discovery tool. This tool works to continuously map often-neglected assets in your cloud environment such as old backups for databases, snapshots for old servers, and more.

When the bot uncovers these assets, it will automatically attempt to share those assets with its own cloud account. If successful, this data can then be extracted and scraped to be sold on the Dark Web. As of 2020 alone, 15 billion stolen login details, stemming from 100,000 breaches circulating on the Dark Web.

If you think it can't happen to you – you're wrong.

03. Cross-account Access Breach

Our last example involves cross-account access. Imagine your platform engineering manager uses cross-account access as a legitimate part of his job, in this case to integrate a solution to a data analytics platform. The account is used for private data only, and none of the assets are publicly exposed to the web. Secure? Maybe not.

In this case, the data analytics solution itself is breached because of an application vulnerability in the SaaS platform. This may have come from human error on the vendor's part – but there is nothing that your business, or your platform engineer could have done to prevent it. The attacker abuses the cross account access to its customer, extracts all data, and then encrypts it – holding your business to a heavy ransom to release the data, and throwing your organization into a media spotlight nightmare.

It's Not Called "Public" Cloud for Nothing!

Using private networks and then patting yourself on the back that your public cloud environment is secure is not a sound cloud security strategy. It's important to recognize that on the public cloud – nothing is ever really private when a sophisticated attacker (or plain old human error) gets involved! There are several routes to data leaks, ransomware attacks, or public access for your private data. Don't believe us? Check all your cloud subdomains on Panoptica's tool, Recon.Cloud. Here, we can show you simply by scanning your AWS public cloud footprint, how much of your ***internal*** network a potential attacker can actually see.

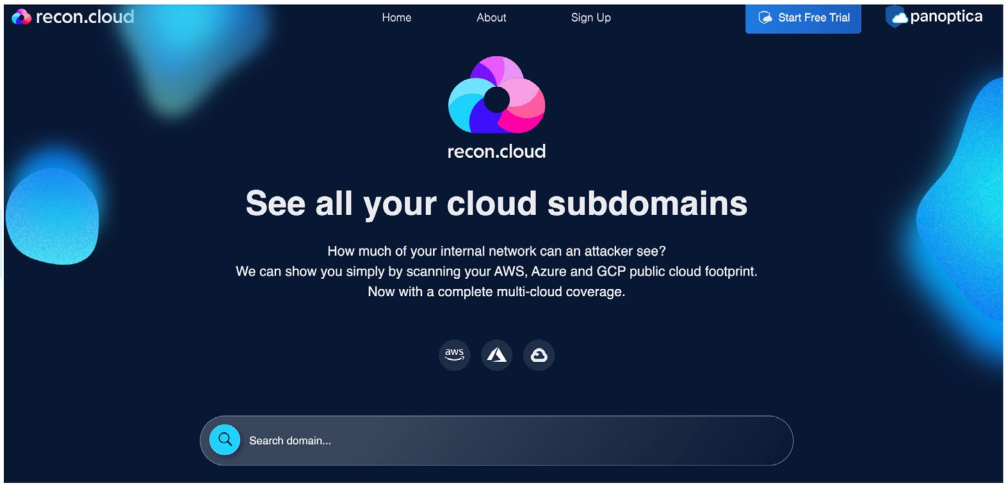

See all your cloud subdomains

How much of your internal network can an attacker see?
We can show you simply by scanning your AWS, Azure and GCP public cloud footprint.
Now with a complete multi-cloud coverage.

Don't ignore private networks – because we can guarantee that the attackers won't either.

Conclusion

The cloud is a dynamic and complex space.

There are many things that can and do go wrong as organizations expand their multi-cloud presence and attempt to secure their cloud environments.

- 99% of companies will have cloud misconfigurations that they don't even know about[1]
- Until 2025, 99% of cloud failures will be due to human error[2]
- SMBs can face up to 5,000 security alerts per day, yet only 55.6% of them investigate these[3]

Even with the best of intentions, cloud security misconfigurations are pervasive and create the potential for high impact exploits to occur. The common cloud security misconceptions we have detailed and hopefully, busted(!), are just the tip of the iceberg. It is essential that organizations take a proactive stance and work from the inside-out of their environments to ensure that they do their utmost to secure their cloud.

Unfortunately, there is no silver bullet. The cloud space is continually shifting – cloud services are expanding, and opportunistic attackers are always on the hunt, looking for new ways to breach your assets and gain access to those crown jewels.

So how can you ensure your organization is taking the right steps to stay ahead of these nefarious actors?

The Way Forward

A single, comprehensive platform to meet dynamic shifts in cloud security requirements

With shifting cloud security requirements and new assets constantly being added to the cloud, organizations need a complete cloud security platform to protect and secure their cloud stack and business assets, from build to runtime.

At its core, best practices indicate that organizations should partner with vendors who offer:

- **Graph-based technology at the heart of the system**: Through graph-based algorithms, surface all potential steps vulnerable to attack. Visualize a complete flow of a potential attacker from the network entrance point all the way through the workload risk and reveal the final impact.
- **Critical attack path approach**: The attack path combines multiple singular findings into a path on top of the cloud environment topology. This approach exposes the real risk of these security findings as each attack path reflects the probable chain of attack that a nefarious actor would likely leverage in the specific environment.
- **Prioritization**: Looking at a variety of layers within multi-cloud environments builds the story behind each disparate finding. This multi-faceted approach builds the context or story by connecting between findings to surface what really matters. This provides the ability to accurately build prioritization of findings.
- **Remediation**: Out of the box remediation will reduce the time it takes for your Sec/DevOps teams to better secure any discovered vulns. Dynamic deny guardrails that can be tailored to your organization's specific requirements, means no one size fits all. Instead, a customized solve for your specific risks.

Through leveraging an offensive perspective, you can visualize your environment from the eyes of an attacker, and better proactively protect your most valuable assets while reducing the effort required from your team.

Panoptica is Outshift by Cisco’s cloud native application protection platform (CNAPP) that uncovers and remediates vulnerabilities during development through to production, ensuring your applications and cloud environments are secure and compliant. Through graph-based technology, the platform is able to unlock visual insights, critical attack paths, and speed up remediation to safeguard your modern apps across multiple hybrid cloud platforms.

To learn more visit: **https://www.panoptica.app/**

[1] https://www.zdnet.com/article/99-percent-of-all-misconfiguration-in-the-public-cloud-go-unreported/

[2] https://www.gartner.com/smarterwithgartner/is-the-cloud-secure

[3] https://www.spiceworks.com/it-security/network-security/articles/network-security-engineer-key-job-role-and-skills/#:~:text=The%20same%20report%20also%20points,them%20investigate%20these%20security%20warnings.