2024

State of Cloud Security Report

Uncovering what is lurking in the depths of cloud environments

Inside This Report

2024 STATE OF CLOUD SECURITY REPORT

Foreword

The past year has been impacted by a challenging economic climate and shrinking budgets, putting cybersecurity defenders on the back foot.
At the same time, attackers are becoming increasingly sophisticated - finding new attack vectors and leveraging AI to generate malware and automate attacks. Ransomware attacks in particular, reached record numbers in 2023.

In the world of cloud security, the increasing adoption of cloud services and cloud-native technologies is heightening both the possibilities and risks. With most organizations now using three or more cloud service providers, cloud environments have become more complex than ever before.

Despite these challenges, we believe that when security teams are focused on their most critical risks and are equipped to remediate these quickly, they can stay one step ahead of their attackers. We hope that this report prepared by the Orca Research Pod will be a valuable resource to help organizations do just that.

Gil Geron
CEO and Co-Founder of Orca Security

About the Orca Research Pod

The Orca Research Pod is a group of cloud security researchers that discover and analyze cloud risks and vulnerabilities to strengthen the Orca Cloud Security Platform and promote cloud security best practices.

Research Methodology

This report was compiled by analyzing data captured from billions of cloud assets on AWS, Azure, Google Cloud, Oracle Cloud, and Alibaba Cloud scanned by the Orca Cloud Security Platform.

Report Data Set:

- Cloud workload and configuration data
- Billions of real-world production cloud assets
- Data referenced in this report was collected in 2023
- AWS, Azure, Google Cloud, Oracle Cloud, and Alibaba Cloud environments

25+ vulnerabilities discovered on AWS, Azure, and Google Cloud

2024

+ Sys:All GKE Loophole ⇒
+ Three new Azure HDInsight vulnerabilities ⇒

2023

+ Azure Digital Twins SSRF ⇒
+ Azure Functions App SSRF ⇒
+ Azure API Management SSRF ⇒
+ Azure Machine Learning SSRF ⇒
+ Azure Storage Account Keys Exploitation ⇒
+ Azure Super FabriXss ⇒
+ Two Azure PostMessage IFrame Vulnerabilities ⇒
+ Bad.Build Supply Chain Risk in GCP ⇒
+ 8 Cross-Site Scripting (XSS) vulnerabilities on Azure HDInsight ⇒
+ Unauthenticated Access Risk to GCP Dataproc ⇒

2022

+ AWS BreakingFormation ⇒
+ AWS Superglue ⇒
+ Databricks ⇒
+ Azure AutoWarp ⇒
+ Azure SynLapse ⇒
+ Azure FabriXxs ⇒
+ Azure CosMiss ⇒

⚠ The report underscores a further need to cover the basics and for context-driven risk prioritization so the most critical issues are fixed first.

Executive Summary

Leveraging unique insights into current and emerging cloud risks captured from the Orca Cloud Security Platform, this report reveals the most commonly found, yet dangerous, cloud security risks. Summarizing the results from our research, these are our main findings:

- **Basic security practices still lacking**
  We are still still seeing many risks that stem from not following foundational cybersecurity principles, and changing this remains pivotal to strengthening cloud security postures.
- **Many risks found on exposed and public assets**
  Worryingly, several of the risks are related to cloud assets in which security should be prioritized the most - those that are exposed and public facing, especially when they store sensitive data, or provide a path to sensitive data through lateral movement.
- **Some improvement in overall cloud security postures**
  More encouraging however, is that we found some reduced risk numbers compared to our previous 2022 report, with many fewer Log4Shell-vulnerable assets and a 1-5% improvement in security postures across industries.

Key Findings

81%

of organizations have public-facing neglected assets with open ports

Attackers routinely scan for open ports and known vulnerabilities, making these assets—which often do not have a patch available—prime targets.

21%

of organizations have at least one public-facing storage bucket with sensitive data

Misconfigurations in sensitive data storage increase the risks of exposed customer data, ransomware, reputational damage, and regulatory penalties.

61%

of organizations have a root user or account owner without MFA

Because the root user has complete access to all services and resources in the cloud account, it is highly recommended to use MFA as an extra layer of security.

82%

of organizations have a Kubernetes API server that is publicly accessible

Unrestricted access to the API server opens paths for attackers to reach underlying Kubernetes infrastructure and workloads, which could lead to data exposure and supply chain attacks.

59%

of organizations still have at least one asset vulnerable to Log4Shell

Despite Log4Shell being discovered 2+ years ago and described as the riskiest zero-day vulnerability in a decade, many organizations have still not completely eradicated the vulnerability from their cloud estate.

62%

of organizations have severe vulnerabilities in code repositories

These vulnerabilities exist in code that could imminently be pushed to production environments and cause data breaches and system compromises.

82%

of AWS SageMaker users have a notebook exposed to the internet

Amazon Sagemaker, a cloud-based AI platform, often contains sensitive training data, and bad actors are aware of this. This makes it important to sufficiently protect and limit access to these resources.

1-5%

improvement in industry cloud security scores

The good news is that all industry averages improved over time in 2023 - an encouraging sign that cloud security, despite a difficult economic climate, is being prioritized and leaving an impact.

01 Neglected Assets

NEGLECTED ASSETS

1.1 The Threat of Exposed Neglected Assets

A neglected asset is a cloud asset that uses an unsupported operating system or has remained unpatched for 180 days or more. As applications and operating systems reach EOL, vendors stop offering support, causing security and stability to decrease over time.

When neglected assets are public-facing, especially through the widely targeted ports **80, 443, 8080, 22, 3389** or **5900**, the risk of exploitation escalates. Attackers routinely scan for open ports and known vulnerabilities, making these assets prime targets. Upgrading these systems becomes imperative, as they serve as easy entry points for cyber attackers. Agent-based workload protection may struggle to address this risk as neglected assets are less likely to have an agent installed and maintained.

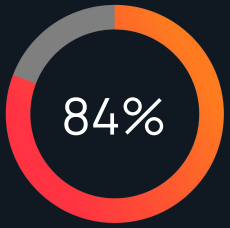

84%

of organizations have at least one public-facing neglected asset.

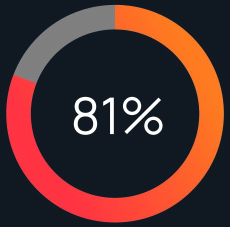

81%

of organizations have public-facing neglected assets with commonly exploited open ports.

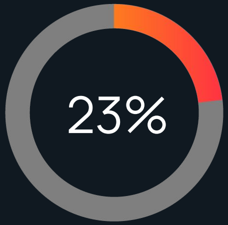

23%

of organizations have a subdomain at risk of takeover

NEGLECTED ASSETS

1.2 Subdomain Takeover Risk

A subdomain takeover is one of the easiest attack vectors for adversaries to exploit, and can be used for phishing campaigns, credential theft, or distribution of malware.

Subdomain takeover, of which we found that 23% of organizations are at risk, is when a subdomain's CNAME record points to a non-existent cloud service and a bad actor uses it to serve malicious content.

These loopholes usually appear when cloud services are shut down and organizations neglect to update or remove the corresponding DNS records, leading to unprotected domains and subdomains.

02

Data Exposure

DATA EXPOSURE

2.1 Exposed Sensitive Data

Storage buckets and databases are designed with security configuration settings to ensure that sensitive data such as personally identifiable information (PII), credit card data, healthcare information, as well as dev keys, secrets, and tokens, can be stored securely.

However, it is important to ensure that these assets are configured correctly. Making these assets public can put organizations at risk of data exfiltration, ransomware, reputational damage, and regulatory penalties. In general, data stores that contain sensitive data should **never** be publicly accessible. The fact that one-fifth of organizations have their data storage set up this way should be a wake up call.

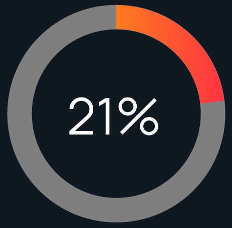

21%

of organizations have at least one public-facing bucket that contains sensitive data.

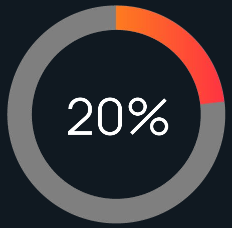

20%

of organizations have a public-facing database with sensitive data.

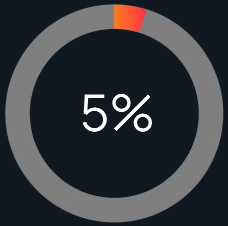

5%

of organizations have an Amazon S3 bucket with public ‘Write’ access.

DATA EXPOSURE

2.2 Public Write Permissions

An S3 bucket that allows public ‘Write’ access allows attackers to add, delete and replace objects within the bucket. This can lead to unintended changes, data leaks, and malware infiltration, as well ransomware attacks that overwrite or encrypt files.

Having an S3 bucket that allows public ‘Write access’ should be avoided at all cost. Even though only 5% of organizations have such an asset in their cloud, it is still very worrying due to its inherent high risk.

Even though by default an S3 bucket is always created as “private,” misconfigurations can, and do happen. On AWS, you can use IAM policies, bucket policies, and access control lists (ACLs) to define access policies for your buckets to help avoid this.

DATA EXPOSURE

2.3 Threat on Data in AI Models

Cloud-based AI platforms, such as Amazon SageMaker, are crucial for organizations wanting to leverage artificial intelligence, offering tools for building, training, and deploying machine learning models. However, since AI models use large amounts of training data, which can include proprietary and sensitive data, these assets are at higher potential risk.

Through accessing exposed SageMaker notebooks, malicious actors can gain unauthorized access to the AI model's underlying code. This can lead to the theft of proprietary algorithms or, in extreme cases, enable remote code execution (RCE). These risks emphasize the need for stringent security protocols in managing cloud-based AI models to safeguard data integrity and protect intellectual property.

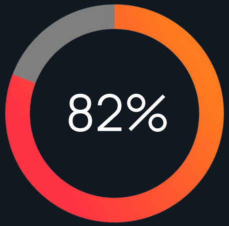

82%

of Amazon SageMaker users have at least one notebook exposed to the internet.

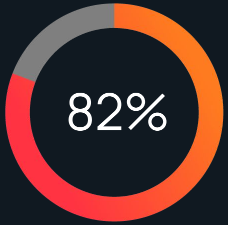

82%

of organizations have a publicly-accessible Kubernetes API server

DATA EXPOSURE

2.4 Exposed Kubernetes API Servers

A Kubernetes API server, facilitating all communication within clusters and with external components, becomes a prime target when exposed. While the majority of Kubernetes API servers require authentication, unrestricted network access makes it easier for attackers to modify resource states and potentially breach underlying infrastructure, containers, and other workloads. This is even more dangerous if a bad actor has managed to obtain credentials as part of a larger campaign.

This finding marks a **12% increase** from our 2022 report, underscoring the urgency to bolster security measures as Kubernetes adoption surges. While intentional public access exists for testing, the majority of publicly accessible API servers stem from misconfigurations.

The default in cloud-managed Kubernetes platforms is to expose the API server. This serves as a reminder to use secure, rather than merely default, settings.

03 Vulnerabilities

VULNERABILITIES

3.1 Two Decades of Vulnerabilities

While it may seem stunning to see vulnerabilities in cloud environments that actually **predate cloud computing**, it is perhaps not that surprising. As organizations move applications from on-premises environments to the cloud (also known as 'lift and shift'), existing vulnerabilities are often moved with them. Even though these vulnerabilities may be old, this doesn't mean there's no risk: **22%** of the 10+ years old vulnerabilities have known exploit code in the wild. Although some of the lower severity vulnerabilities remain unresolved, we found that **83%** of these actually do have a fix available.

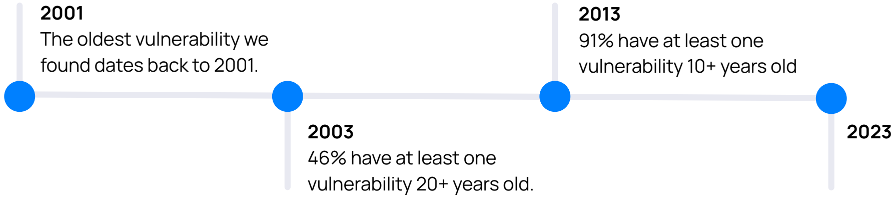

2001
The oldest vulnerability we found dates back to 2001.

2003
46% have at least one vulnerability 20+ years old.

2013
91% have at least one vulnerability 10+ years old

2023

While the weight of 20+ years of vulnerabilities can seem overwhelming, a context-focused approach to vulnerability management, focusing on the vulnerabilities most likely to be actively exploited and do the most damage in your environment, is recommended, as well as re-architecting legacy applications.

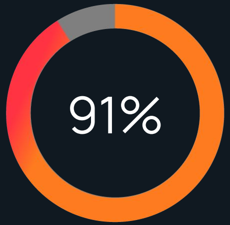

91%

of organizations have at least one vulnerability older than 10 years

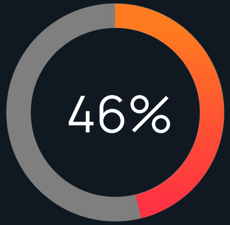

46%

of organizations have a vulnerability 20+ years old

The oldest vulnerability we found dates back to **2001**

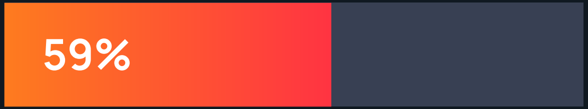

of organizations still have at least one asset vulnerable to Log4Shell.

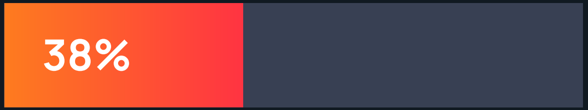

of organizations have a Log4Shell vulnerable asset that is public facing.

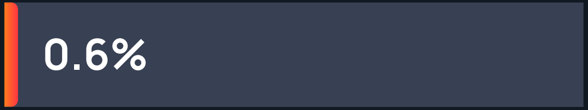

of all public-facing assets are vulnerable to Log4Shell.

3. Log4Shell is still a problem

Even though Log4Shell was initially discovered 2+ years ago and is described as one of most dangerous software vulnerabilities ever, the security flaw in Log4j is still present in many cloud environments. Apparently, attackers are also aware of this, as shown in recent reports highlighting that Log4Shell is still being actively exploited.

While less than 1% of public-facing assets are vulnerable—a **10% decrease** compared to our 2022 report—each instance represents an easily exploitable opportunity for a bad actor to launch an attack.

There's some indication that these Log4j vulnerabilities are being reintroduced through 3rd-party libraries used in the application development process. Therefore, it is important to scan development pipelines to prevent reintroduction of Log4Shell and other vulnerabilities.

04 Identity & Access

IDENTITY & ACCESS

4.1 Weak Authentication on Public-facing Assets

Public-facing workloads, accessible from the internet either accidentally or intentionally, constitute a significant portion of an organization's attack surface and are among the first targets that threat actors will attempt to compromise. When a public-facing workload allows for password authentication, and accounts have passwords that are common or have been part of credential dumps from previous public breaches, attackers have an easier time compromising the asset with dictionary and password-spraying attacks.

Adopting best practices such as multi-factor authentication, avoiding password-only authentication, and implementing appropriate controls for non-human/non-interactive authentication is crucial to safeguarding these systems, reducing the risk of exposure and potential compromise.

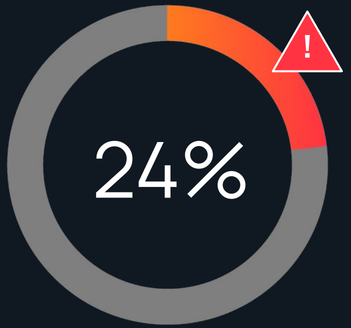

of organizations have at least one public-facing workload with a weak or leaked password.

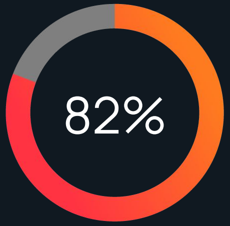

82%

of organizations have IAM user credentials that haven’t been used for 90+ days.

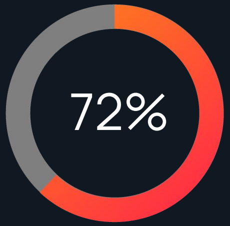

72%

of organizations have unused IAM roles.

IDENTITY & ACCESS

4.2 Unused IAM Users and Roles

Identifying unused users and roles, and deactivating those identities or revoking their privileges may seem straightforward, but it’s definitely not. Managing IAM configurations in the cloud at scale poses a significant challenge, due to complexity and ever-changing cloud environments, especially when managing large multi-cloud environments.

Roles in particular, often overlooked, present a unique challenge. Unattached to individuals, they're easily forgotten, created automatically during asset setup, and neglected during removal.

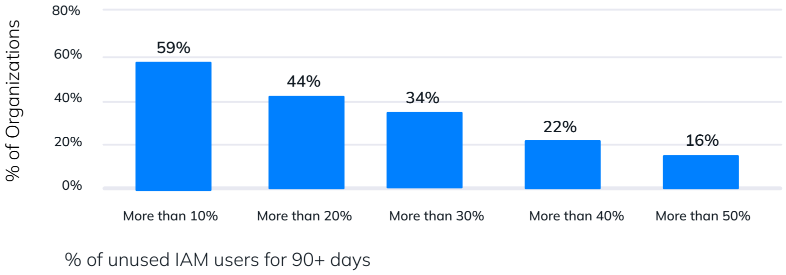

IDENTITY & ACCESS

4.3 Root Risks: Highlighting the MFA Gap

When creating a cloud account, one user is granted complete access to all services and resources in the account (i.e. the AWS root user, Azure Account Administrator, etc.). If this user lacks Multi-Factor Authentication (MFA) requirements, bad actors can potentially try to obtain login credentials using dictionary and password spraying attacks. MFA adds an extra layer of authentication assurance beyond traditional credentials, reducing the risk of unauthorized access.

While there can be many reasons why an organization hasn't implemented MFA for their account's root user, the most likely reason is because it's hard to manage MFA for shared accounts. However, this should not be seen as a blocker. Larger organizations for instance, may choose to store a hardware token in a safe or other controlled space. Distributed and cloud-native organizations can use software tokens that can be shared in a controlled way. Amazon plans to make MFA mandatory for AWS root user in accounts by mid-2024.

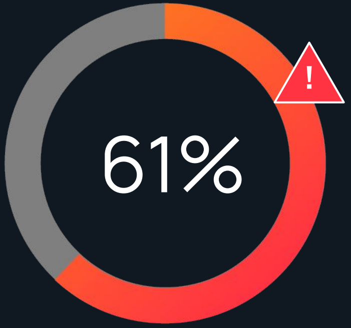

of organizations don't apply MFA on their Root/Account Owner user.

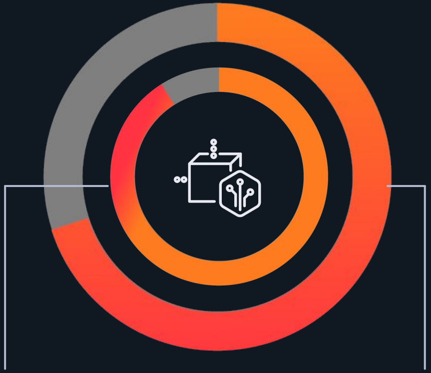

90%

of organizations have at least one asset that enables lateral movement.

72%

of organizations have public-facing assets at risk for lateral movement.

4.4 Lateral Movement Exposure

Lateral movement, defined in the Orca Platform as configurations, secrets, roles, and other risks that would make it easier for an attacker to pivot from one compromised asset to another, poses a significant threat to organizations, potentially allowing attackers to reach critical systems.

The ability to identify the cloud assets at greater risk for lateral movement, particularly in public-facing workloads, as well as the ones that form dangerous attack paths to business critical assets, helps security teams focus on fixing the most dangerous issues first.

A Zero Trust approach, enforcing stringent identity authentication and authorization, is valuable in mitigating lateral movement risks. By treating every user and device as untrusted, regardless of their network location, Zero Trust minimizes damage scope in the event of a breach.

05 Malware

5. Malware

As cloud services become increasingly interconnected and data transfers between various cloud platforms occur frequently, the risk of cloud-native malware infections rises. Cloud-native malware is specifically designed to target cloud environments, exploiting vulnerabilities in cloud infrastructure and applications, or propagating on cloud storage and collaboration tools.

A Trojan, by far the most common type of malware that we see in the cloud, is a type of malware attack that disguises itself as legitimate code or software in order to gain access to a system. Once inside, they grant attackers the same capabilities as legitimate users, including the ability to export, modify, or delete files. This ability to masquerade as legitimate user activity makes Trojans especially dangerous in cloud environments, where they can leverage existing cloud functionalities and tools to evade detection.

The fact that we are seeing malware in virtual machines, storage buckets, and containers underscores the need for vigilant security measures across all aspects of cloud infrastructure to safeguard against these pervasive threats.

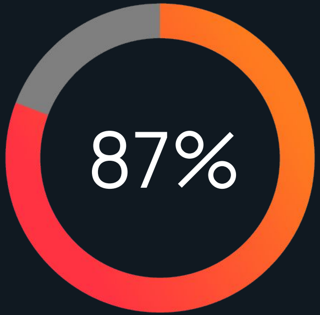

of cloud malware attacks are via known Trojans.

Malware by asset type

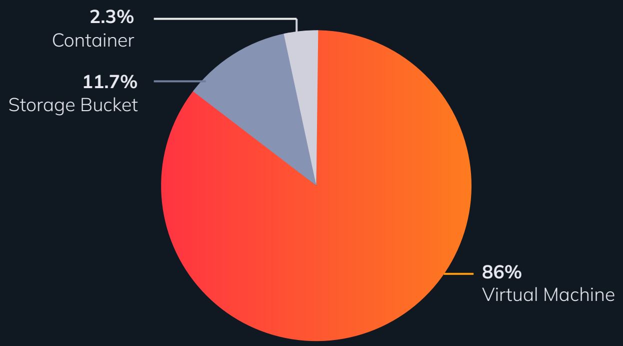

06 CI/CD Security

6. CI/CD Security

62%
of organizations have severe vulnerabilities (CVSS > 7) in code repositories.

Incorporating third-party and open-source components into applications is common, but carries risks. Nearly two-thirds of organizations encounter severe vulnerabilities in code within repositories like GitHub, Bitbucket, and GitLab. These vulnerabilities, potentially escalating in larger codebases, can lead to data breaches and system compromises.

Software Composition Analysis (SCA) is crucial for security, particularly for external code, including third-party and open-source elements. When integrated into the CI/CD pipeline, SCA helps identify problematic dependencies, preventing vulnerabilities in operational applications.

Implementing automated "Shift Left" practices in cloud security strategies is essential for early vulnerability detection in these repositories, safeguarding software integrity before deployment.

70%
of organizations have unencrypted secrets in code repositories.

Almost three quarters of organizations have unencrypted secrets, such as login credentials and API keys, in their code repositories. This lack of encryption poses significant risk of security breaches in cloud environments, where secret management is complex. Attackers can exploit these exposed secrets rapidly: Orca's 2023 Honeypotting in the Cloud Report found that it takes merely 2 minutes to exploit keys on GitHub.

Utilizing cloud provider secrets management services is crucial for effective secrets management throughout the development lifecycle.

Early detection of secrets and vulnerabilities enables a holistic security strategy, covering everything from initial code submission to ongoing monitoring in production stages.

07 Cloud Security By Industry

7. Cloud Security By Industry

1-5%

increase in industry average security scores in 2023.

Public Sector & Education improved by

5.2%

Financial services and healthcare

are the most secure.

While perhaps it is unsurprising that heavily regulated industries like financial services and healthcare providers scored strongest, it is noteworthy that all industry averages improved over time in 2023, an encouraging sign and strong evidence of the value that a comprehensive cloud security platform can bring.

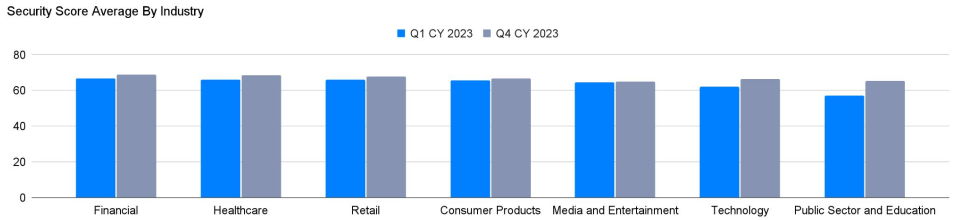

Orca calculates the security score based on performance in these five categories:

- Suspicious activity and detected malware
- IAM misconfigurations
- Exposed and vulnerable sensitive data
- Assets with critical vulnerabilities
- Level of responsiveness to security risks

08 Key Recommendations

Key Recommendations

Based on our findings, we have summarized key recommendations on how to reduce your cloud attack surface and harden your environment.

Patch strategically

While identified assets with known vulnerabilities should be patched as soon as possible, it's important to know where your efforts will result in the biggest security improvement. Therefore it is important to understand the cloud risk context to know which vulnerabilities enable dangerous attack paths and then patch them first.

Don’t neglect workloads

Ensure that only software applications or operating systems that are currently supported and receiving vendor updates are added to the organization’s authorized software inventory. Eliminate unsupported workloads.

Maintain an updated cloud asset inventory

Maximize cloud security by opting for agentless-first solutions, which offer complete visibility and deep insights into all assets, avoiding blind spots common in agent-based systems. Prioritize protection for critical assets with sensitive data.

Principle of Least Privilege (PoLP)

Give users only the minimum level of access privileges that they need to do their job. Ensure that regular users don’t have the power to escalate their own privileges or create new accounts.

Maintain IAM hygiene

Monitor and deactivate unused identities and access keys after a set period. Revoke access for ex-employees. Implement Privileged Access Management approaches to provide just-in-time access for users and generate temporary credentials for third parties that only require access for a limited time.

Implement strong user authentication

Always implement Multi-Factor Authentication (MFA) where possible and use strong, unique passwords.

Key Recommendations (continued)

Know where your crown jewels are

Make sure you know where your most critical business assets are located in the cloud. This includes sensitive shadow data that cloud security teams may not even be aware of. Apply the most stringent security measures and prioritize blocking attack paths to these assets.

Monitor, mitigate web and API risks

Implement robust security monitoring of domains and subdomains and regularly audit configurations to prevent mismanagement and misuse.

Leverage malware detection

Make sure the tool you use not only finds known malware by using signatures, but also uses heuristic scanning to detect unknown malware and zero-day threats.

Automated IaC configuration

Utilizing declarative Infrastructure as Code methods, as opposed to manual configuration and deployment, minimizes human error and provides an earlier checkpoint in the process for risk scanning and monitoring within your cloud estate.

Regular audits and continuous monitoring

Audit and enforce security policies to avoid misconfigurations. Monitor access logs, set alerts for unusual activities to swiftly detect and respond to active security threats.

Backup often and regularly

Regularly backup vital data, storing offline if feasible, and protect backups from deletion. This reduces ransomware impact by allowing system restoration without paying ransoms.

About Orca Security

Orca's agentless-first Cloud Security Platform connects to your environment in minutes and provides 100% visibility of all your assets on AWS, Azure, Google Cloud, Kubernetes, and more.

Orca detects, prioritizes, and helps remediate cloud risks across every layer of your cloud estate, including vulnerabilities, malware, misconfigurations, lateral movement risk, API risks, sensitive data at risk, weak and leaked passwords, and overly permissive identities.

To find out more, schedule a personalized demo of the Orca platform.

orca
security