WIZ

2022 cloud security threats

As more organizations move to the cloud, so do hackers. What can you do to better protect your cloud environment in 2022? Wiz Research has compiled the most pressing cloud risks and how you can protect against them.

January 2022 Report

Table of contents

The current cloud security landscape	3
Cloud security threats	4
Data exposure	4
Supply chain risks	5
Identity based	6
Software based	7
Cloud native threat campaigns	9
Cracks in the shared responsibility model	10
Cloud security checklist for 2022	11

The current cloud security landscape

Cloud adoption is growing quickly. However, such a rapid adoption rate of new technology has its downsides as well. Before addressing the risks to the cloud, let's first understand the current cloud security landscape and complexity:

Lack of visibility

Most organizations are not even aware of their entire cloud footprint – how many assets they actually have running in the cloud, who is using them, or whether they are vulnerable in any way. Lack of visibility is the most basic problem and a huge pain point for organizations, and one that only increases as cloud technologies are more widely adopted.

Lack of security awareness

We also need to understand that a quick adoption rate leaves no time to educate users properly about cloud usage best practices or the cloud-native risks that derive from misusage. For example, a developer can inadvertently push into the CI/CD pipeline an asset with secrets, that is then deployed to an externally exposed resource which is easily abused by hackers. Lack of user knowledge leads to increasing security risks, which are harder to mitigate in a complex, agile, and constantly evolving environment like the cloud.

While more and more organizations move to the cloud, so do the hackers. They are quicky adopting this new technology (much faster than the users are), and exploiting the lack of security awareness and cloud knowledge to wreak havoc.

Lack of responsibility

In the pre-cloud era, the responsibility for security was fully in the hands of the security teams, not the users. As we uncover new types of cloud vulnerabilities, we discover more and more issues that do not fit the current model. In the cloud, things operate differently and security does not fall on the shoulders of a single entity: it is a delicate balance. Some of these vulnerabilities require a unique remediation process with varying responsibilities from CSPs and customers, and there is no standardized way for addressing cloud vulnerabilities. Other vulnerabilities originate from developers using the cloud independently disregarding the security team and guidelines.

"Gartner Forecasts Worldwide Public Cloud End-User Spending to Grow 23% in 2021... Emerging technologies such as containerization, virtualization and edge computing are becoming more mainstream and driving additional cloud spending. Simply put, the pandemic served as a multiplier for CIOs' interest in the cloud."

Gartner Forecasts Worldwide Public Cloud End-User Spending to Grow 23% in 2021

Cloud security threats

This report is based on our research and experience from the past year in securing enterprise cloud environments. We chose to focus on novel, notable, and high-impact risks that we believe you should be familiar with and include in your cloud security strategy for 2022.

Data exposure

Challenge

Data exposure is still shockingly common. According to our research, over 55% of companies have at-least one database that is currently publicly exposed to the internet. Many of these databases use weak passwords or do not require authentication, making them an easy target for attackers who are continuously scanning the internet in search of such exposed databases. Considering that an unsecured Elasticsearch server can be breached in eight hours flat, such exposures must be fixed as soon as possible.

If we look at data breaches specifically, it is astonishing to see the scale of data breaches caused by exposed databases. Effectively, unintentionally exposed databases are one of the most common sources of data breaches and it is becoming a huge challenge for security teams to keep up, especially as cloud environments are becoming more and more complex. Here is only a partial list of breaches caused by misconfigured databases between 2020-2021:

Best practices

Visibility into all your cloud databases

- ☑ Review your PaaS databases and compute workloads, including: VMs, containers, and the software installed on them. Better yet, choose tools that enable you to do this efficiently.
- ☒ Configuration-based solutions have limited capabilities to provide the necessary visibility and cannot inspect or scan workloads and are not recommended.

Identify exposure from your cloud environment

- ☑ Choose exposure engines that have full visibility of your cloud environment in order to identify any routing or network services that allow traffic to be exposed externally. This includes load balancers, application load balancers, CDNs, network peering, cloud firewalls, etc. Furthermore, in order to assess external exposure from a Kubernetes cluster, the exposure engine must factor a large number of Kubernetes networking components including cluster IPs, Kubernetes services, ingress rules, and many more.
- ☒ CSPM tools focus on object configuration, preventing them from assessing exposure across multiple paths effectively.

Company	Breach Scale	When
Microsoft	250M records	January 22, 2020
Estee Lauder	440M records	February 11, 2020
U.K. Security firm	5B+ records	March 19, 2020
KeyRing	14M user records	April 6, 2020
Clubillion	200M records per day	July 7, 2020
VIP Games	23M records	January 26, 2021
Reverb	5.6M records	April 26, 2021
The Telegraph	10TB records	September 14, 2021

Supply chain risks

Supply chain risks are becoming a major attack vector, and possibly the most famous attack vector thanks to the notorious SolarWinds attack 2020 and the recent Log4Shell vulnerability in 2021 that affected thousands of applications.

To start with, just a quick look at supply chain attacks at a high level. Supply chain attacks occur when attackers gain access to a 3rd party vendor's software or environment. By breaching the 3rd party and inserting malicious code, it lets attackers get into more secure targets and get access to thousands of organizations that use that vendor solely through their one breach.

There are two major types of supply chain risks in the cloud:

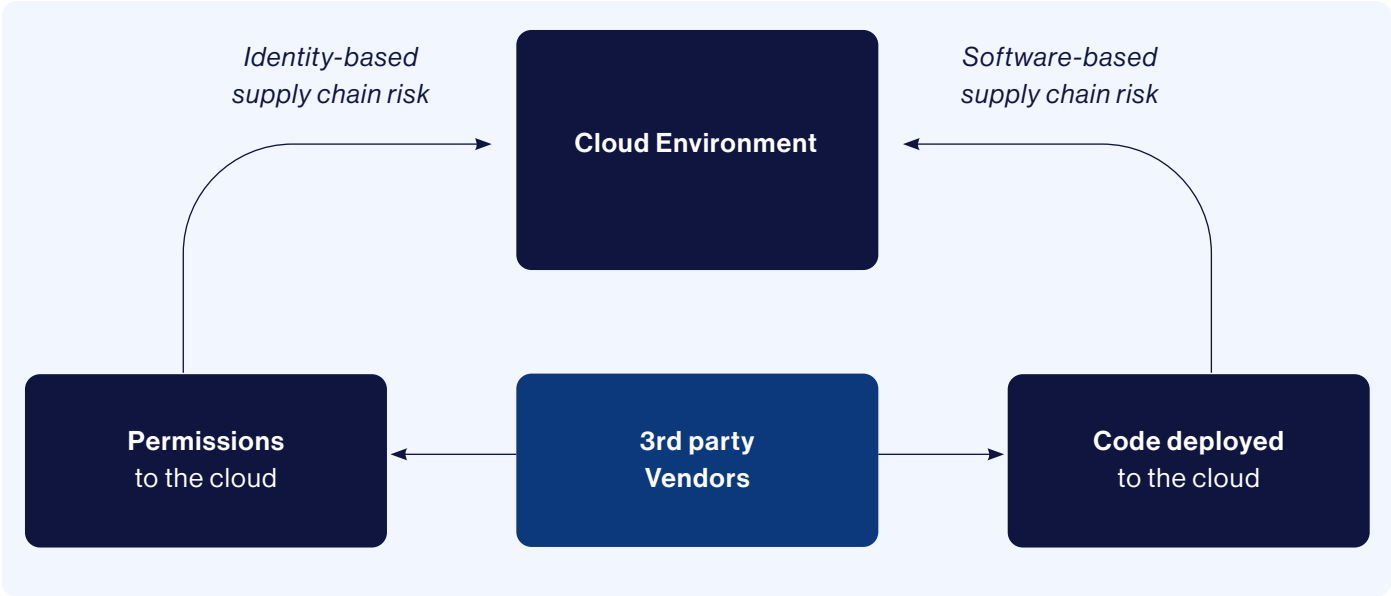

Identity-based supply chain risk challenges

- New type of risk in cloud via 3d party permissions
- Widely used in almost all cloud environments
- Lacking awareness and control by security teams

Software-based supply chain risk challenges

- Software risk in the cloud is much more complex
- Compute types: VMs, containers, serverless
- Lack of unified cloud asset inventory

Identity based

Challenge

This type of risk originates from granting 3rd party vendors permissions within your cloud environment, instead of installing 3rd party software. Cloud identity permissions are complex. When granted without proper security awareness, innocent looking permissions could lead to unintended exposure. For example, if a 3rd party service with high privileged access to your account is breached, then your data and infrastructure could be at risk too.

In 2021, Microsoft published that the Nobelium threat actor group (the one behind the SolarWinds hack in 2020) targeted multiple cloud service providers (CSPs) and managed service providers (MSPs) that have been granted administrative or privileged access by other organizations. Their goal was to leverage the trusted relationships to move laterally in cloud environments, from the service providers to their customers. This 2021 Nobelium campaign marks that threat actors are well aware of the fragile identity supply chain, and are aiming to gain access to downstream customers through other cloud service providers in order to carry out further attacks or access targeted systems. Such operations are extremely effective for attackers. They can obtain privileged access to thousands of targets by choosing their CSP/MSP target wisely.

To evaluate the potential magnitude of the issue, the Wiz Research Team conducted extensive research. The findings are alarming:

In most cases, vendors are granted over-permissive privileges. The most common example is the AWS "ReadOnlyAccess" policy, which is extremely popular amongst 3rd party vendors (a default for 25% of vendors included in our research). Vendors and customers believe it is a harmless policy, but instead it provides wide read access to many of your databases, such as DynamoDB, S3 buckets, SQS queues, and more.

Adding to this that security teams rarely monitor these permissions after they were granted, vendors can easily abuse them to do as they will, and continue doing so for a very long period of time.

Best practices

Reduce the permissions you provide to 3rd parties to the minimum that is required

- ☑ Certify any 3rd party vendor that requires access to your environment and ensure they have the most limited form of access that still lets them accomplish their function.

Identify and monitor all effective permissions within your cloud

- ☑ You need clear visibility of all effective permissions in your cloud environment.

82% of companies provide 3rd party vendors highly privileged roles

76% of companies have 3d party roles that allow a full account takeover

Over 90% of cloud security teams were not aware they gave high permissions to 3rd party vendors

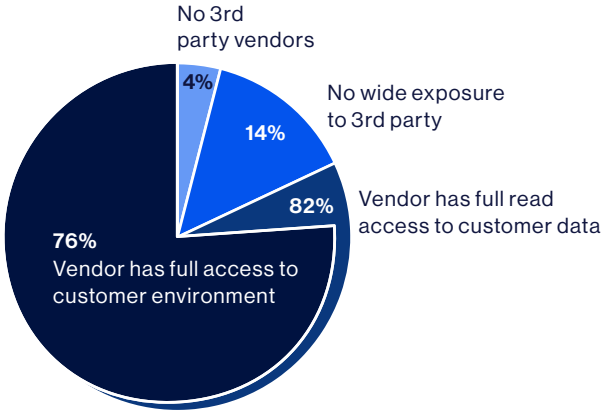

- The research covers 1,500 AWS accounts
- We focused on the 40+ most popular third-party vendors
- Analyzed roles using our identity and entitlements engine:
  1. Privilege escalation paths
  2. Permission boundaries
  3. Deny statement and NotActions
  4. Advanced AWS constructs like conditions and SCPs

Software based

Challenges

The cloud has an additional risk factor that originates from cloud providers themselves, who are also using pre-installed software on their workloads. Cloud providers often don't share what 3rd party software they are running for you, in your cloud environment. How can you protect your environment from software you are not even aware of?

Let's look at selected examples of this use case:

1. **OMIGOD: Azure customers are unknowingly exposed due a "secretly" installed agent**

   The Open Management Infrastructure (OMI) agent is an open source project. This agent is embedded in many popular Microsoft Azure services without the customer's knowledge, for example, when you set up a Linux virtual machine in you cloud.

   The Wiz Research Team discovered a chain of 4 critical/high vulnerabilities in OMI, dubbed as OMIGOD. Unless a patch is applied, attackers can easily exploit these vulnerabilities to escalate to root privileges and remotely execute malicious code (for instance, encrypting files for ransom). However, since most customers were not aware they were running an OMI agent, they did not apply the patch. In a sample of Azure tenants we analyzed, over 65% were unknowingly at risk. The magnitude of the risk, along with the lack of awareness to this silently installed agent, led Microsoft to develop an auto-update mechanism that patched the vulnerability on Azure services' machines.

2. **Log4Shell: Customers struggle to identify a highly popular compromised logging app in their cloud environment**

   Log4Shell (CVE-2021-44228), published in December 2021, is a critical unauthenticated Remote Code Execution (RCE) vulnerability in a highly popular Java library, Log4j. A day after the vulnerability was published it had already been exploited in the wild. According to Wiz and EY research more than 93% of enterprise cloud environments were vulnerable to the Log4j vulnerabilities. This vulnerability demonstrates how one critical vulnerability in a single library immediately and directly puts thousands of products and dozens of cloud services at risk.

   The first step for security teams is to identify all applications using Log4j running across their environments. In order to get full coverage, it is important to scan all workloads, including VMs running legacy apps, containers running on Kubernetes or other orchestration platforms, and even serverless code running on cloud functions. Since Log4j can be deployed as a package or embedded into the app itself, the scanner must support both.

Hackers are exploiting faster than we patch. Both OMIGOD and Log4Shell were being exploited in the wild just a few hours after publication, which is a lot quicker than it takes security teams to patch the environment. According to a Wiz research, organizations patched only 45% of their vulnerable cloud resources by day 10

3. **Shift-left attacks: New risks affecting the CI/CD pipeline**

   As the developers shift-left, so do malicious attackers. A new vector of supply-chain attacks is now targeting resources through infecting packages in popular package managers, such as:

   - PyPI — The PyPI open source package manager was used to host malware with multiple evasion techniques.
   - npm libraries — An increase in attempts to compromise workloads through npm libraries that were used to install different types of malware, including cryptominers.

   This year, we've also seen how insecure use of secrets in your code leads to multiple breaches in git repositories. For example, attackers exploited a flaw in the Codecov analysis tool to gain access to their git credentials and in-code secrets. In another case, attackers accessed an unintentionally exposed git repository to obtain database secrets that allowed them to access private United Nations employee

data. Both cases demonstrate why secrets should not be hardcoded into your code when possible. To mitigate such issues before they happen and detect potential vulnerabilities in your code or dependencies, use proprietary tools to scan the code and alert on security issues before the code is deployed to production.

Best practices

Proactively improve your cloud security hygiene

- ☑ Remove all secrets, passwords, and sensitive data. For example, use secret storage services (Key Vault), keep your git repositories clean, and refrain from using production credentials in the CI environment where possible.

Maintain a complete cloud asset inventory

- ☑ Security teams must have visibility into all 3rd party software assets and continuously analyze them to ensure that they are properly patched and performing as expected.
- ☒ Without such an inventory, it's easy to miss insecure third-party software in your cloud environment.

Operationalize a shift-left cloud strategy

- ☑ Fix vulnerabilities and misconfigurations in the CI/CD pipeline before deployment. Choose tools that scan your entire cloud stack, across any architecture: VM/images and Container/images for misconfigurations, vulnerabilities, network, and IAM, both at runtime and in the CI/CD pipeline.
- ☒ Silos between the pipeline and runtime cause a fragmented view of the security posture, but even worse, this disintegration extends across architectures where different policies are set up to control Infrastructure-as-Code (IaC), containers, PaaS, etc.

Cloud native threat campaigns

Challenge

Attackers also shift to the cloud, as more of their targets move there. We see more actors that target cloud accounts and workloads and leverage cloud native features for execution, privilege escalation, data exfiltration, and more – taking advantage of the fact that many cloud environments are not protected well enough due to security knowledge gaps. According to our research:

- **70% of cloud resources are not protected by any endpoint protection product.**
- **An average enterprise environment has 41 malware instances on average.**

Two capable malicious actors demonstrated how actors embrace cloud native capabilities in 2021:

1. TeamTNT, who recently targeted exposed Docker APIs to execute their own malicious code with root privileges on a targeted host. This group has been known to specifically target containers and cloud instances using Docker, Kubernetes, or Amazon Web Services that have been misconfigured and are publicly exposed to the internet. The group uses their access to deploy cryptocurrency miners on target workloads. Vx-underground, the largest open malware samples website, published some of the team's tools and confirmed that other ransomware groups have already adopted them.

Agents come in many flavors and forms. Regardless of the purpose they serve, an agent is yet another resource running in your environment. Besides performance and maintenance considerations, there are also security considerations when you run a 3rd party product. OMIGOD was a reminder that agents, even from reputable companies, can be installed with high privileges and come with vulnerabilities ranging from Local Privilege Escalation (LPE) to unauthenticated Remote Code Execution (RCE). Another recent example is CVE-2021-34522, a remote code execution vulnerability in Microsoft Defender that compromised your environment instead of protecting it.

2. Nobelium, the same actor we've mentioned in the identity-based supply chain attacks. As Microsoft reported, the actor targeted CSPs and privileged service accounts in order access their downstream customers. The TTPs (Tactics, techniques and procedures) that Microsoft detected indicate this actor is familiar with Azure services and Azure environments architecture. For example, ROADTools and AADInternals were used to enumerate Azure AD users, service principals credentials were created to enable persistence or additional access, and Azure RunCommand paired with Azure admin-behalf-of (AOBO) were used to obtain access to on-premises environments from the cloud. This cloud native knowledge of the APT group demonstrates how threat actors become more familiar with the cloud, just like their targets. We saw new cloud TTPs this year, and defenders should keep following them to ensure their hunting efforts are effective.

Best practices

Monitor & assess

- ☑ Continuously assess the risks to your environment to ensure coverage of all workloads with malware detection capabilities.
- ☑ Monitor cloud activity and network logs to detect suspicious activities

Harden

- ☑ Patch resources, maintain a least-privilege approach, and reduce unnecessary internet exposure

Cracks in the shared responsibility model

Challenge

The Wiz Research Team has discovered and disclosed several critical vulnerabilities this year – such as AWS cross-account vulnerabilities, ChaosDB, OMIGOD, and we've found that these vulnerabilities don't fit into the shared responsibility model, and enumeration and response are falling flat as a result.

Each of the vulnerabilities required a unique remediation process with varying responsibilities from CSPs and customers. We realized that currently there is no standardized way to address cloud vulnerabilities. Unlike other vulnerabilities that require user intervention like software vulnerabilities where we have CVEs, these cloud vulnerabilities don't have any identifier or enumeration, no standard format, no severity scoring and no proper notification channel. The response actions are a mix of efforts from the CSP and the user. Lack of responsibility and clarity around this handoff for cloud vulnerabilities is leading to missed opportunities and decreased security.

Best practices

A centralized database for reporting, enumerating, and remediating cloud vulnerabilities.

- ☑ Band together to put pressure on our CSPs to request CVEs for cloud services. We need to ask for more transparency, identification, and severity information for each vulnerability.
- ☑ In the meantime, refer to the public GitHub list of security mistakes by cloud service providers (AWS, GCP, and Azure).

If you are convinced, we invite you to join our global Slack channel to further discuss and shape the Cloud vulnerabilities database initiative

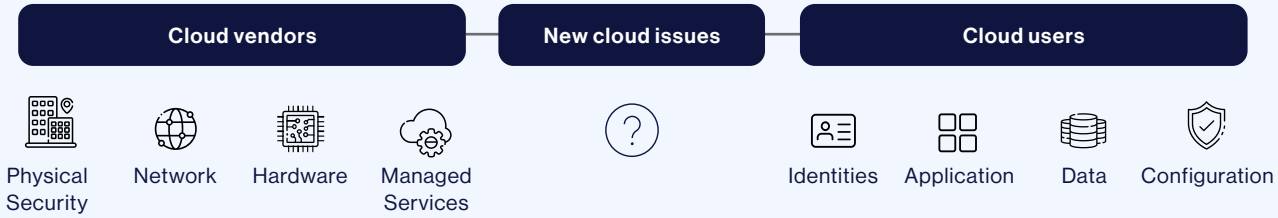

CSPs have security issues on their side as well, some of which require user intervention in order to patch (such as configuration issues or key leaks). However with no formal notification channel and no vulnerability and remediation instructions database, how can this be achieved? Currently, these vulnerabilities fall under a new "middle zone" in the model. Customers sometimes do not address the threat and leave their environment vulnerable, either because they are unaware of it or do not know how to fix it.

As we uncover more cloud-native types of vulnerabilities, we discover more issues that do not fit the current model.

Your cloud security checklist for 2022

Follow these guidelines to significantly improve your cloud security posture in 2022:

- ☐ **Increase visibility** into your entire cloud stack (including all VMs, containers, serverless, and PaaS) and maintain a complete cloud asset inventory to help you detect and patch vulnerabilities effectively

---

- ☐ **Monitor and assess** continuously to detect malware and suspicious activity

---

- ☐ **Minimize internet exposure** by identifying which resources are effectively exposed and uncovering unintentional exposure paths

---

- ☐ **Enforce least privilege principals to minimize the access 3rd party vendors have to your environment** by reviewing any 3rd party vendor that requires access to your environment and data, and ensuring they have the most limited access required

---

- ☐ **Shift-left your security strategy** by integrating your CI/CD pipelines with your security policies to resolve security issues before they are deployed

About Wiz Research

We are a group of cybersecurity veterans who have been in this space for over 10 years, work around the clock to identify potential security issues with CSP services, and help our customers to identify and address them. Sometimes, we find vulnerabilities in the platform itself. In 2021 alone, we have helped AWS, Azure, and Google Cloud to fix over 20 security vulnerabilities across multiple services. We also track emerging cloud security threats and add detections to Wiz Threat Center.

If you have any questions about our research, please reach out to research@wiz.io

About Wiz

We're on a mission to help organizations rapidly identify and remove the critical risks in their Cloud environments. Purpose-built for the unique complexities of multi-environment, multi-workload, and multi-project cloud estates, Wiz automatically correlates the critical risk factors to deliver actionable insights that don't waste time.

Wiz connects in minutes using a 100% API-based approach that scans both platform configurations and inside every workload. Our full security stack context surfaces the toxic combinations that show the attackers' view to a breach. Security and DevOps teams use Wiz workflows to proactively remove risks and prevent them from becoming breaches.

Visit us at wiz.io