



Cado Security Labs

2023 Cloud Threat Findings Report

Note from the Cado Security Founders

As experienced incident responders, James and I have provided crucial support to numerous large enterprises in their response to significant attacks. Time is of the utmost importance in incident response. And, as organizations increasingly adopted cloud technologies, we encountered growing challenges in assisting our clients with swift incident response. Traditional forensics tools and approaches were no longer sufficient, compelling us to seek a better solution. Our frustrations and personal experiences paved the way for the founding of Cado Security, where we developed a platform to revolutionize incident response for the cloud era.

At **Cado Security**, our mission extends beyond serving enterprises by offering a platform to facilitate efficient cloud forensics and incident response. As the founder and a threat researcher myself, my vision for Cado involved investing in initiatives aimed at empowering the broader security community. In pursuit of this goal, we established an internal threat research division dedicated to monitoring the latest attack trends and cloud-focused tactics, techniques, and procedures (TTPs). The following report provides a summary of our team's significant discoveries in 2022.

Our intention in sharing these findings is to equip fellow incident responders and security personnel with the knowledge they need to remain at the forefront of securing organizations.

Chris Doman

Cado Security, CTO & Co-Founder



By leveraging the insights contained within this report, we hope to foster a collective effort to safeguard both large enterprises and small businesses alike, ensuring we can all stay ahead in the ever-evolving landscape of threats.

Table of Contents

4	Introduction
5	Who is Cado Security Labs?
8	Cloud Attack Trends
9	Key Findings
12	Observations & Predictions
16	Conclusion & Recommendations
19	About Cado Security

Introduction

The cloud has become an integral part of modern business, but with its increased adoption comes an increased risk of cyber attacks and data breaches. Cado experts continuously track emerging cloud trends and this report delves deep into the noteworthy discoveries unveiled during the past year. To provide a concise glimpse into their findings, here's a brief summary of a selection of key insights uncovered by the Cado Security Labs team:



Attackers are rapidly evolving their tactics to newer cloud services.

In April 2022, we discovered [Denonia](#), the first publicly-known case of malware specifically designed to execute in an AWS Lambda environment. Although the first sample was fairly innocuous in that it only runs crypto-mining software, it demonstrates how attackers are using advanced cloud-specific knowledge to exploit complex cloud infrastructure. Further, it indicates that we are only scratching the surface of more nefarious attacks that may leverage a much wider range of cloud services.



Misconfigurations and associated credential theft have long been a leading cause of security breaches and that remains true in cloud environments, too.

Consequences of credential theft in the cloud thus far have been generally limited to spinning up cryptominers. However, in more recent research, we have seen attackers use credentials for a variety of much more nefarious purposes, such as stealing data for later attack phases like phishing or data exfiltration. As we have seen from recent reports, such attacks can be highly effective. For example, in [a recent incident](#) analyzed by the Sysdig threat research team, attackers successfully stole AWS keys which resulted in intellectual property loss. Interestingly, the attackers used cryptojacking as a distraction during the period of data exfiltration.



Internet-facing protocols and services remain low-hanging fruit for cloud threat actors.

For cloud-focused malware campaigns, internet-facing protocols such as SSH are commonly abused to propagate payloads. Because SSH is a protocol used across the internet, not just in cloud infrastructure, this finding is not entirely surprising. However, this means that SSH can pose an easy target if inadequately secured. We encourage SSH users to implement basic hardening. For example, disabling password authentication and restricting access to the service for predefined IPs is a best practice to consider. Cloud Service Providers (CSPs) make this easy via the use of technologies such as security groups (AWS) or VPC firewall rules.

Who is Cado Security Labs?

Cado Security Labs is the research and development division within Cado Security's engineering team, responsible for conducting industry-leading threat intelligence and cloud security research since their inception.

Cado Security Labs' analysis of the cloud threat landscape plays a pivotal role in driving the evolution of the Cado platform. The team actively contributes to the development of new features and product ideas, often prototyping them before transitioning them to the broader engineering team. This approach ensures that the Cado platform remains at the forefront of emerging cloud technologies, fulfilling its core purpose of streamlining incident response in the cloud. Through these advancements, security professionals gain the ability to investigate and respond to intricate attack patterns employed within cloud infrastructures.

The team's research also serves as the foundation for an array of valuable resources provided by Cado Security, including technical playbooks, cheat sheets, blog posts, conference talks and other content. By creating these materials, Cado Security endeavors to empower the security community with up-to-date knowledge of the latest trends and **Tactics, Techniques, and Procedures (TTPs)**.

Such achievements are made possible by the collaborative efforts of a diverse and highly skilled team. Each member of Cado Security Labs possesses a unique technological skill set, collectively contributing to the team's vast capabilities. **Here's an overview of the core capabilities of the Cado Security Labs team:**



Threat Intelligence

Cado Security Labs acquires threat intelligence data from a variety of custom sources, such as honeypots and client engagements. Frequently, work performed to establish such data sources contributes to the wider engineering effort, as these are typically complex engineering projects in themselves.

In addition to these custom sources, Cado Security Labs' threat intelligence engineers conduct routine monitoring of public malware and threat intelligence repositories. When new threats are discovered, they are analyzed by the Cado Security Labs team to understand their behaviors and indicators. These insights are then translated into detections that are built into the Cado platform.

By constantly updating and incorporating new threat intelligence, Cado Security Labs not only strengthens the security posture of Cado customers but also empowers the wider security community. Through this collaborative effort, the team strives to disseminate knowledge of the cloud threat landscape, enabling organizations and security professionals to better defend against evolving threats.



Malware Analysis

Once threat intelligence has been conducted, malware samples are quickly triaged and any novel malware is analyzed using a combination of off-the-shelf and custom tooling. This typically begins with an initial triage using a sandbox. If interesting TTPs or attributes are observed in the sample in question, Cado Security Labs malware analysts will move on to static analysis using a disassembler.

Malware samples or campaigns with a clear cloud focus are of particular importance to Cado customers. Any such samples are analyzed in-depth and their behaviors and indicators are documented and published for use by the broader security community.

The Cado platform supports malware detection through the use of pattern matching technologies such as YARA. The platform also has its own proprietary behavioral detection mechanism, allowing analysts to define malicious behaviors of both malware and human adversaries. Threat intelligence research directly informs the creation of detections for these technologies, allowing the Cado platform to alert users when such threats are discovered during evidence processing.



Research & Development

Cado Security Labs collaborates closely with Cado's engineering team to seamlessly integrate threat findings into the Cado platform. Leveraging their cloud-specific knowledge and advanced programming skills, the team frequently prototypes new features and enhancements based on threat intelligence projects or novel Tactics, Techniques and Procedures (TTPs) employed by cloud threat actors. An example of an engineering project heavily informed by Cado Security Labs' research is [VARC](#) - Cado's Volatile ARTifact Collector. VARC is a free tool available for use by the security community to streamline the process of collecting volatile data.

In addition, Cado Security Labs engineers are also responsible for maintaining repositories of proprietary detection rulesets. These rulesets contain malware and behavioral definitions which are then integrated into the platform, allowing for the detection of malicious behaviors and serving as key pivot points for analysts during an investigation.

The task of detection engineering extends beyond the mere creation of detection rules; the team is also responsible for ensuring the ongoing effectiveness and relevance of detections as complex malware campaigns and attack patterns evolve. Cado detection rules undergo continuous revisions to adapt to the evolving threat landscape. In addition, rigorous testing mechanisms are implemented to minimize false positives and identify any potential regressions.

This aspect of the Cado platform holds immense significance, as solid detection engineering is paramount to providing users with the ability to quickly pivot an investigation based on key malicious activity and gain an in-depth understanding of cyber security incidents.

Cloud Attack Trends

Background

Cado Security Labs operates honeypot infrastructure across four distinct geographical regions for the purposes of collecting cloud attacker telemetry. In each region, a sensor node runs a number of services known to be targeted by cloud-focused threat actors. With this data, Cado Security Labs researchers can examine attack patterns in near real-time and use information derived from this approach to identify and report on emerging cloud security threats.

The honeypot sensors run a variety of services. For example, Cado previously reported on the targeting of Redis by groups like [WatchDog](#). Similarly, threat actors such as [TeamTNT](#) are known to target the Docker Engine API and abuse it to spawn malicious containers. These services have been defined as initial access vectors in a number of cloud-specific attacks.

Key findings from attacker telemetry which we will cover in more detail include:


- SSH is the most commonly targeted service
- Opportunist threat actors generally scan for vulnerabilities in a single service
- Threat actors continue to prioritize botnets and DDoS attacks



Key Findings

SSH is the Most Commonly Targeted Service

Number of Unique IPs per HP Service



Internet-facing protocols and services remain low-hanging fruit for cloud threat actors

Of the services simulated by Cado Security Labs across honeypot sensors, SSH was by far the most commonly targeted. Since SSH is a protocol used across the internet, not just in cloud infrastructure, this statistic is unsurprising. SSH allows secure communication between clients and servers, and is typically used for server administration. This often means that SSH servers are internet facing and can pose an easy target if inadequately secured.

Accounting for just over a quarter of all traffic to the Cado Security Labs honeypot is traffic destined for Redis. Redis is an in-memory data store and is frequently deployed as part of a distributed application in cloud environments. The developers of Redis strongly discourage exposing the data store to the internet, as their [security model](#) is designed with trusted clients in mind. Despite this, Redis is frequently [observed](#) as an initial access vector for cryptojacking groups, such as 8220 Gang, TeamTNT and WatchDog.

A mere 4.3% of traffic was defined as scans for Log4Shell, a vulnerability in Apache's Log4J logging library (tracked as [CVE-2021-44228](#)) that made headlines in December 2021. Shortly after discovery and public disclosure, Log4Shell was almost immediately exploited by threat actors to distribute [ransomware](#).




For cloud-focused malware campaigns, SSH is commonly abused to propagate payloads. This is done either through enumeration of related hosts via the `known_hosts` file or by leveraging scanners such as [zmap](#) and [pnsan](#). Campaigns from threat actors such as [Diicot](#) rely heavily on SSH misconfigurations.

The low levels of Log4Shell traffic seen by Cado's infrastructure indicate that cloud-focused threat actors are no longer prioritizing this vulnerability as a means of initial access.

This could be due to the high levels of press coverage that the vulnerability received at the time, with multiple private and public sector organizations providing guidance on how to remediate it for users.

Opportunist Threat Actors Generally Scan For Vulnerabilities in a Single Service




Opportunistic threat actors generally pick a service and conduct mass scanning to identify vulnerable instances of it deployed in the wild.

Of the traffic reaching Cado's honeypot sensors, the vast majority (97.5%) scanned for a single specific service. This concurs with Cado Security Labs research, as most campaigns analyzed internally rely on a sole initial access vector.

In a small minority of cases (1.9%), the same source IP was observed conducting scans for more than one service. In our experience, threat actors will target one specific service rather than multiple (for example, SSH alone versus SSH and Redis). This could be due to the fact that attackers are aware of a specific vulnerability in that particular service or they have development experience in that area.

Threat Actors Continue to Prioritize Botnets and DDoS Attacks



Analysis of traffic dedicated to the propagation of malware has resulted in the breakdown of categories seen in the pie chart above.

The vast majority of observed traffic is dedicated to spreading common botnet families, these include Mirai, XorDDoS and IRCbot - a generic name for botnets making use of the IRC protocol. It's worth noting that samples categorized as Mirai may actually be one of the many existing variants of this malware.

Such variants could be considered commodity malware at this stage, since threat actors actively repurpose the Mirai source code and adapt it to their needs. This has resulted in families such as Cayosin and Qbot, which are sold and redistributed as a service. **From the telemetry analyzed by Cado Security Labs, it's clear that threat actors still place significant value in botnets and their usefulness in DDoS attacks.** This has been especially evident throughout the Russia-Ukraine war, where vigilante groups have carried out DDoS attacks in retaliation to [real-world events](#).

The remaining categories consist largely of generic trojans, shell script downloaders and cryptocurrency miners. Cado Security Labs has written extensively about cloud-focused [cryptojacking](#) groups and from this summary of attacker telemetry, they remain a prevalent threat. Such groups typically deploy stagers and initial access payloads in the form of shell scripts against misconfigured services. These payloads then prepare the compromised system for mining and retrieve additional payloads for propagating the malware and carrying out additional objectives.

From the telemetry analyzed by Cado Security Labs, threat actors still place significant value in botnets and their usefulness in DDoS attacks.



Observations & Predictions



Serverless Function Misuse

In April 2022, Cado Security Labs researchers [discovered](#) the first malware specifically targeting serverless environments. The malware was named Denonia, based on a string used in the command and control infrastructure, and was determined to have been written specifically for execution in AWS Lambda. Since then, a number of related [samples](#) have been found, suggesting this campaign is ongoing.

Prior to the discovery of Denonia, at least one [instance](#) of serverless targeting had been discussed publicly, it's likely that many more went unreported. To date, all of the publicly-discussed incidents have involved compromising serverless functions for cryptojacking purposes. This is an obvious use-case for unauthorized hijacking of resources.

However, as attackers gain familiarity with cloud infrastructure, it is expected that serverless functions will be leveraged for more destructive purposes. A similar evolution has been observed in the objectives of cloud threat actors such as TeamTNT, whose primary objective was cryptojacking before they widened scope to include [credential exfiltration](#) and the deployment of botnet malware.

Serverless functions remain a fruitful target for cloud threat actors, and a pain point for defenders, due to their ephemerality and the lack of observability into serverless execution. Often administrators are unaware of serverless functions in their estate being compromised until they receive an unexpected bill from their cloud provider.

Cado Security Labs Prediction

Attacks leveraging serverless functions will increase in severity and sophistication. Cloud-native threat actors will continue to invest effort into compromising serverless functions to utilize scalable infrastructure and evade detection.



Ransomware Groups Focus on Linux and ESXi

A worrying trend over the past year is an uptick in the discovery of ransomware samples targeting Linux and VMware ESXi servers. Previously, high-profile ransomware groups concentrated efforts on Windows desktops due to the prevalence of Windows in commercial environments. This appears to be changing as ransomware groups develop payloads that are cross-platform from the outset or are ported from Windows to Linux. Modern compiled languages, such as Rust and Golang, facilitate cross-platform compilation and have been heavily-utilized by ransomware groups.

One such example of non-Windows targeting is the recent (February 2023) ESXi ransomware - [ESXiArgs](#). This ransomware family exploits a heap overflow vulnerability ([CVE-2021-21974](#)) in OpenSLP (a service discovery protocol used by ESXi) to implant ransomware and encrypt system files required by ESXi guests, rendering the guests themselves useless.

In July 2022, Kaspersky [reported](#) on the emergence of Luna - a Rust-based ransomware family targeting Linux servers. Kaspersky researchers note that Luna appears to have been cross-compiled for Windows and Linux from the beginning, suggesting that the developers had Linux systems in mind when originally conceptualizing the malware.

Similarly, NAS manufacturer QNAP disclosed [details](#) of a campaign carried out by the DeadBolt ransomware group to infect QNAP NAS devices. The group exploited a vulnerability in a QNAP feature known as Photo Station to deliver the ransomware and encrypt files stored on the NAS. The internet-facing nature of such devices, and their use for on-prem data storage in small businesses, makes them a sensible target for smaller or emerging ransomware groups.

Considering the speed of adoption of cloud technologies and the fact that cloud infrastructure is underpinned by Linux systems, ransomware groups targeting Linux is understandable. ESXi servers are also key assets for many organizations and hypervisor guests will often be running mission-critical applications. Disruption to these services by ransomware is costly and destructive.

Cado Security Labs Prediction

Ransomware groups will continue to develop non-Windows ransomware for hypervisors like ESXi. This widening of scope is also likely to involve targeting of Linux-based cloud infrastructure in the future.



Threat Actors Hijack Cloud Services for Spamming

Cado Security Labs recently [reported](#) on a new family of hacktools designed to harvest AWS credentials for the purpose of SMTP and SMS abuse. The family of tools is known as Legion, and is distributed as a paid product on groups within the Telegram messaging service. Legion automates the process of exploiting misconfigured web applications running PHP frameworks such as [Laravel](#). The malware will hunt for publicly-accessible environment variable files (.env files) hosted by these applications, and run a series of searches over their contents to retrieve credentials for various cloud and SMTP services.

If AWS credentials are discovered in these environment variable files, Legion will attempt to use these to gain persistence in AWS environments by creating a malicious user account and assigning the AdministratorAccess managed policy to it. This managed policy allows access to all services available to the AWS account and includes access to the console. Legion then proceeds to set up [Amazon Simple Email Service](#) (SES) and will even send a test email to confirm access.

Tools such as Legion are highly desirable for groups engaged in phishing attacks or general spamming. Using another organization's AWS account to carry out such attacks confuses attribution and avoids having to pay for and manage infrastructure for the purpose of carrying out nefarious actions. Once this activity is identified and remediated, Legion then automates the process of discovering a new target and the attackers can easily pivot to another account.

Cado Security Labs Prediction

Attackers will continue to see the benefit of exploiting the cloud. Moving beyond compute functionality, we anticipate attackers will increasingly leverage services like SES that can aid phishing and spam campaigns.





Conclusion & Recommendations

As we look towards the future, it's important for organizations to consider the evolving trends that shape the threat landscape, and take steps to secure their systems and protect against these types of threats. With the rapid migration to the cloud, it's critical that organizations invest in evolving their current security strategies to ensure the same level of security is extended to the next-generation of technologies.

As organizations gear up to secure the cloud and enhance their ability to efficiently identify, investigate, and respond to cloud threats, we urge them to consider a few key predictions for the upcoming year:

- **Threat actors will continue to leverage Botnets as a means to execute DDoS attacks.**
This has been especially evident when monitoring hacktivist groups throughout the Russia-Ukraine war. Although it requires relatively low skill, DDoS attacks are disruptive for organizations and can act as a decoy for more serious attacks.
- **There will be an increase in cross-platform ransomware.**
Over the past year, we've seen an uptick in ransomware samples targeting Linux. Since cloud services are largely based on Linux, this is likely to result in increased targeting of the cloud by ransomware groups.
- **Attackers with a range of motives will utilize the cloud to carry out their objectives.**
As noted, we've seen spamming groups targeting and leveraging cloud resources and native capabilities.
- **Attackers will continue to invest significant effort into cryptojacking campaigns.**
Often the same groups deploying botnet malware will also deploy miners.

In addition, there are several key recommendations we believe should be considered by organizations to ensure effective and efficient incident handling in the cloud:

- **Understand the Cloud Service Provider (CSP) shared responsibility model.**

The first step toward security in the cloud is to carefully understand that security and compliance is a shared responsibility of the cloud service provider and its customers. To summarize, customers are responsible for everything “IN” the cloud, whereas the CSP is responsible for security “OF” the cloud. For example, customers are responsible for maintaining security of their own data, operating systems, network and firewall configurations, identity and access management, and more. On the other hand, the CSP is responsible for securing the overall hardware and global infrastructure.

- **Ensure you have access to the right data.**

Data capture in the cloud is different from traditional infrastructure, therefore it's important to ensure you have enabled sufficient logging to support an incident investigation. Additionally, you should regularly test that you are able to acquire, process and analyze data from CSP log sources, disk images and other key artifacts from deployed resources. Another key point is to ensure you have the right level of automation in place so that you are able to capture evidence from ephemeral environments before it disappears.

- **Avoid unnecessarily exposing services like Docker and Redis to the internet.**

Performing regular scans of your entire estate, both cloud and on-premises, for exposed services that could be leveraged by an attacker is key. If any are found, then you should perform a triage check to determine if they have been compromised.

- **Check public repositories for any cloud credentials.**

Breaches of cloud accounts are typically from credentials being found from publicly accessible repositories such as github. If any exposed credentials are found, then you should take the necessary corrective actions and perform an investigation to determine if they have been used by an attacker.

- **Implement principle of least privilege.**

By implementing the principle of least privilege, you can protect higher value resources if lower value resources are compromised. Lateral movement is a key stage of the attack lifecycle and can result in significant damage in the event an account breach does occur.

About Cado Security

Cado Security is the provider of the first cloud forensics and incident response platform. The platform leverages the scale and speed of the cloud to automate the end-to-end incident response process – from data capture and processing to investigation and response. Cado enables security teams to gain immediate access to forensic-level data in multi-cloud, container, and serverless environments. Evidence items extracted from cloud-provider logs, disk, memory and more, are processed in parallel to drastically reduce time to investigation. The platform was built to empower security analysts of all levels by automatically highlighting the most important events related to an incident including its root cause, scope and impact. Cado also supports remediation actions so that organizations can quickly contain active threats.

If you're interested in learning more, [contact us](#) or take advantage of a [14-day Free Trial of the Cado Platform](#).