Cado Security Labs

H2 2023 Cloud Threat Findings Report

CADO

Table of Contents

3 Introduction

4 Who is Cado Security Labs?

7 Key Technical Findings

16 Observations

18 Conclusion & Recommendations

19 About Cado Security

Introduction

We are very excited to deliver our second cloud threat findings report which provides a detailed overview of significant discoveries made by the Cado Security Labs team over the last six months. The goal of this report is to help security professionals remain at the forefront of securing organizations, as its content is based on real-world techniques employed by attackers to target cloud-based environments. As commercial adoption of cloud technologies continues, cloud-focused malware campaigns have increased in sophistication and number – a collective effort to safeguard both large enterprises and small businesses alike is key. We're stronger together.

Some key insights uncovered by the Cado Security Labs team which are analyzed in depth within this report include:

Attackers continue to exploit web-facing services in cloud environments

A common theme in the malware campaigns we've reported on over the past six months is the exploitation of various web-facing services to achieve initial access in cloud environments. This includes services like Docker, Redis, Kubernetes, and Jupyter, all of which are frequently deployed in the cloud. It has become increasingly evident that attackers have developed advanced technical knowledge around how these services are deployed and how to exploit them. An example of an emerging malware campaign targeting Jupyter Notebooks is Qubitstrike, which was discovered in October 2023.

Cryptojacking is no longer the sole focus of cloud attackers

The vast majority of cloud and Linux campaigns we analyze at Cado involve hijacking resources for the purpose of mining cryptocurrency, usually Monero. While cryptojacking is a very real and significant threat, we've started to see a diversification in objectives displayed by recent Linux and cloud malware campaigns. One such campaign that demonstrates this is Legion, a cloud-focused hacktool written in Python. Legion was first discovered in April 2023 and its primary objective is to automate the hijacking of various cloud SMTP services, for the purpose of carrying out mass-spamming attacks.

Rust malware continues to proliferate

Modern compiled programming languages, like Rust and Golang, offer a number of benefits to malware developers, such as allowing them to compile for multiple operating systems and architectures, without having access to the operating systems or architectures themselves. In last year's cloud threat findings report, we discussed the emerging trend of ransomware developers targeting Linux systems, a trend aided by the ease in which modern compiled programming languages allow cross-platform development. Recent cross-platform ransomware campaigns, such as Blackcat, have utilized Rust for their primary payloads. We've also seen Rust leveraged in other recently discovered non-ransomware campaigns, such as P2Pinfect.

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

Malware Analysis

Research & Development

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.

Key Technical Findings

Background

Cado Security Labs operates honeypot infrastructure across four distinct geographical regions, for the purpose of collecting cloud attacker telemetry. Traditionally, Cado honeypots consisted of simulated services, such as those found in T-Mobile's open source Tpot project. While such projects are excellent in that they allow you to quickly deploy believable honeypots for a number of services, their low-interaction nature soon became a limitation.

The latter half of 2023 saw the introduction of a new, more sophisticated, high-interaction honeypot system. This new system, dubbed "Cloudypots", allows Cado Security Labs researchers to honeypot real services quickly and in a safe manner. To achieve this, the system leverages OpenStack's KVM-based virtualization - allowing sufficient isolation between guest virtual machines and the hypervisor itself. For more Cloudypots implementation details, we highly recommend reading our blog on the topic.

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

- Attackers target cloud services that require specialist technical knowledge to exploit
- Docker is the most commonly-exploited "cloud-native" service for initial access
- Threat actors leverage hosting companies across the globe for their infrastructure
- Cloud/Linux malware campaigns continue to diversify from cryptojacking

Attackers Now Possess an Intimate Knowledge of Cloud Services like Docker, Redis, and Jupyter

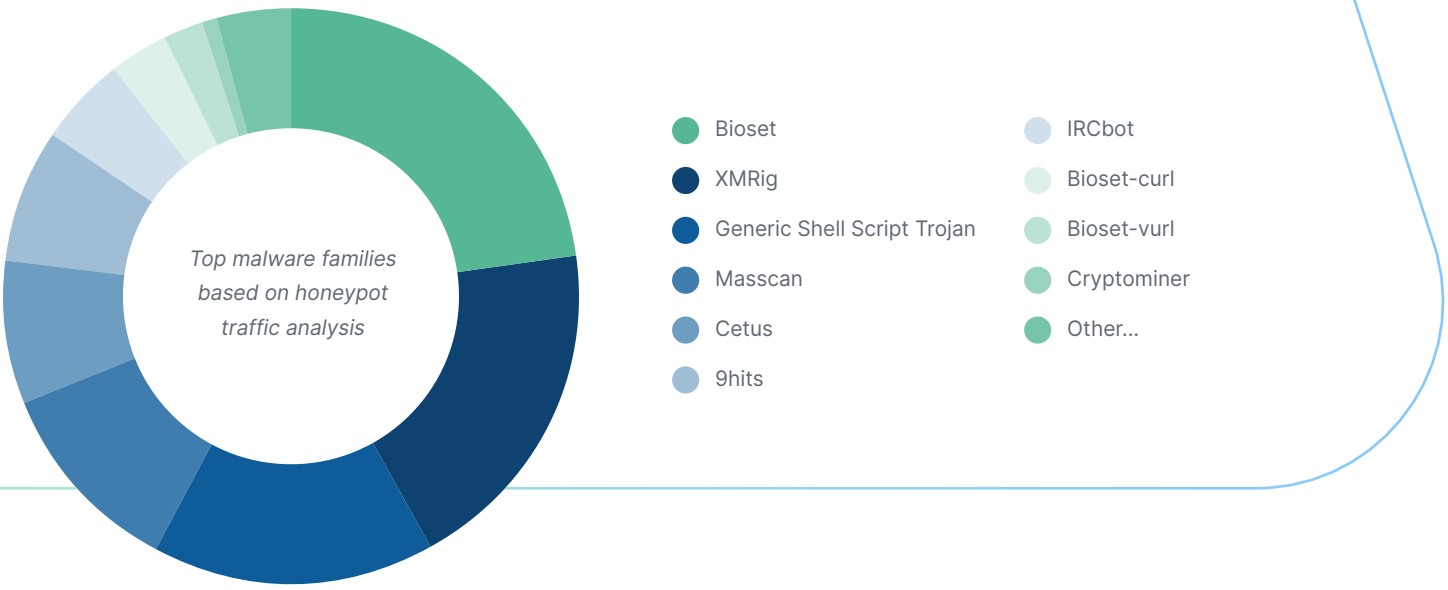

2023 was an interesting year for cloud-focused malware, and Linux malware in general. Attackers continued to target services like Docker and Redis, both of which are frequently deployed in cloud environments. Attacking these services requires specialist technical knowledge, different to what's required for attacking generic Linux servers.

Attacking these services requires specialist technical knowledge, different to what's required for attacking generic Linux servers.

We saw Docker-specific knowledge utilized against our honeypot, with attackers frequently attempting to escape created containers by mounting the host filesystem. In a similar vein, we also witnessed sophisticated exploitation of Redis, including exploitation of the Lua sandbox escape vulnerability (CVE-2022–0543), abuse of the data store's replication feature and continued attempts at conducting the well-known unauthenticated remote code execution attack. This reinforces statements we have made previously about attackers investing significant time into understanding these services and their weaknesses.

This year, we have deliberately omitted SSH traffic from our reporting as we highlighted it as a common attack vector in last year's report. The vast majority of honeypot traffic consists of opportunistic attempts to exploit servers, this year is no different. We opted to avoid covering such attacks, in favor of those that demonstrate a level of technical sophistication and target the services we, as cloud security professionals, are interested in.

Qubitstrike

Qubitstrike is a campaign that was discovered by Cado Security Labs researchers in October 2023. This campaign targeted insecure deployments of the Jupyter Notebook application, a web-based interactive computing platform commonly deployed in cloud environments. Attackers exploited the Jupyter service to spawn a terminal on one of our honeypot hosts.

After gaining access to the host, malicious payloads were retrieved from a Codeberg repository and executed. These payloads conducted a cryptojacking attack on the Jupyter host and attempted to exfiltrate credential files for well-known Cloud Service Providers (CSPs). The attacker then attempted to use one of these credential files to access what they believed to be an AWS account. Instead, it notified our canary service and generated a notification for us.

Qubitstrike established a number of interesting communication channels, including using the Telegram Bot API to exfiltrate the CSP credentials. The campaign also leveraged Discord as a C2 communication mechanism, with one of the payloads acting as an agent on the infected host. This agent allowed the attacker to send shell commands via Discord to the infected host, and the agent would return the output of the terminal session. All of this was viewable in an attacker-controlled Discord channel.

While not the first campaign to target Jupyter Notebooks, it's definitely not a common initial access vector. Qubitstrike highlights the lengths attackers will take to gain an understanding of web-facing services.

Docker Continues to be a Lucrative Target for Cloud Attackers

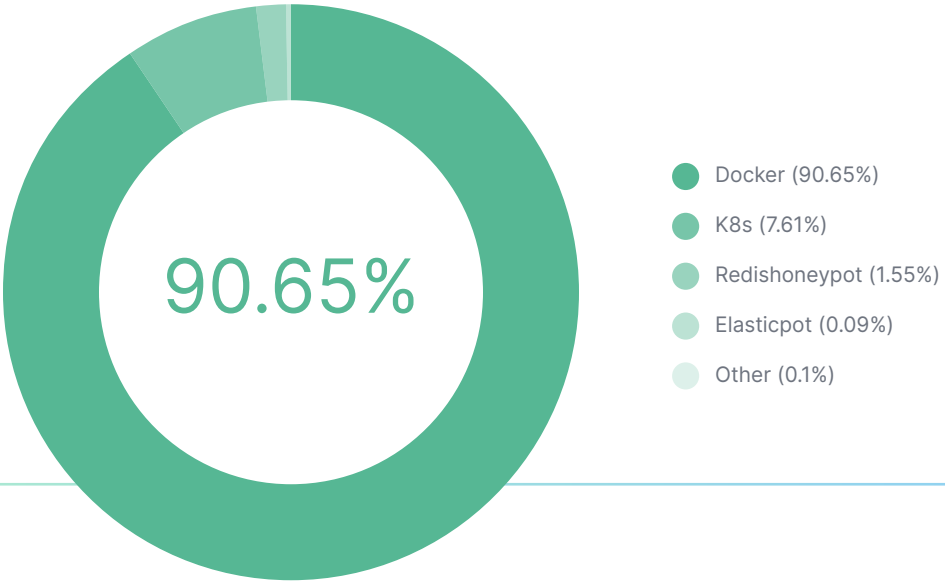

Although cloud-focused attackers aim to exploit a variety of services typically deployed in cloud environments, Docker remains the most frequently targeted for initial access (90.65% of honeypot traffic when discounting SSH). Several novel Docker malware campaigns were identified by Cado Security Labs researchers throughout 2023, including 9hits (detailed in a later section).

Although cloud-focused attackers aim to exploit a variety of services typically deployed in cloud environments, Docker remains the most frequently targeted for initial access.

2023 also saw the evolution of various long-running Docker-specific campaigns, including Bioset and Cetus, both of which abuse exposed Docker API endpoints to execute malicious code in the context of a Docker container. Clearly, attackers place value in their ability to compromise the Docker service for the purpose of carrying out their objectives.

Bioset

Bioset is a cryptojacking campaign first discovered in 2020. The Bioset name is taken from a legitimate Linux system process. When a malicious process is spawned by the attacker, it appears as “bioset” in process listings, making it look benign. The Bioset name is used in this manner across a number of different cryptojacking campaigns.

Due to their similarity, we have clustered these campaigns into a single detection, as they have a common objective of cryptojacking and target exposed Docker Engine API deployments. There has been an increase in the number of Bioset detections observed over the past year. Excluding generic SSH key upload events, Bioset initial access events are the most common sighting in our Cloudypots infrastructure.

Cetus

Cetus is another cryptojacking attack that targets misconfigured Docker Engine API deployments. It was also first discovered in 2020 and has undergone incremental updates since. In keeping with similar attacks, Cetus campaigns are particularly noisy and should be easy to detect.

This is largely due to the deployment of XMRig and masscan on the host, resulting in the increased resource usage typically associated with cryptocurrency mining, along with increased outbound network traffic as the malware scans a random subnet for new hosts.

Linux Attackers are Having More Success in Compromising Nodes Across the Globe

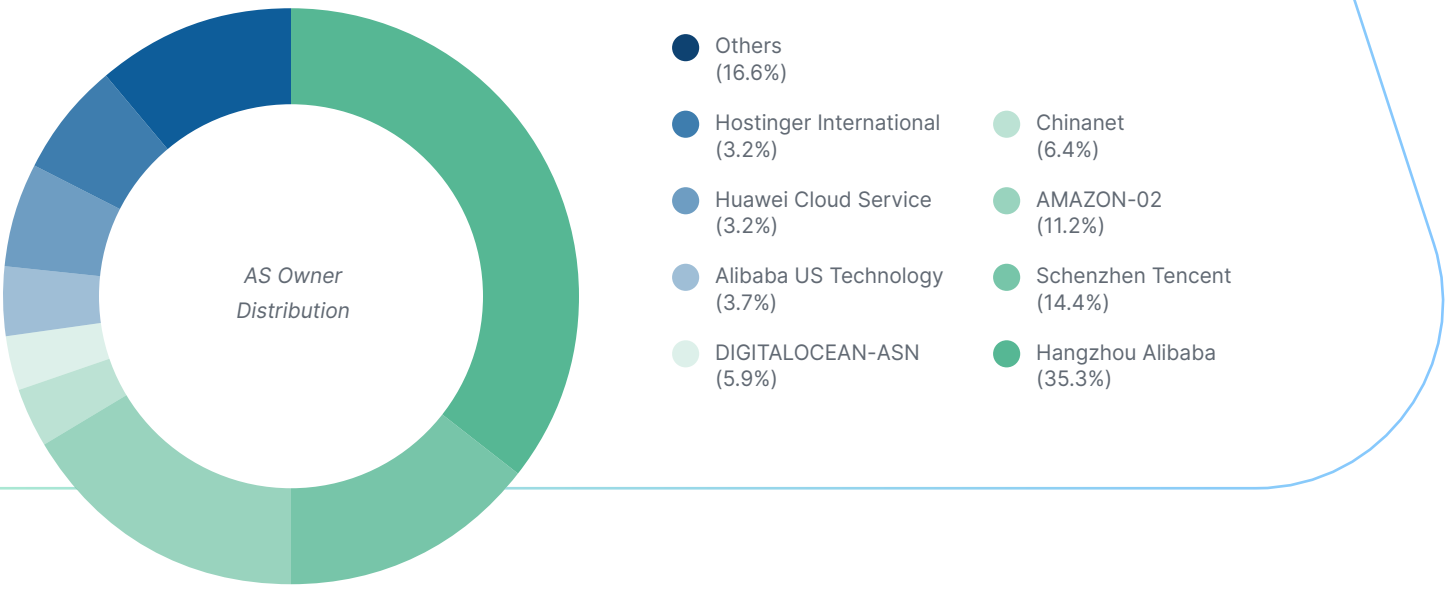

Evidently, cloud and Linux attackers are targeting services typically deployed in the cloud for initial access, but how successful are these attacks? The pie chart above shows the distribution of Autonomous System owners for IP addresses known to be infected with P2Pinfect.

As can be seen, the campaign has a wide geographical distribution with nodes belonging to providers in China, the US, and Germany. It's believed that P2Pinfect originated in China, due to the high concentration of infections there, but the above chart shows that the campaign has achieved success in compromising nodes across the globe.

For a more detailed geographical breakdown, please see the graph below.

Regardless of where your infrastructure is located, it is likely still susceptible to Linux and cloud-focused attacks.

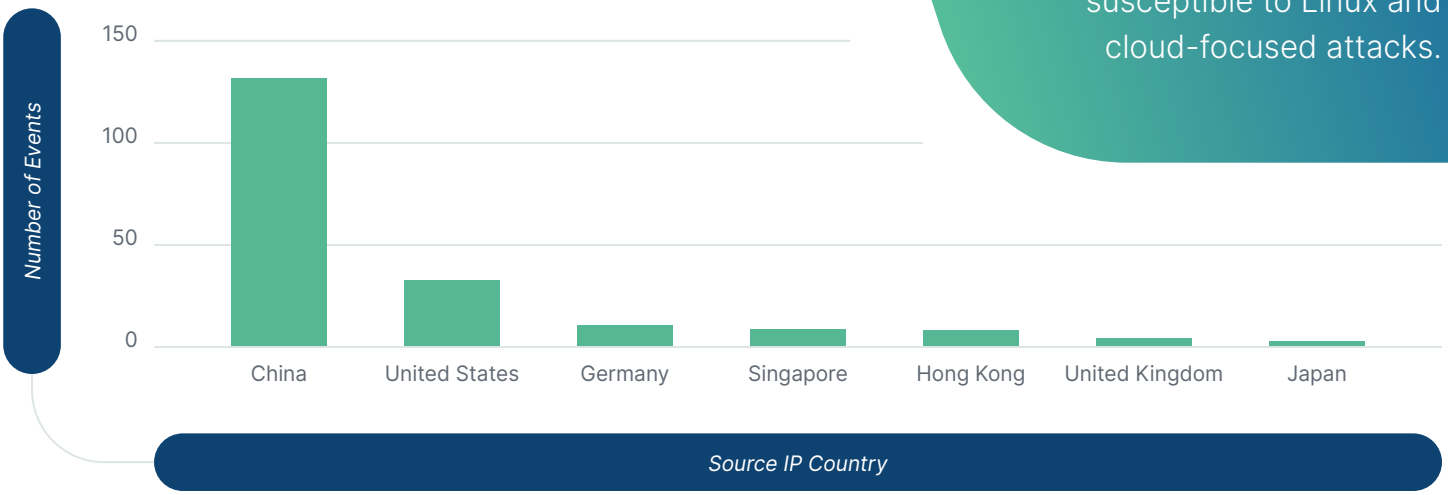

This demonstrates that regardless of where your infrastructure is located, it is likely still susceptible to Linux and cloud-focused attacks.

P2Pinfect

P2Pinfect is a cross-platform botnet agent written in Rust. The malware initially targeted insecure deployments of Redis, leveraging exploits and abusing native Redis features to gain a foothold in the target system.

Payloads delivered to the system demonstrated a high level of technical sophistication, including the use of multiple evasion techniques, the delivery of a monitoring agent capable of updating the main payload and a dynamic node list complete with heavy obfuscation.

Cloud & Linux Infrastructure is Now Subject to a Wider Variety of Attacks

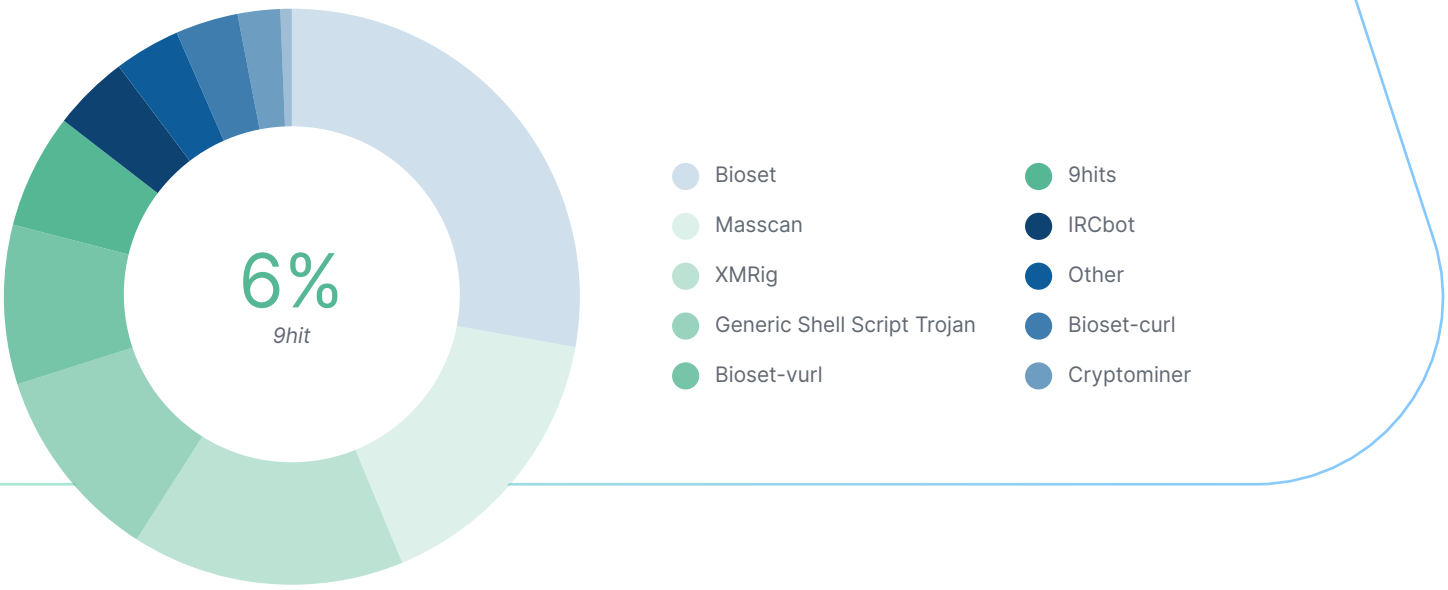

Despite a continued emphasis on targeting Docker and Redis to conduct cryptojacking attacks, the latter half of 2023 demonstrated that attackers are diversifying from this objective.

In last year's cloud threat findings report, we discussed the worrying trend of ransomware on Linux and ESXi systems. This continued throughout 2023, with the discovery of new Linux variants of ransomware families such as Abyss Locker.

> With the discovery of new Linux variants of ransomware families, there is a worrying trend of ransomware on Linux and ESXi systems.

From a cloud perspective, campaigns with a variety of objectives were discovered throughout the course of 2023 and into early 2024. One such campaign was 9hits, where attackers hijacked Docker endpoints for malicious traffic generation.

Another example of a diversification in objectives is the P2Pinfect campaign. While the primary objective of P2Pinfect remains unclear, the malware has yet to exhibit any TTPs consistent with cryptojacking campaigns and has yet to be observed deploying a cryptocurrency miner.

9hits

9hit refers to a new campaign discovered by Cado Security Labs at the beginning of 2024. The campaign involves compromising Docker hosts for the purpose of deploying the 9hits viewer application. This application, essentially a headless version of Chrome, is used to generate traffic to websites of the attacker's choosing - resulting in the generation of credits for the 9hits traffic exchange platform.

It's expected that the attacker will either resell these credits, or sell their traffic generation services more generally to individuals. This is a similar monetization approach as the one taken by botnet developers, initial access brokers (IABs), and other cybercrime service providers.

This campaign is particularly noteworthy as it's the first time the 9hits platform has been implicated in a malware attack, suggesting attackers with a variety of objectives are now targeting Docker.

Observations

Attackers Continue to Exploit Web-Facing Services in Cloud Environments

Much like traditional on-premise cyber attacks, cloud attacks require an intimate knowledge of the target environment. Our research has continuously highlighted the fact that adversaries possess this knowledge and are willing to use it to their advantage.

A common theme in the malware campaigns we report on at Cado Security Labs is the exploitation of various web-facing services to achieve initial access in cloud environments. By this, we mean services such as Docker, Redis, Kubernetes, and Jupyter, all of which are frequently deployed in the cloud.

Attackers have developed advanced technical knowledge around how these services are deployed and, concerningly, how to exploit the services to achieve initial access. Often, after such a service has been exploited, the attacker will hunt for cloud service provider (CSP) credentials, and attempt to move laterally to other resources.

An example of this is the Qubitstrike campaign, discovered in October 2023. In this campaign, the threat actor exploited a Jupyter Notebook instance running on one of our honeypot sensors and spawned a Bash terminal using Jupyter's terminal feature. They then used this feature to run additional payloads on the underlying host, including a credential exfiltration script that specifically hunted for CSP credentials.

Since the credential exfiltration code contained hardcoded CSP credentials, this attack demonstrates that attackers believe exploitation of services like Jupyter can help them gain access to cloud environments. In addition, they are investing significant time into hunting for misconfigured deployments of these services.

Cloud and Linux Threat Actors Continue to Diversify From Cryptojacking

The vast majority of cloud and Linux campaigns we analyze at Cado involve hijacking resources for the purpose of mining cryptocurrency, usually Monero. This has led to cryptojacking becoming almost synonymous with Linux malware, giving the impression that it is the sole objective of most Linux malware campaigns.

While cryptojacking is a very real and significant threat, we've started to see a diversification in objectives displayed by recent Linux and cloud malware campaigns. One such campaign that demonstrates this is Legion, a cloud-focused hacktool written in Python. Legion was first discovered in April 2023 and its primary objective is to automate the hijacking of various cloud SMTP services, for the purpose of carrying out mass-spamming attacks.

Since its discovery, we've learned that Legion is part of a wider family of hacktools, with high-profile variants known as Fbot and AndroxGh0st. AndroxGh0st was recently the subject of an advisory released by the Cybersecurity & Infrastructure Security Agency (CISA), in which they discussed Indicators Of Compromise (IOCs) exhibited by the malware in a number of ongoing investigations both the CIA and FBI were involved in.

Clearly, attackers see cloud SMTP services as a valuable resource for carrying out spamming attacks. This is likely due to the scalability and speed offered by cloud services, the same reasons why legitimate organizations see the cloud as beneficial.

Rust Malware Continues to Proliferate

Modern compiled programming languages, like Rust and Golang, offer a number of benefits to malware developers. Primarily, these languages allow developers to compile for multiple operating systems and architectures, without having access to the operating systems or architectures themselves.

This has been of great benefit to ransomware groups, many of which have traditionally targeted Windows. In last year's annual cloud threat findings report, we discussed the emerging trend of ransomware developers targeting Linux systems, a trend aided by the ease in which modern compiled programming languages allow cross-platform development.

Much like Golang before it, Rust is gaining popularity with malware developers. Several recent cross-platform ransomware campaigns, such as Blackcat, have utilized Rust for their primary payloads.

Cado Security Labs researchers encountered their first Rust payload during analysis of P2Pinfect - a peer to peer botnet targeting Redis and SSH servers.

Since usage of Rust is a relatively new phenomenon in the malware development community, very few malware analysis tools are capable of handling Rust binaries in an effective manner. Similarly, due to how new the language itself is, many malware researchers and reverse engineers are unfamiliar with Rust internals. This is another likely reason why the language is chosen for malware development.

We anticipate that we'll continue to see malicious payloads developed in Rust, as the language gains popularity in general software development.

Conclusion & Recommendations

As organizations increasingly adopt cloud technologies, it is critical that security teams reassess their internal tools and approaches in order to ensure their ability to properly identify, investigate, and respond to emerging cloud threats.

With this report, we aim to help security professionals gain a better understanding of how attackers are exploiting cloud-based technologies, and in turn, enable them to build a more robust internal security program.

Here are a few key recommendations we believe should be considered by security teams to ensure effective and efficient incident handling in the cloud:

- Establish a policy of regularly reviewing the security of deployed services in your cloud estate, particularly if they include the services described here.
- Consider reducing your attack surface by only deploying public-facing services when necessary and making use of networking security features (security groups, etc.) that your CSP provides.
- Ensure you are collecting and aggregating logs from both your CSP's control plane and for the individual services you intend to run in your account. Establish periodic review and automated alerting for anomalies found in these log sources.

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 to see a demo.

See a Demo