Lacework

Lacework 2021 Cloud Threat Report

VOLUME 2

INDEX

Introduction

In this report we share what we have learned over the past three months in the ever evolving cloud security threat landscape. Our aim is to provide defenders, researchers, and anyone curious about cloud threats with actionable information they can take back to their organizations. In this report we discuss the most recent attacker techniques we have observed, share tools and processes we have developed, and also present some data on attacker trends. Happy Hunting.

Executive Summary

Organizations should start thinking of cybercriminals as business competitors. They are working hard to profit either directly (through ransom and extortion) or indirectly (by stealing resources). Over the past few months we saw many interesting techniques, but the most interesting trend is rising demand for access to cloud accounts. We see this evidenced by the sale of admin credentials to cloud accounts from Initial Access Brokers. We see continued increases in scanning and probing of storage buckets, databases, orchestration systems, and interactive logins. Additionally, multiple threat actors continue to invest in evolving their cybercrime campaigns targeting cloud services. All this coupled together shows the increasing threat to businesses today. Whether it's access to your data or access to your resources, there are multiple ways to capitalize.

Notable Attacker Techniques

Lacework Labs designs, builds, and tracks threat activity in a methodology based around the MITRE ATT&CK® techniques on top of our own expertise of adversary activity. This section details the report's most noteworthy Tactics, Techniques, and Procedures (TTPs). As reported in our Crimeware Findings section of this report, we've observed TeamTNT making use of compromised Docker Hub images for target initial access. While our disruption effort was successful in this case, we're certain more have yet to be discovered. These three TTPs are associated with this increasingly common threat.

User Execution: Malicious Image [T1204.003]

T1204.003 is when an attacker benefits from a target running a malicious image to facilitate execution, often as an initial access method. As noted by MITRE, "Amazon Web Services (AWS) Amazon Machine Images (AMIs), Google Cloud Platform (GCP) Images, and Azure Images as well as popular container runtimes such as Docker can be backdoored."

Persistence: Implant Internal Image [T1525]

Adversaries implant cloud or container images with malicious code to establish persistence after gaining access to an environment. AWS AMIs, GCP Images, Azure Images, and Docker can be implanted or backdoored. With the T1525 technique, the backdoored image is in a registry within a victim's environment. This could provide persistent access if the infrastructure provisioning tool is instructed to always use the latest image. Adversaries make use of this to automate the deployment of their image to retain access or expand capabilities.

Execution: Deploy Container [T1610]

As we observed with TeamTNT and others abusing Docker Hub, adversaries can deploy a container into an environment to facilitate execution or evade defenses during an intrusion. T1610 can be used in opportunistic cryptojacking and as an access method for a targeted attacker.

Lacework Labs Research

The Lacework Labs team continually conducts security research focused on risks and threats relevant to cloud services, containers and container orchestration systems, and new attack surfaces the public cloud exposes through services or deployment methods. In this volume of our cloud threat report, the team has observed and investigated new crimeware incidents, vulnerabilities, and attacker opportunities. Additionally, we believe in contributing back to the security community, so we're also sharing a collection of security practitioner focused educational guides. Lacework Labs has:

- Identified and disrupted TeamTNT's Docker Image abuse, analyzed the malicious use of the cpuminer utility, linked a new "Tsunami-Ryuk" malware variant to Keksec group, and detailed an 8220 Gang campaign using a new custom miner and IRC bot.
- Researched potential Canary Token abuse to aid ransomware attacks, shared an analysis on Initial Access Brokers (IAB) focused on cloud networks, and identified a collection of vulnerabilities in XMRigCC.
- Released educational materials and tools automating the analysis of Katien/Tsunami IRC bot variants, and the process of threat hunting SSH keys in malicious bash scripts.

Crimeware Findings

The crimeware threat landscape continues to thrive into mid 2021, and our recent observations showcase various interesting research projects initiated through intrusion investigations and proactive threat monitoring.

As you may recall, in our inaugural volume of the Cloud Threat Report, we dove extensively into the adversary commonly known as TeamTNT. More recently, in May of 2021 we observed TeamTNT target exposed Docker APIs to deploy malicious images. During our investigation, we discovered Docker images containing TeamTNT malware were being hosted in public Docker repositories accomplished through malicious account takeovers. The research led us to multiple cases in which TeamTNT leveraged exposed Docker Hub secrets on GitHub to abuse for staging the malicious Docker images. In the end, Lacework Labs was able to contact the owners of compromised accounts, in addition to the Docker Hub security team, to take down the abused accounts. For more on this story, see Taking TeamTNT's Docker Images Offline.

Lacework Labs has also made additional discoveries related to popular cloud relevant crimeware and actors, such as cpuminer-linked malware and Keksec group.

Cpuminer, the open-source multi-algorithm miner, has been legitimately used for years. However Lacework Labs has observed an increase in its illicit use for cryptojacking altcoins. Our research dives into the cpuminer forks and details activity in the wild which includes propagation via Jupyter command execution and a variant of the WSO webshell, which infected numerous wordpress installations.

Keksec, also known as Necro and Freakout, is now leveraging a new Tsunami DDoS malware variant named "Ryuk." Note this is unrelated to the popular Ryuk ransomware family, so we've dubbed it Tsunami-Ryuk. The group continued to remain prolific in their opportunistic targeting of cloud infrastructure for the purposes of carrying out cryptocurrency mining and DDoS campaigns. Our research took a comprehensive look at Keksec's intrusion infrastructure and includes previously unreported content such as the use of the new Tsunami variant named Ryuk, new DGA algorithm, and new persona details.

Lastly, we've also conducted a large research effort into the adversary most commonly referred to as 8220 Gang. Our research has led to the discovery of a new cluster of activity linked to a 8220 Gang campaign of infecting hosts, primarily through common cloud services, with a custom miner and IRC bot for further attacks and remote control. PwnRig, the custom XMRig-based miner variant, attempts to conceal its configuration details and makes use of a mining proxy to prevent the public from monitoring its pool details. The modified IRC bot is also installed on victim machines. In order to assist defenders, we provided a dump of all samples, their build IDs, associated C2 server(s), and mining wallets for the PwnRig family. We also provided a dump of the illicit miner samples, their configured C2 server(s), in addition to the IRC channel configured for automated botnet access. For this Cloud Threat Report, we're also releasing a Ghidra script for easier automation of PwnRig intelligence extraction - Download Here.

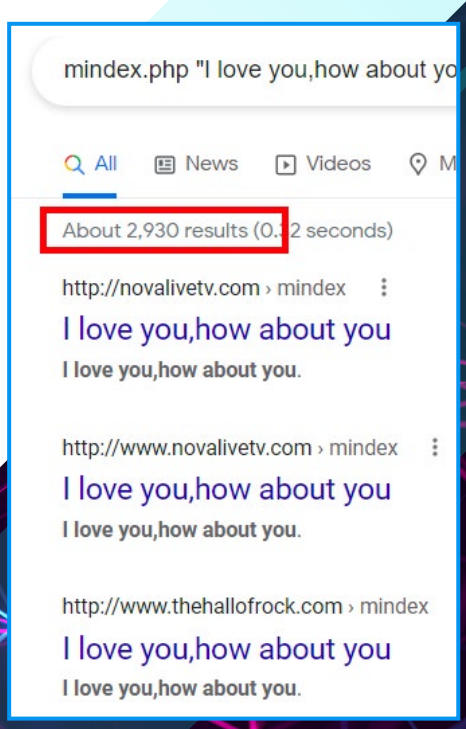

Vulnerabilities and Attacker Opportunities

Lacework Labs has conducted research and published findings on a variety of opportunities in which attackers are ripe to benefit from, and how they can and do operate in novel attack techniques.

First up, our team took a look at the opportunities for potential Canary Token abuse to aid ransomware attacks. Based on our findings, a new technique was discovered in which the legitimate canarytokens.org service was potentially abused as a notification service for ransomware execution. At the time of our blog, we assessed the actor(s) were still developing this technique. It's important to note we have not seen this technique actively used in the wild. However due to it linking to a potential emerging ransomware attack technique, we decided it was important to share with the community at large. The research team at canarytokens.org was contacted and the abusive canary tokens were promptly deactivated.

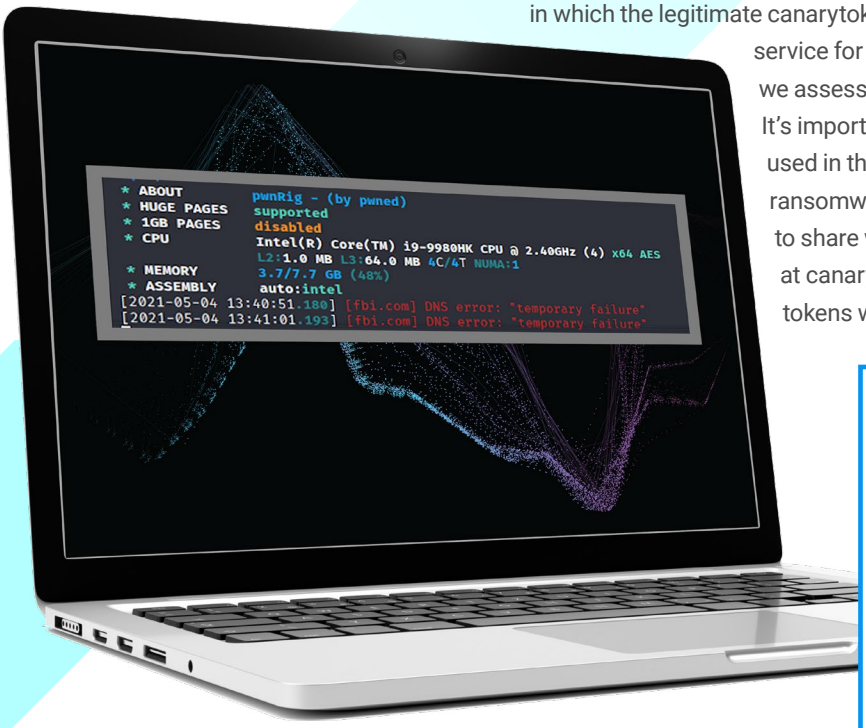

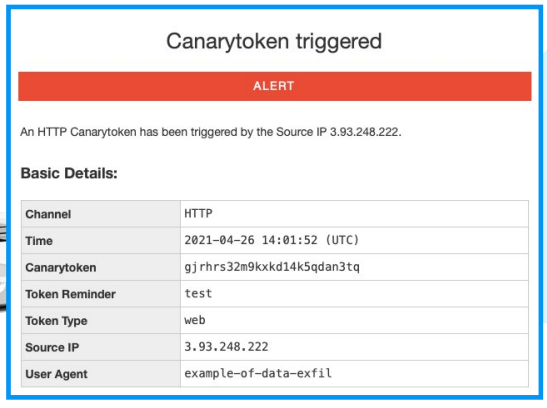
Canarytoken triggered

ALERT

An HTTP Canarytoken has been triggered by the Source IP 3.93.248.222.

Basic Details:

Channel	HTTP
Time	2021-04-26 14:01:52 (UTC)
Canarytoken	gjrhrs32m9kxkd14k5qdan3tq
Token Reminder	test
Token Type	web
Source IP	3.93.248.222
User Agent	example-of-data-exfil

In early June we released an analysis on Initial Access Brokers (IAB) and their expansion into offering access into cloud networks. IABs have evolved from the opportunistic compromise of one-off internet-facing assets for resale as mere proxies, to now including the targeting of corporate networks, assessing the environment for business value, and ultimately selling access into the organization for usage by ransomware gangs, espionage, and everything else in between. AWS, Google Cloud, and Azure administrative accounts have quickly become the new hotness in underground marketplaces. What started as one-off marketplace postings continues to escalate as criminals begin to understand and operationalize the utility of access to cloud services above and beyond cryptocurrency mining. Our analysis shares a peek into some of these marketplaces, the impact they can inflict, and steps we can take to mitigate this risk.

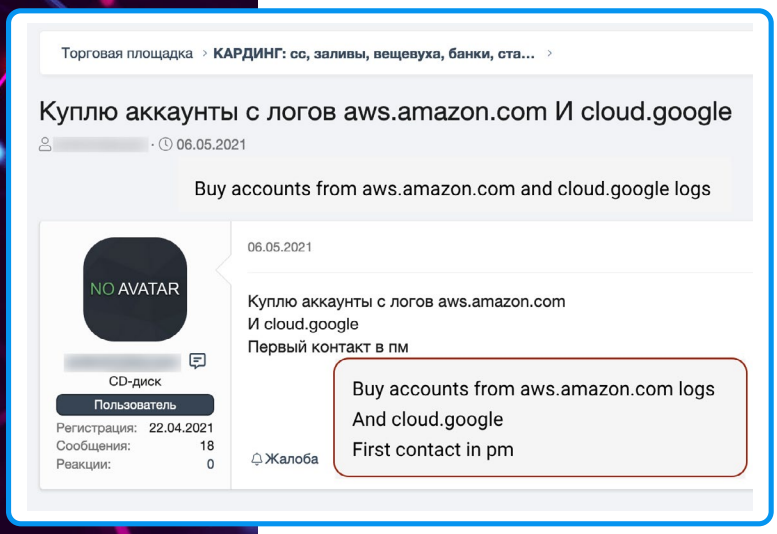

Lastly, our team also identified a collection of vulnerabilities in the XMRigCC following various cryptojacking intrusions. XMRigCC is a fork of XMRig that offers a UI for those managing numerous Cryptominers, along with the ability to start, stop, restart and execute commands on the remote miners. The XMRigCC vulnerabilities we've identified enable rogue clients/compromised hosts/ victims to attack upstream servers, ultimately adding additional risks to victims and legitimate users. Lacework Labs disclosed these vulnerabilities, proof-of-concept code and recommended mitigations to the developer of XMRigCC, and confirmed the fixes, prior to public release.

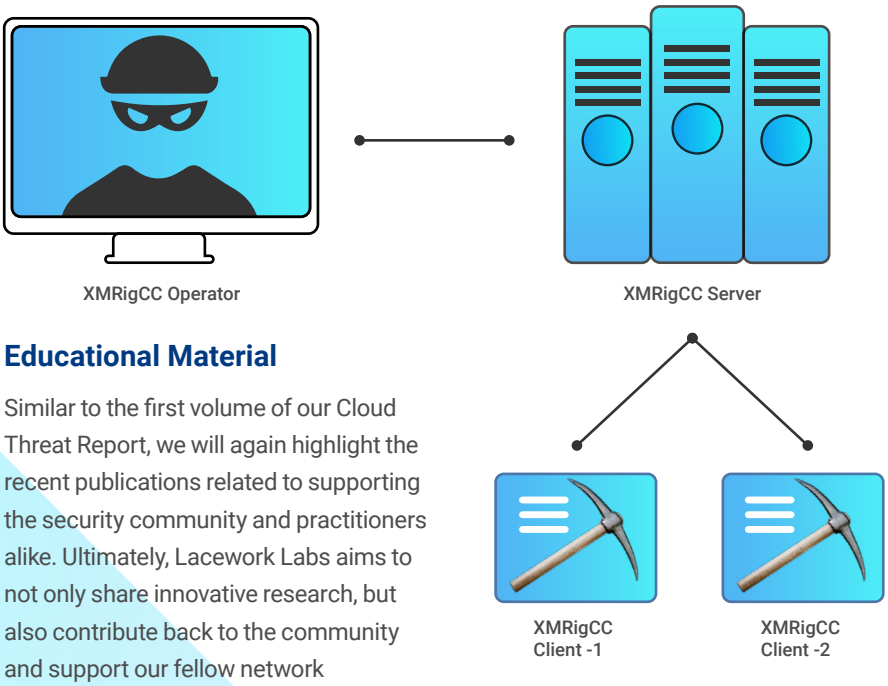

Educational Material

Similar to the first volume of our Cloud Threat Report, we will again highlight the recent publications related to supporting the security community and practitioners alike. Ultimately, Lacework Labs aims to not only share innovative research, but also contribute back to the community and support our fellow network defenders. With that in mind, we've released two publications on this topic.

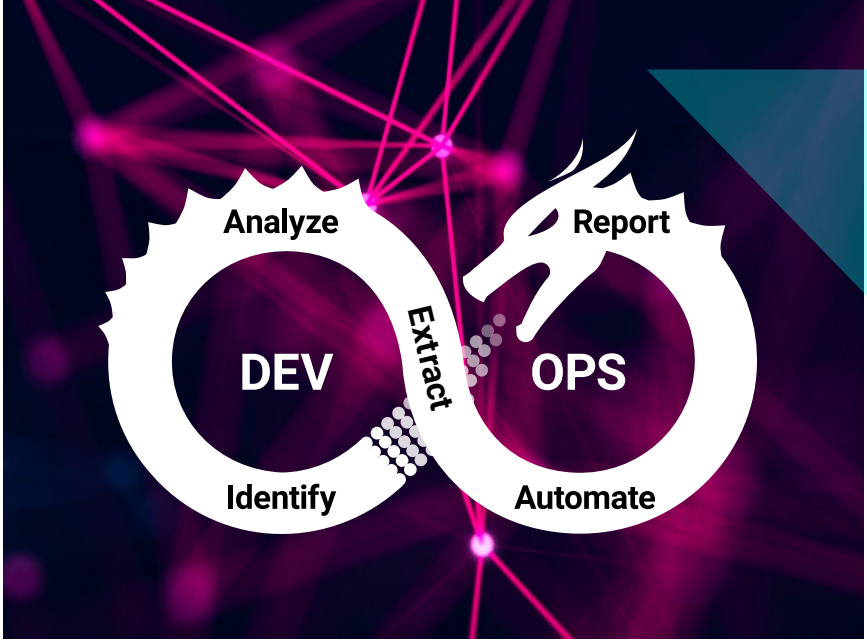

First, "Hacking Like its 1999 – Automating Analysis Like its 2021" dives into research and the release of a Ghidra script in effort to assist researchers and incident responders in automating the extraction of critical information on Katien/Tsunami IRC bot variants. These IRC bots provide attackers with remote access to the victim hosts via the IRC protocol. Lacework Labs continues to observe large quantities of opportunistic cloud-targeting attacks using IRC bots in conjunction with Cryptojacking objectives.

Additionally, we released "Threat Hunting SSH Keys – Bash Script Feature Pivoting." Malicious actors often add SSH keys to victim hosts for persistence, and in this publication we will show you how to hunt with that knowledge. Threat hunting is a vague term with many different areas of focus. However if we focus on the key feature of many bash scripts adding SSH keys, we can identify a large number of malware samples for collection and analysis. Defenders supporting a single network or many can make use of this process in widening their detection capabilities.

Cloud Services Probing

Lacework Labs captures a range of telemetry in both product deployments and custom honeypots. This allows us to see trends relevant to cloud defense purposes. The following data provides insights on traffic captured from May 1 to July 1, 2021.

For honeypots, many cloud-relevant applications are continually targeted. However SSH, SQL, Docker and Redis were by far the most popular based on our telemetry. Noteworthy findings can also be found during analysis of SSH and AWS Cloudtrail logs.

AWS Cloudtrail

Cloudtrail logs can present many insights into AWS environments. For information on AWS reconnaissance, it's often useful to examine S3 activity. Lacework Labs analyzed S3 probing across our customer base and while there was some overlap with honeypot and brute-forcing traffic, the overall traffic profile for the quarter was distinctive. One notable finding was that Tor appeared to be more heavily utilized in AWS reconnaissance relative to other activity. There were definitive trends, however, to the specific Tor networks with the majority originating from the following sources:

- 60729:"Zwiebelfreunde e.V."
- 208294:Markus Koch"
- 4224:"CALYX-AS"
- 208323:"Foundation for Applied Privacy"
- 62744:"QUINTEX"
- 43350:"NForce Entertainment B.V."

The following shows the top observed S3 scanning hosts for the quarter. The most active was 190.2.132[.]207 (ASN 49981:"Worldstream B.V"), which was observed probing the cloud environments for a large portion of Lacework monitored networks.

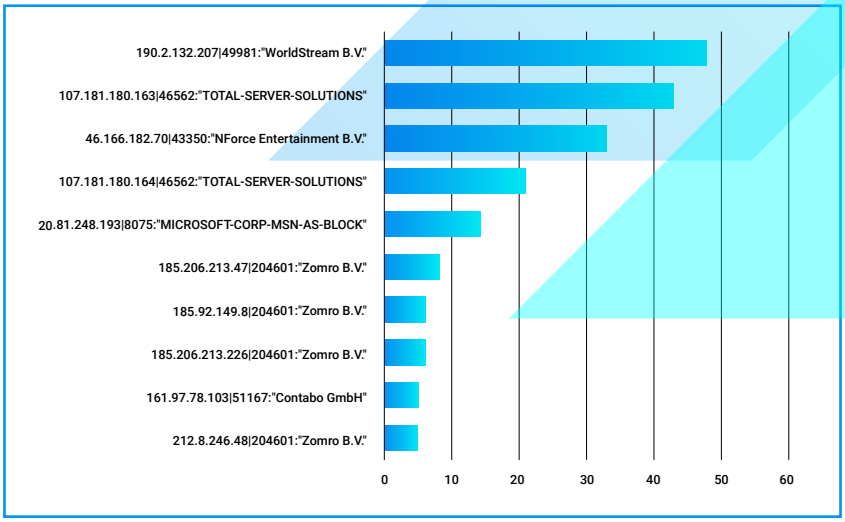

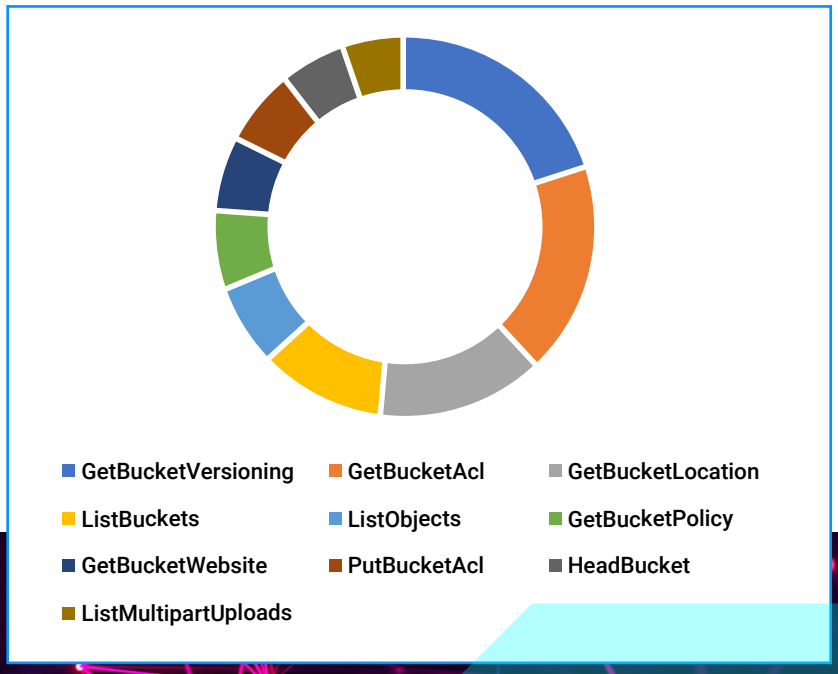

The top observed S3 APIs include GetBucketVersioning, GetBucketAcl, and GetBucketLocation. PutBucketAcl was also seen from many of the Tor nodes.The PutBucketAcl API overwrites the permissions on an S3 bucket so the occurrence of this request from an unknown host is more likely to be offensive as opposed to passive scanning.

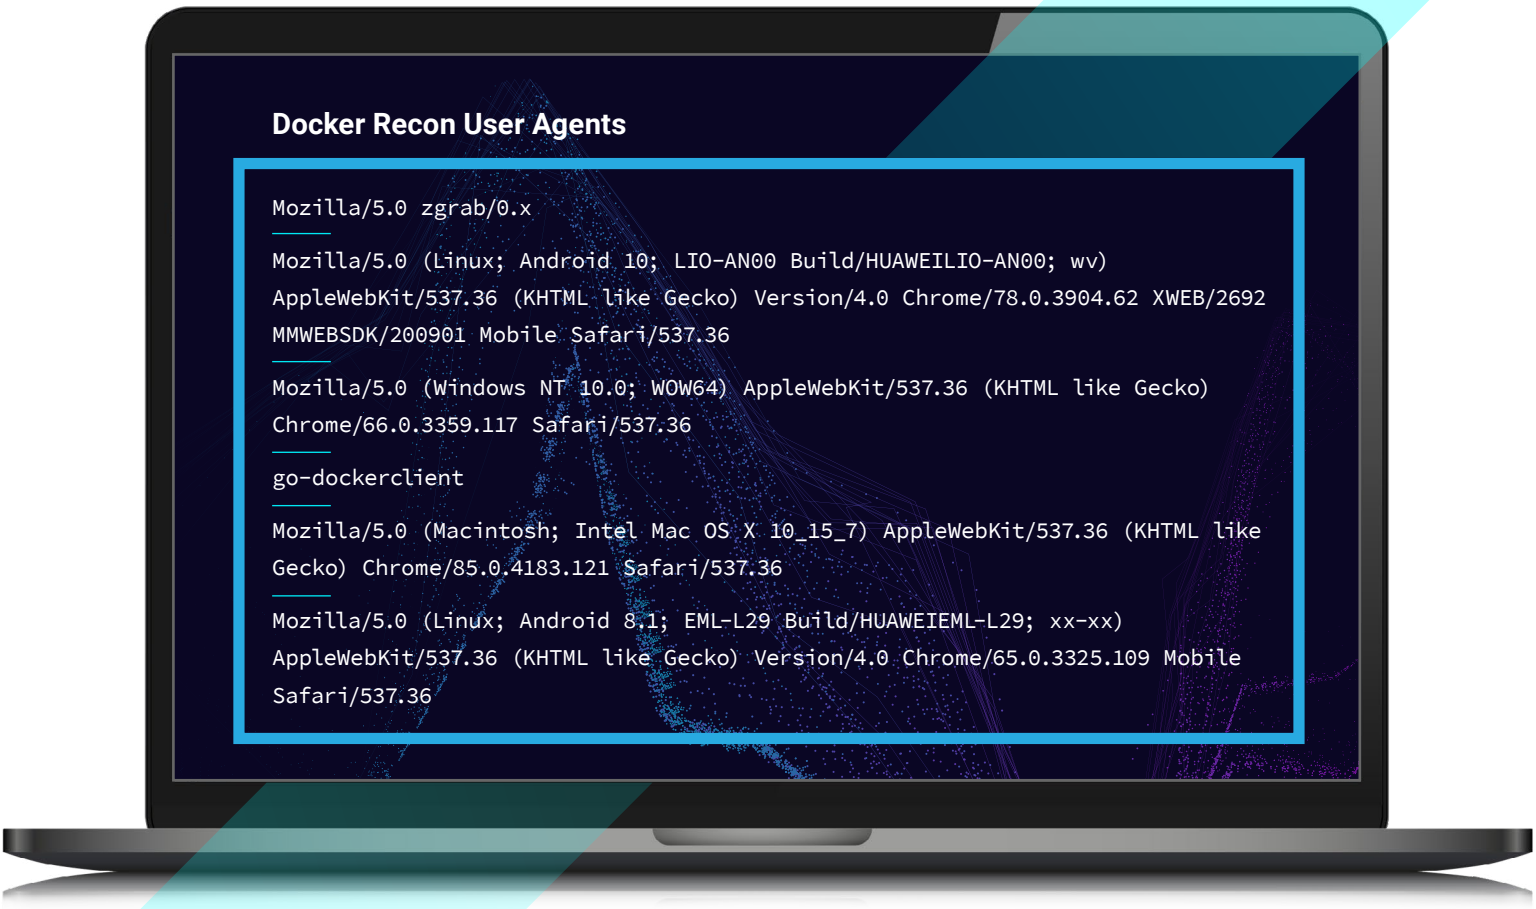

Docker

Roughly 93% of the observed hosts probing Docker APIs leveraged the popular application scanner zgrab. The Go Docker client, as well as several mobile user-agents were also seen suggesting possible mobile-botnet activity. User-agent examples:

Redis

Hosts probing Redis primarily used the Redis command line interface INFO command which returns basic information and statistics about the server. This tactic is used for both malicious recon and by popular legitimate scanning services such as Shodan. Redis has no security out-of-the-box which makes it a low-hanging fruit for attackers.

The following shows top traffic sources by ASN for associated traffic from May 1 to July 1, 2021. Over 27K unique hosts were observed during this timeframe.

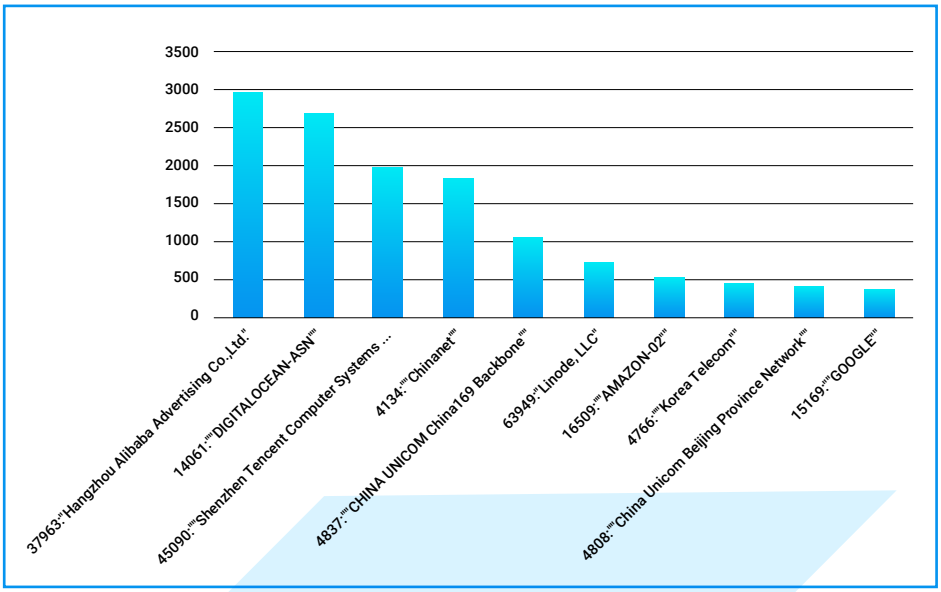

SSH - Secure Shell Protocol

Port 22 (SSH) remains the most attacked port in the wild as it is the most popular service for remote access and administration. Lacework Labs gathers statistics on SSH brute-forcing activities against production environments. This is often carried out using botnets such as Groundhog which was documented by Lacework in January 2021. The Groundhog botnet continues to rapidly expand and has infected over 70,000 cloud servers since tracking.

From May 1st to July 1st 2021 over 137,000 unique brute-force hosts were observed with the majority originating from Asia and India. The top ASN is 4134 Chinanet which was also one of the top networks observed in our honeypot telemetry and has historically been among the top sources of botnet activity across the board.

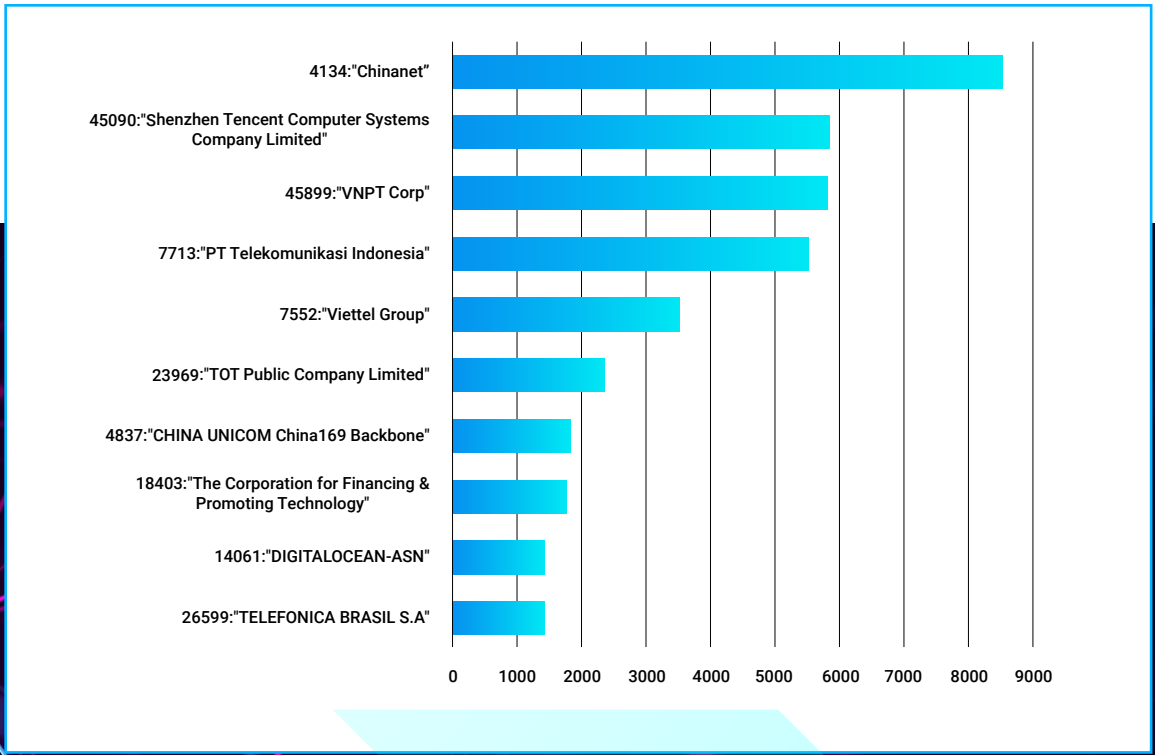

Recommendations

- Ensure your Docker sockets are not publicly exposed and appropriate firewall rules/ security groups and other network controls are in place to prevent unauthorized access to network services running in your organization.
- Ensure your base images are coming from trusted upstream sources and audited appropriately.
- Implement Key-based SSH authentication.
- Ensure the access policies you set via console on your S3 buckets are not being overridden by an automation tool. Frequent auditing of S3 policies and automation around S3 bucket creation can ensure your data stays private.
- Enable protected mode in Redis instances to prevent exposure to the internet.

Connect With Us

The Lacework Labs team continues to build and expand our online presence in effort to contribute back to the security community. Below are areas you can find and follow us delivering excellence in efficacy and innovative threat research. Click below!

FIND OUT MORE

lacework.com

Lacework
6201 America Center Dr
Suite 200
San Jose, CA 95002

08/21

Lacework