UNIT 42

THREAT FORECAST:

CLOUDY WITH A CHANCE OF ENTROPY

Summer 2019 Unit 42 Cloud Threat Risk Report

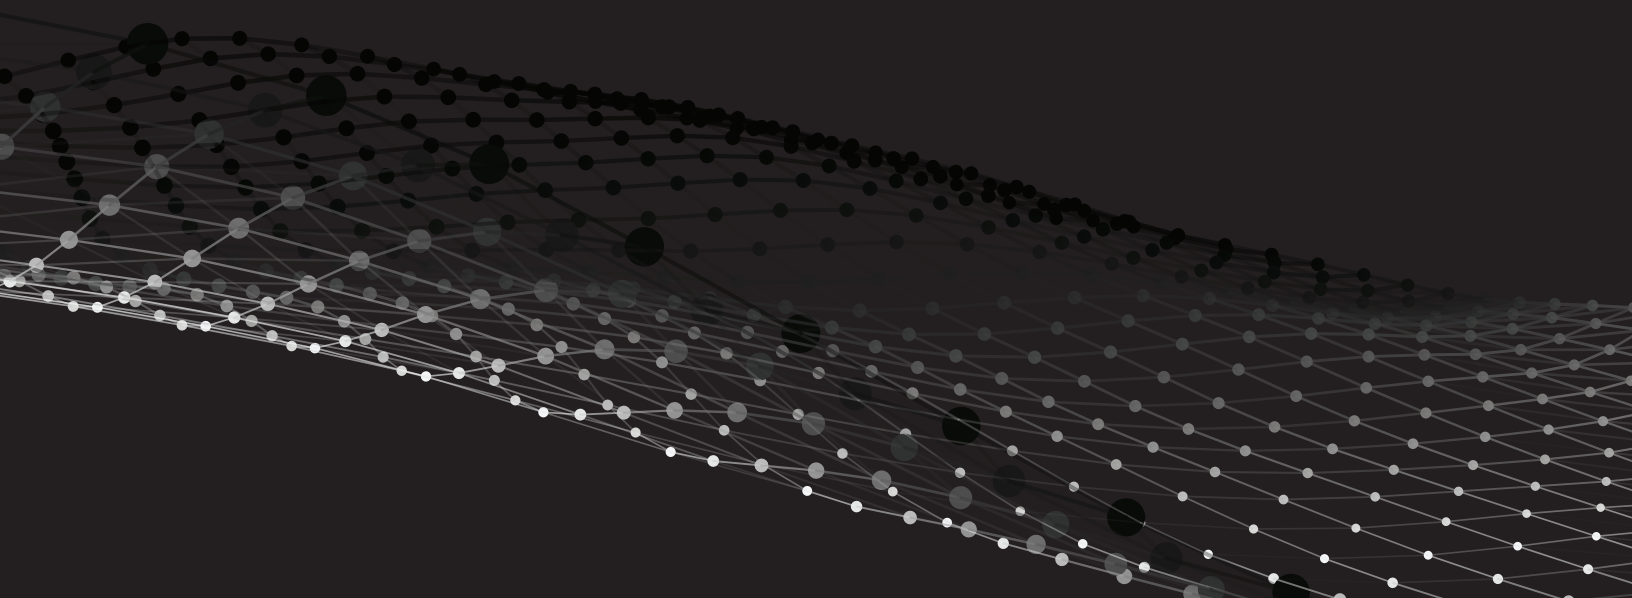

Introduction

In the first half of 2019, there were 21 headline-grabbing incidents involving public cloud platforms. Cloud service providers (CSPs) maintained their sterling reputation for platform security as only a very small percentage of the incidents could be directly attributed to the providers. However, consumers of infrastructure- and platform-as-a-service (IaaS and PaaS) cloud offerings continue to struggle with getting the basics of security right.

"The information showcased in the following report is based upon threat intelligence from multiple data sources and is not based upon surveys. Surveys are excellent indicators of how organizations think, while threat data shows what they actually put into practice."

If there's one thing cloud providers have done extremely well, it's innovate. Unfortunately, this torrent of new, innovative features—often released on a near-daily basis—has led to exponentially more complexity. Although many IT and security organizations conceptually understand the Shared Responsibility Model, our research shows there is a breakdown when putting this concept into practice.

This report highlights key insights from incidents spanning the first half of 2019 and presents original research from the cloud-focused division of the Palo Alto Networks Unit 42 threat intelligence team. We've designed the report to empower your business with the required security knowledge, tools, and best practices to fulfill your role in the Shared Responsibility Model.

More than

40,000

container systems operate under default, insecure configurations

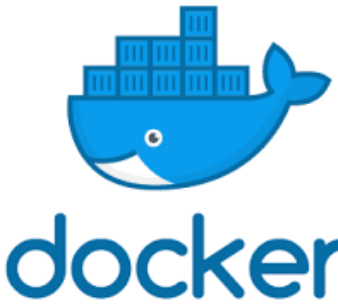

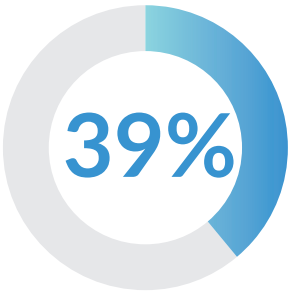

of organizations publicly expose RDP (port 3389) on cloud hosts

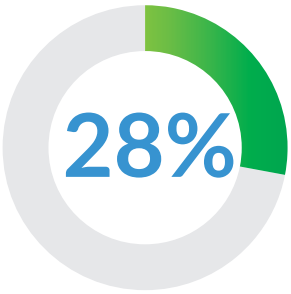

of organizations communicate with known malicious cryptomining C2 domains

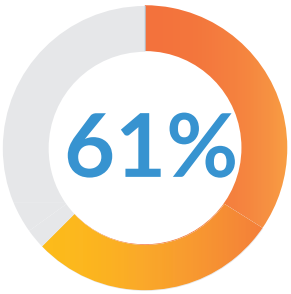

of organizations use unsecured TLSv1.1 and older protocols

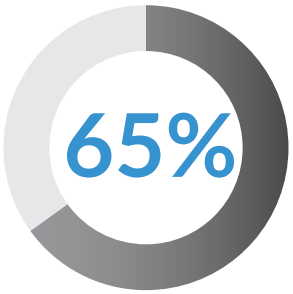

of reported cloud incidents were due to misconfigurations

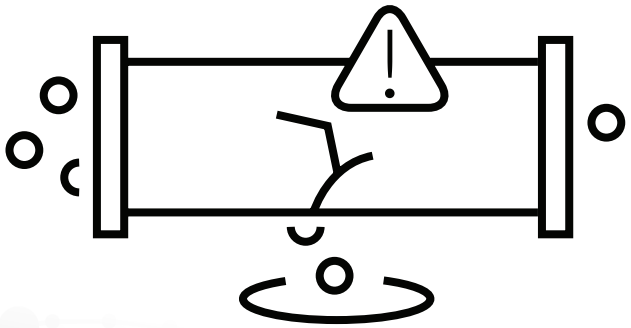

Data leakage is the

#1 outcome of attacks on public cloud infrastructure

Figure 1: Key trends and takeaways

Summary of Impact

40,000 Exposed Containers
Research shows more than 40,000 cloud container systems, such as Kubernetes® and Docker, that have default configurations. These containers were easily identified using basic keyword searches. Several also hosted default-configured applications, such as ElasticSearch® and MySQL® databases, as well as database user interfaces like Kibana®. Attackers frequently target default-configured systems due a higher potential for successful compromise as these systems are less likely to be patched or updated.
Exposed RDP
Remote Desktop Protocol (RDP) allows a user to establish a graphical user interface with a remote system. Traditionally, IT departments use this functionality to manage production systems within a data center. Should a malicious actor take advantage of RDP services, the actor would have control over the cloud system as if physically sitting in front of it.
Cryptomining Operations
Malicious cryptomining, also known as cryptojacking, is the process of secretly mining digital currency, such as bitcoin, on systems owned and operated by another party. Researchers identified 28% of organizations had established communications with domains known to initiate cryptojacking operations.
Insecure Protocol Usage
Secure Sockets Layer (SSL) are security technologies used to establish encrypted links between web servers and browsers. SSL was deprecated in June 2015 and replaced with Transport Layer Security (TLS). TLS version 1.1 was expired and replaced in 2008 with version TLSv1.2. But researchers have identified that 61% of organizations still enable and use TLSv1.1/1.0 and deprecated versions of SSL within their environments.
Reported Cloud Incidents
Sixty-five percent of reported incidents involving cloud infrastructure were found to have resulted from misconfiguration. Resulting in an increased chance of data leakage from these organizations.

Attackers Are Opportunistic

Attackers want your data, wherever it sits. For an attacker looking to exfiltrate data and make a profit, common cloud misconfigurations make for easy targets. Over the last 18 months (as of this writing), 65% of publicly disclosed cloud security incidents were due to misconfigurations, and 25% were due to account compromises.

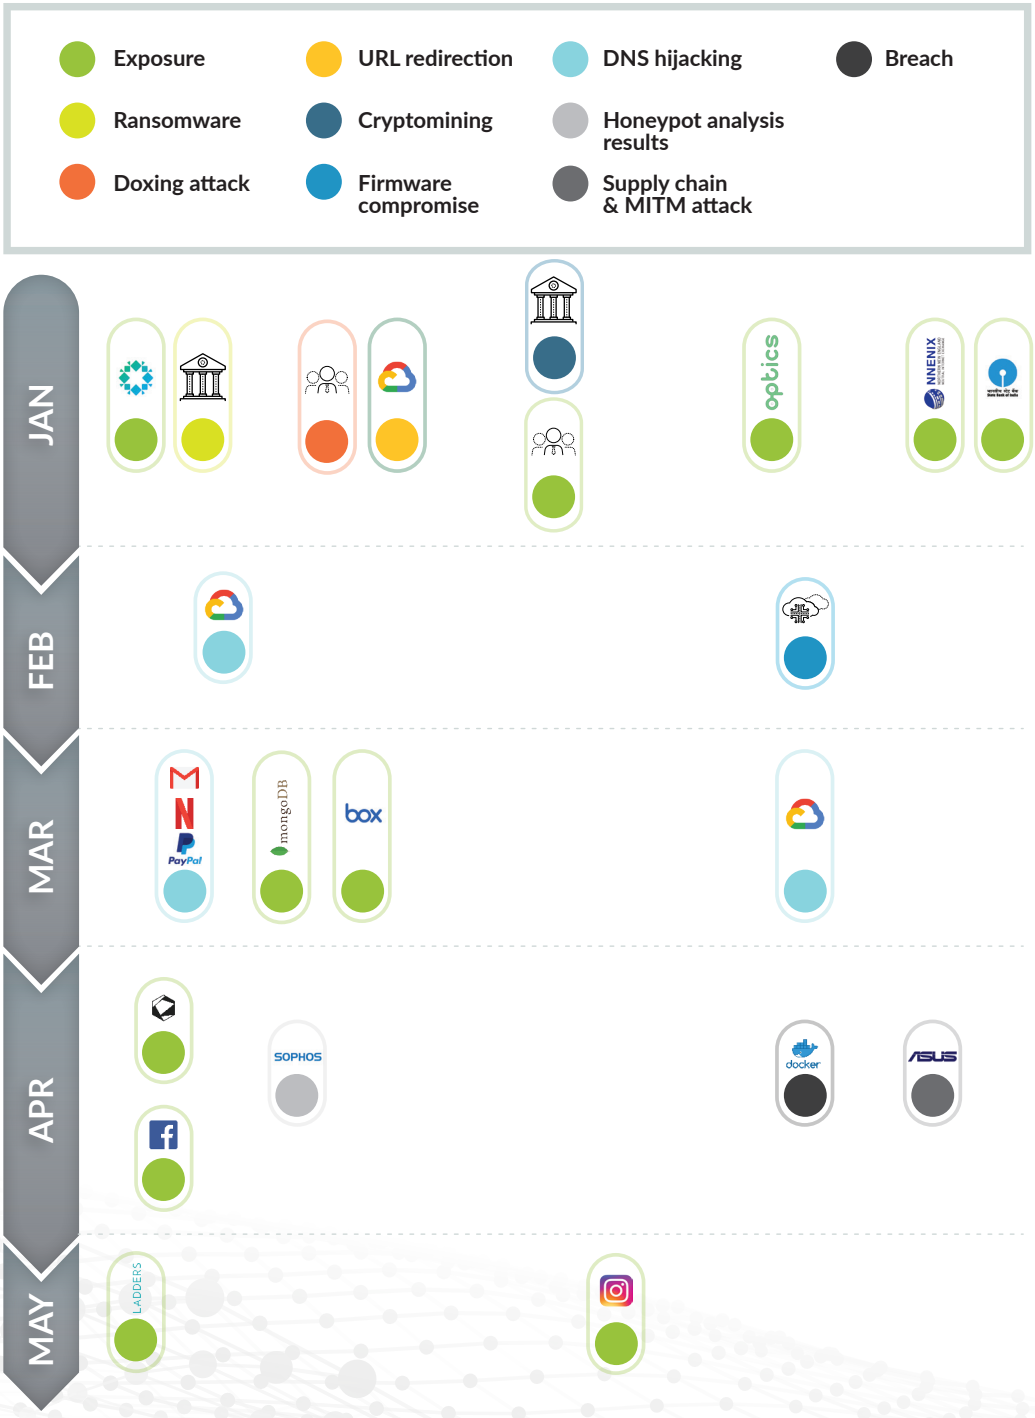

Figure 2: Cloud incidents between January and May 2019

At-a-Glance Summary

Exposure Rises as Container Adoption Surges

Docker and Kubernetes have taken the enterprise by storm over the past 18 months. Unit 42 research shows more than 40,000 unique containers have default configurations, allowing for quick identification. Many organizations also appear to use default-configured applications hosted on default containers. Security teams need to embrace containers as they are key to enabling DevSecOps. However, teams also need to ensure that the applications and hosts are securely configured and monitored.

Malware Extends Its Reach to the Cloud

Cloud-based malware attacks are becoming increasingly common. In early 2019, Unit 42 reported on a campaign by the Rocke cybercrime group that specifically targeted public clouds. Detecting these types of attacks has proven difficult for many organizations as security tools deployed in cloud environments often lack integrated threat intelligence feeds.

Cloud Complexity Yields Low-Hanging Fruit for Attackers

Despite CSPs continuing to bolster their native security capabilities, Unit 42 research shows misconfigurations as the most common type of security incident. Organizations need to focus on shifting security left by better understanding how code is pushed to the cloud. This will help security teams codify secure configurations in infrastructure as code (IaC) and prevent many of these types of exposures before they make it to production.

On-Premises Vulnerability Management Habits Carry Over to the Cloud

Some habits die hard, and patching is one of them. Even novice security professionals know on-premises patching is a challenge. Unfortunately, the challenge has spread to the cloud. Organizations struggle with this as many standalone vulnerability management tools lack cloud context and remain scattered across multiple consoles. Security teams would be wise to consolidate tools to create a cloud-centric view.

Detailed Summary

Exposure Rises as Container Adoption Surges

Containers are not immune to misconfigurations or data exposure simply because they are hosted in the cloud. Container platforms like Docker and orchestration platforms like Kubernetes bring security concerns of their own when used in production environments. More than 40,000 container platforms exposed to the internet were found to be using out-of-the-box configurations that allowed them to be identified using the simplest of search terms: the platforms' names themselves, "Docker" and "Kubernetes" (see Figures 3 and 4).

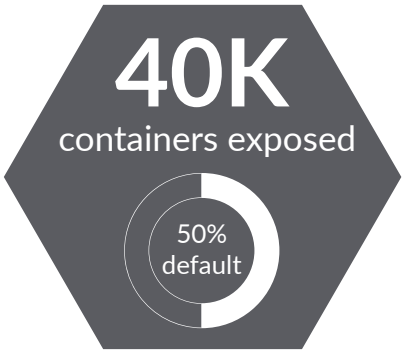

Total results

23,354 exposed Docker containers

Top countries

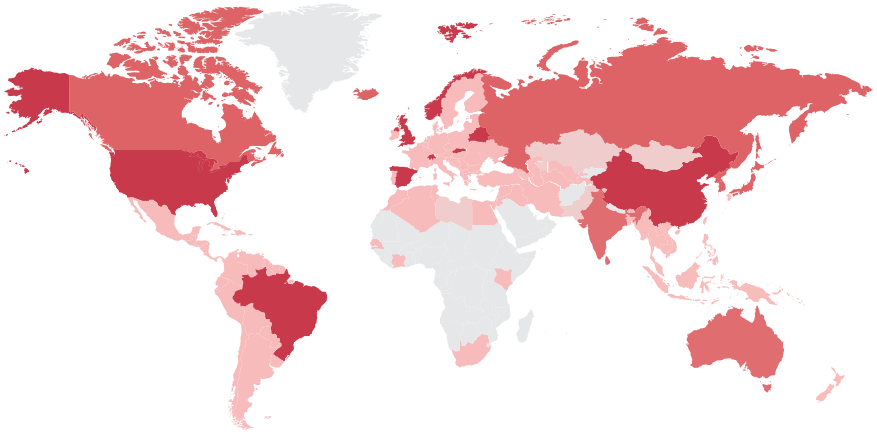

CHINA
6,015

UNITED STATES
4,617

GERMANY
2,119

HONG KONG
1,960

FRANCE
1,639

Figure 3: Docker containers exposed to the internet

Total results

20,353 exposed Kubernetes containers

Top countries

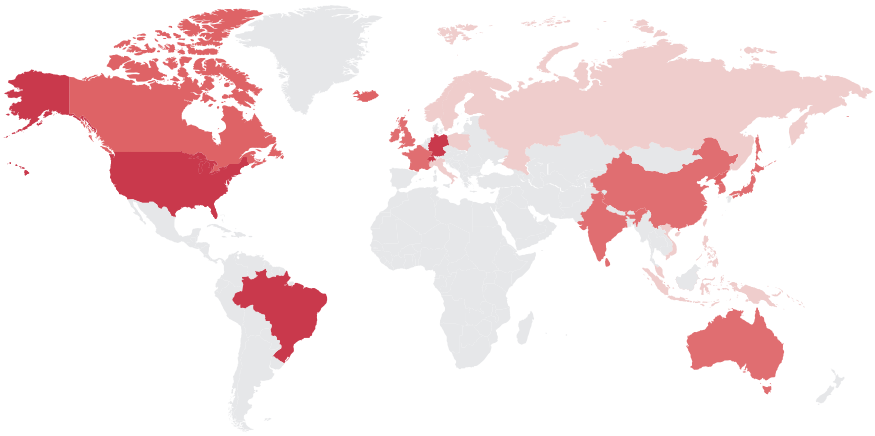

UNITED STATES
11,425

IRELAND
2,834

GERMANY
2,529

SINGAPORE
793

AUSTRALIA
522

Figure 4: Kubernetes containers exposed to the internet

The researchers discovered that both the containers and hosted services were using default configurations. Not all identified systems allowed for unauthenticated access to the data they contained, but a surprising number did. A select few of the default containers were also hosting default-configured applications, which allowed researchers to enter those systems without authentication. The apps, identified as ElasticSearch and Kibana, were found on both Docker and Kubernetes platforms.

Every container with sensitive data should be placed behind a properly configured security policy or external-facing firewall to prevent access from the internet. Additionally, Unit 42 researchers suggest that organizations should not keep default configurations for their container infrastructure. Instead, an organization's security policy should provide guidance on container configuration that is unique to that organization and require authentication before any data can be retrieved.

Malware Extends Its Reach to the Cloud

Unit 42 has been studying, following, and reporting on actor groups and detailed incidents that target cloud infrastructure. Rocke, a Chinese-based cyberthreat actor group also known as The Iron Group, SystemTen, Kerberods, and Khugepageds, is one such group. Rocke was initially reported to perform data destruction ransomware attacks, before pivoting toward cryptomining operations within cloud infrastructures. The group has joined several others in targeting cloud systems by writing cloud-focused malware that disables and uninstalls cloud security tools.

Rocke utilizes relatively unusual tactics, techniques, and procedures (TTPs) when operating: a 12-step operation (see Figure 5) that includes the use of a third-party code repository, such as Pastebin or GitHub®. The repository delivers the initial payload by connecting to a command-and-control (C2) server, and attackers elevate their rights and continue their operational procedures from there.

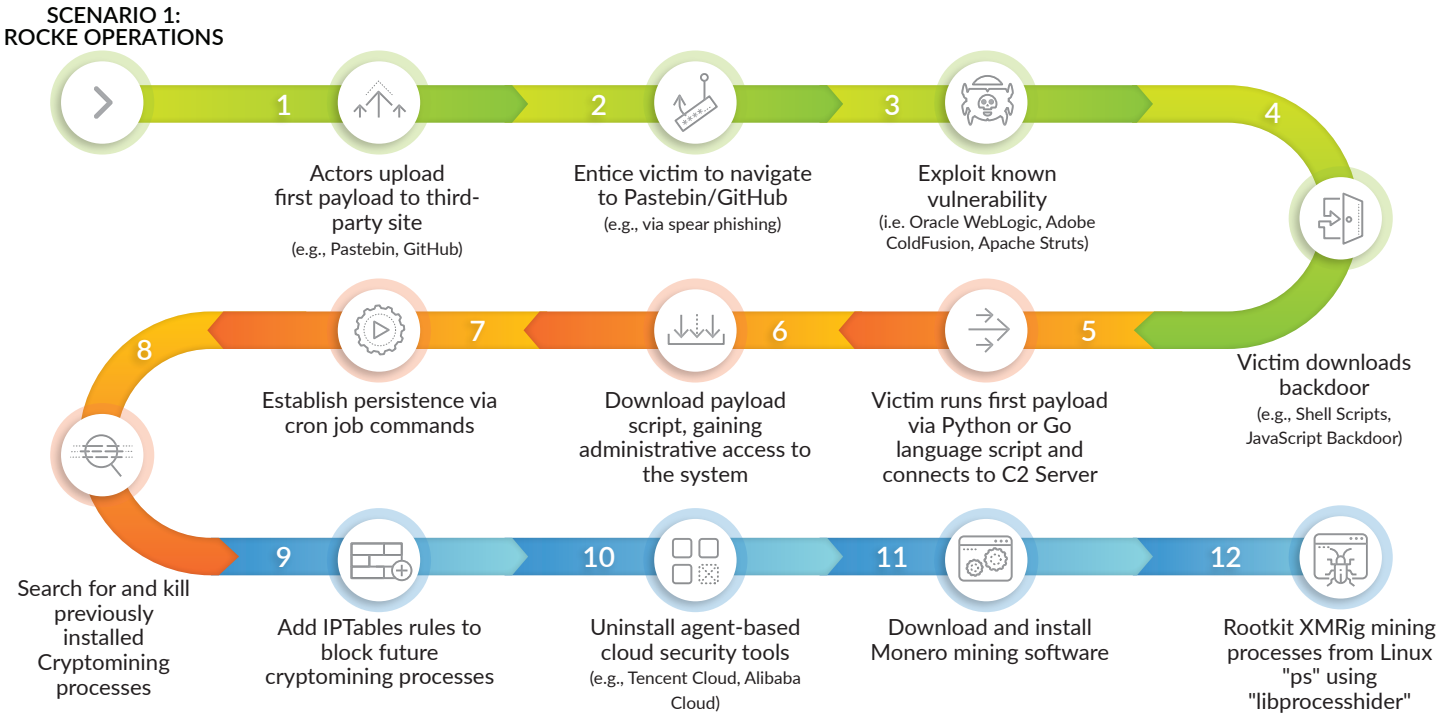

Figure 5: Rocke group's 12-step TTPs

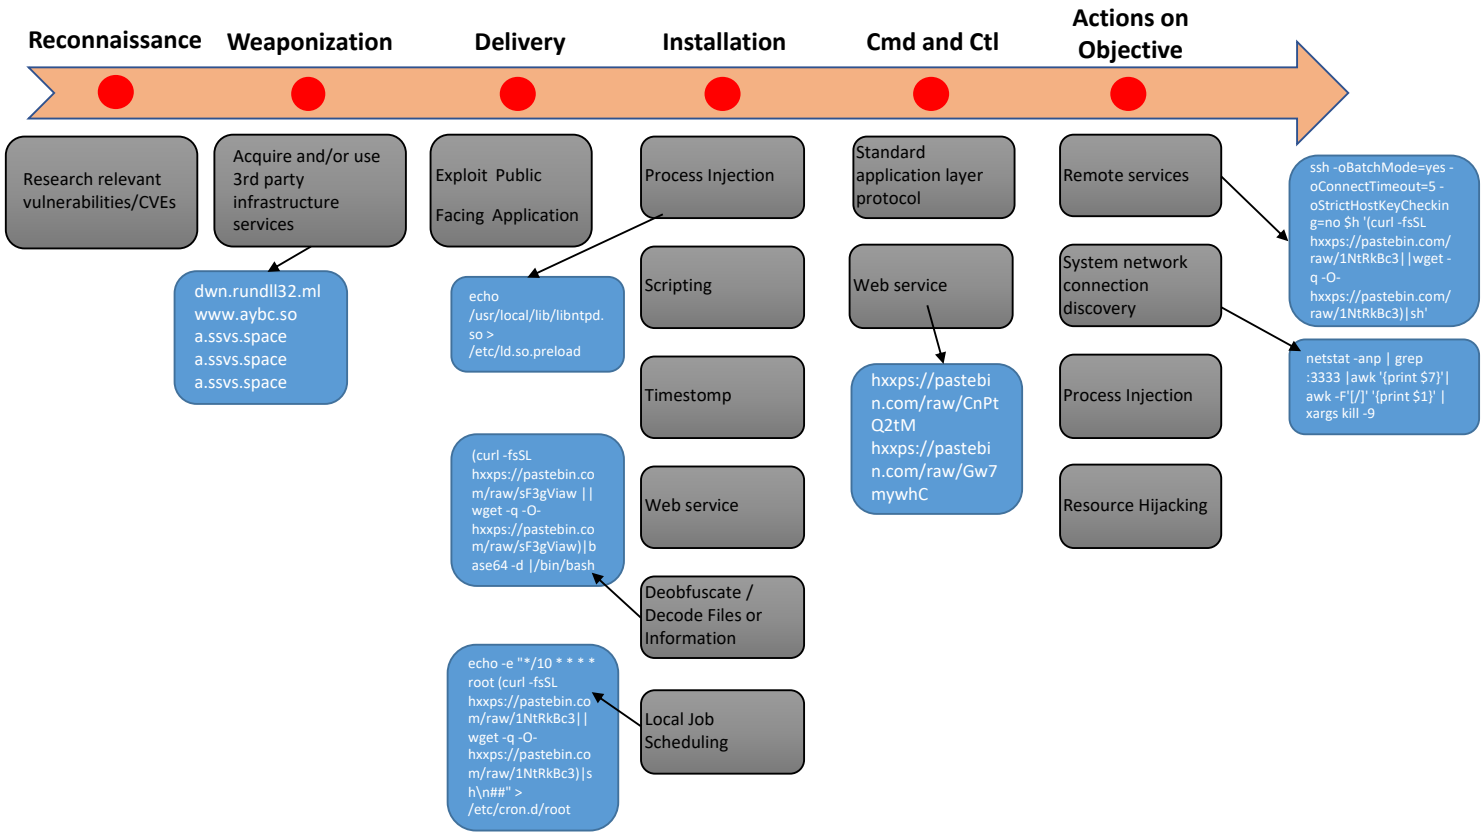

Figure 6: Rocke operations mapped on the cyberattack lifecycle

Rocke is able to perform cyber operations with relative ease considering the vulnerabilities they exploit are not new. The majority of reported vulnerabilities used by the group were released in 2017 (CVE-2017-10271—Oracle WebLogic and CVE-2017-3066—Adobe ColdFusion) and 2016 (CVE-2016-3088—ActiveMQ). This alone highlights that Rocke could be slowed down if organizations maintained timely and effective patching schedules for cloud systems. Additionally, organizations should employ up-to-date blacklisting rules on firewalls to prevent connections to known malicious domains. It is important to note that all of Rocke's known C2 domains have been flagged as malicious and are blocked by PAN-OS®. These domains are known to be part of the Rocke infrastructure and should not be allowed to connect to an organization's cloud systems:

- sowcar[.]com
- w2wz[.]cn
- z9ls[.]com
- thyrsi[.]com
- baocangwh[.]cn
- gwjyhs[.]com

Cloud Complexity Yields Low-Hanging Fruit for Attackers

Data from reported public cloud-related incidents between February 2018 and June 2019 shows that 65% of these incidents resulted from misconfiguration. As a result of these misconfigurations, data leakage is the No. 1 outcome of attacks on public cloud infrastructure (see Figure 7).

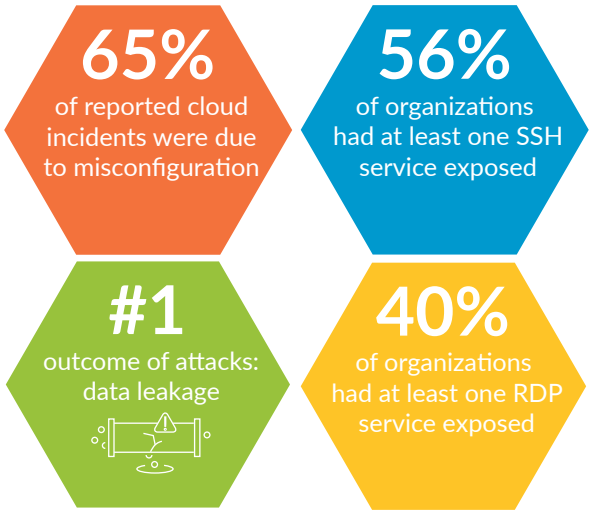

Figure 7: Public cloud incidents, Feb 2018 – Jun 2019

In pairing this data with the alerts gathered from CSP environments, the largest number were "insufficient access control allowing internet access to/from internal networks." Also, the most common services exposed to the internet from cloud infrastructure were SSH (TCP port 22) and RDP (TCP port 3389). Our data reveals that 56% of organizations had at least one SSH service exposed to the internet, and nearly 40% of organizations had at least one RDP service exposed. There is no reason to allow inbound traffic from the entire internet (0.0.0.0/0) to services such as SSH, RDP, or SMB. System administrators should always restrict open ports/services to a small set of whitelisted sources. The inbound traffic control policy can be configured through security groups in AWS®, Network Security Groups in Azure®, and firewall rules in GCP™. These traffic control policies need to be monitored continuously for any deviation due to misconfiguration or attack. Administrators also need visibility into firewall logs to watch for any indicators of malicious activity.

On-Premises Vulnerability Management Habits Carry Over to the Cloud

Patching is an operation that needs critical focus in the cloud. Taking the previous example from Rocke, the group actively exploits vulnerabilities that are more than two years old. At the time of this writing, we have identified 29M vulnerabilities in Amazon EC2®, 1.7M in Azure Virtual Machine, and 4M in GCP Compute Engine. These vulnerabilities target the applications customers deploy to CSP infrastructure, such as outdated Apache servers and vulnerable jQuery® packages. Cloud systems, whether Amazon EC2, GCP Compute Engine, or Azure Virtual Machine, can all be configured to originate from a template. The template option allows

29M
vulnerabilities in Amazon EC2

1.7M
in Azure Virtual Machine

4M
in GCP Compute Engine

organizations to quickly and efficiently spin up systems and containers that match a given configuration. CSPs provide a framework for organizations to quickly and reliably use resources that have the latest configurations. It's up to the organizations themselves to ensure that systems and containers being used in production environments have the latest security patches. The beauty of cloud environments is that the entire environment can be deployed from a template, meaning that every system deployed maintains the same configuration as every other system of its type.

Security teams must ensure that the golden template used by AWS, GCP, Docker, or Kubernetes to deploy production systems is configured to use the latest security patches and versions as directed by the application vendor. This will ensure organizations are performing their due diligence in maintaining secure environments and raising the overall security hygiene of their cloud infrastructure. Many organizations are adopting serverless or function-as-a-service (FaaS) to delegate the platform responsibility to the CSPs and focus only on developing the applications. Although FaaS frees system administrators from many maintenance hassles, the problems of misconfiguration and unpatched applications remain unsolved. For example, an inappropriate identity and access management (IAM) product attached to a function can lead to the compromise of another cloud service. A function without an execution limit can be flooded with requests and result in significant financial costs for wasted resources. A function using outdated third-party libraries is still vulnerable to cross-site scripting or SQL injection. It is imperative to gain visibility and continuously monitor every serverless function.

Conclusion

Security teams need to understand the core requirements for securing modern applications and workloads in the cloud:

1. Ensure your security teams can access a realtime view across virtual machines, containers, and serverless applications. Maintaining visibility into diverse computer paradigms can be a challenge, but it is critical.

2. Integrate security into DevOps workflows to allow your security teams to scale their efforts in an automated way. Developers have a lot of power in the cloud, and your security needs to be able to keep up.

3. Harden your applications and workloads. Although some security requirements fall to CSPs as part of the Shared Responsibility Model, your security teams are still responsible for configuration and compliance of individual workloads, containers, and functions, including platforms like Kubernetes.

4. Maintain runtime protection. As your organization's cloud footprint grows, being able to automatically model and whitelist application behavior becomes a powerful tool for securing cloud workloads against attacks and compromises.

Ready to Identify the Threats in Your Cloud?

Prisma™ Public Cloud (formerly RedLock) analyzes more than 10 billion events every month. That analysis shows us that poor configuration, permissive behaviors and lack of policies lead to many openings for bad actors and unidentified threats to exploit. By proactively detecting security and compliance misconfigurations as well as triggering automated workflow responses, Prisma Public Cloud helps ensure you continuously and securely meet the demands of your dynamic cloud architectures.

Start your 30-day free trial

Prisma Public Cloud Scanning APIs are free public API services designed to help developers, DevOps and security teams detect and address security issues with their containerized applications.

Start scanning for vulnerabilities

About Unit 42

Unit 42 is the global threat intelligence team at Palo Alto Networks. We believe threat intelligence should be free, shared, and available to all for the common good. We deliver high-quality, in-depth research on adversaries, malware families, and attack campaigns. Our analysts uncover and document adversary behaviors, and then share playbooks that give insight into the various tools, techniques, and procedures threat actors execute to compromise organizations.

We share our findings freely so defenders everywhere can access world-class threat intelligence. Unit 42 is a recognized authority on cyberthreats, frequently sought out by enterprises and government agencies around the world.

About Prisma by Palo Alto Networks

Prisma™ provides the industry's most complete cloud security offering. Accelerate your cloud journey with a product suite designed to secure today's complex IT environments.