



Lacework™

2019 MID-YEAR CLOUD SECURITY TRENDS AND TIPS REPORT



The amount of enterprise data and workloads moving to the cloud is on a steady rise. With that, the number and sophistry of threats targeting the cloud are also increasing.

Securing enterprise assets in the cloud involves some unique challenges:

- Cloud security pivots on a shared responsibility model. In this model, cloud service providers such as Amazon Web Services (AWS), Microsoft Azure, Google Cloud, etc. are responsible to protect the cloud infrastructure and services. It is the onus of enterprise subscribers to employ security controls to protect their user accounts, data, and applications.
- On-premise applications, though shrinking, are expected to be around for years to come. As such, an enterprise cloud security strategy needs to be robust enough to protect a combination of on-premise, single-cloud, multi-cloud, and hybrid environments.
- As enterprises shift towards DevOps-style development patterns and serverless PaaS models, workloads are more granular and dynamic with shorter lifespans. This presents new security challenges, as solutions designed to protect static applications and end-user devices would be inadequate.
- In highly scaled and dynamic cloud environments, traditional security measures such as pre-defined rules and signatures of known threats are insufficient to keep pace with the fast-evolving threatscape.

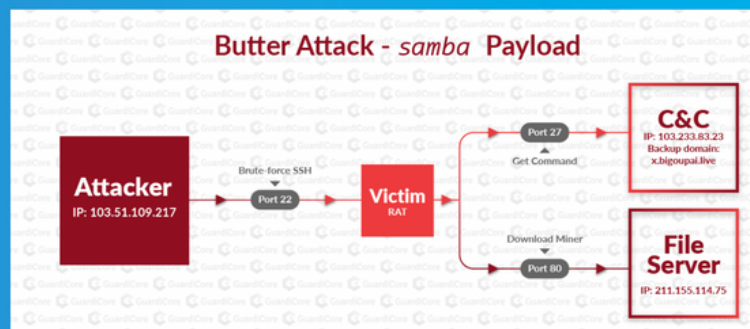
The changing dynamics of the cloud have been a fertile ground for more recent threats and security incidents. Based on industry-wide threat research and analysis, Lacework highlights **ten** major trends in cloud security seen over the past year. Mitigation strategies for these trending threats are also presented.

1) ACCOUNT COMPROMISES

Unauthorized account access, insecure interfaces/APIs, privilege escalations to hijack accounts, services and traffic are some of the trending instances of account compromise. Typical ways in which organizations expose their accounts are by committing hardcoded keys to the repo (e.g. GitHub), yielding to phishing attacks, and exploitation. A common method attackers employ to gain unauthorized access is SSH brute force attack. Repeated attempts to guess username-password combinations for the secure shell is indeed an old tactic, but still effective to brute force on public cloud workloads. In the recent Bread and Butter malware campaign, Butter executes a brute force SSH attack and leaves a backdoor user named butter, together with a Trojan registered as a service on every breached Linux machine.

EXAMPLE – BREAD & BUTTER ATTACKS

- Recent Malware campaign
- Begins with brute force SSH
- Add user “butter”
- Downloads RAT
- RAT communicates with CNC
- RAT downloads XMR miner
- Reported by [Gaurdicore](#)



Another bad access hygiene seen with organizations is the utilization of root user accounts to perform activities. Cloud experts strongly discourage this. Instead of the root user, administrators can create individual users with IAM policies attached, and ensure that the root user access keys have limited access. Organizations that implement key-based authentication were found not to rotate the access keys. This creates greater exposure as keys are typically associated with more permissions.

Mitigations

- Use multi-factor and key-based authentication versus old-school passwords.
- Create individual users with IAM policies attached, limit root user usage.
- Implement security discipline and train users on safe keep of account keys and credentials.
- Use continuous monitoring that provides deep visibility and flags behavioral abnormalities, e.g. alert on successful SSH authentication after a series of failed attempts.

2) MISCONFIGURATION AND ERRONEOUS ORCHESTRATION

Poorly configured security for cloud databases and cloud orchestration platforms continued to be a weak point for organizations. A common source of this problem is insufficient understanding of the cloud providers' (Amazon, Google, Microsoft, IBM, etc.) security capabilities coupled with oversight. Many enterprises attempt to recreate their local, on-premise solutions in the cloud.

Cybercriminals are becoming more adept at exploiting misconfigured accounts. A recent study by IBM X-force found misconfiguration was responsible for the

exposure of nearly 70 percent of the compromised records they tracked. Infrastructure as code means misconfiguration is the biggest risk, as seen with open S3 buckets, Elastic-search and Kubernetes. In 2018, misconfigured S3 buckets turned out to be the most common reason for leaving sensitive data in the open, which was either stolen, leaked, or held as a hostage against a ransom. NoSQL databases like MongoDB is a favorite target for hackers where misconfigurations could be easily exploited to expose sensitive data. In many instances, lack of passwords for the admin accounts were spotted as the root-cause.

Attackers are also targeting containerized resources in the cloud, serverless applications, API services and orchestration platforms like Kubernetes, which are publicly exposed as a result of poor configuration or oversight.

Mitigations

- Use tools to ensure compliance with CIS, SOC 2, PCI, and HIPAA industry standards.
- Leverage the native logging capabilities from cloud providers to analyze and alert against exploits.
- SecOps workflows with configuration change alerts.
- Implement regular audits to keep up with organizational changes and the dynamic cloud environment.
- Enforce sufficient access restrictions and role-based permission controls.

3) CRYPTOJACKING/CRYPTOCURRENCY MINING

Cloud servers proved to be attractive resources for surreptitiously mining cryptocurrency and blockchain algorithms. In cryptojacking (a recent trend fueled by the bitcoin craze), cybercriminals sneak into vulnerable devices to run

coinminers and sabotage CPU usage. Cryptojacking was a major trend in 2017. Although the falling price of bitcoin ramped down its prevalence, it is still a major threat. Currently, Monero is the most popular coin to mine illicitly.

In addition to individual devices (smartphones, PCs, etc.), cryptojacking criminals have targeted enterprise resources in the public cloud. WannaMine (MSH.Bluwimps), a cryptojacking script, is spread in enterprise networks. It uses the Eternal Blue exploit (tied to WannaCry ransomware attack) to render the infected devices unusable due to high CPU usage.

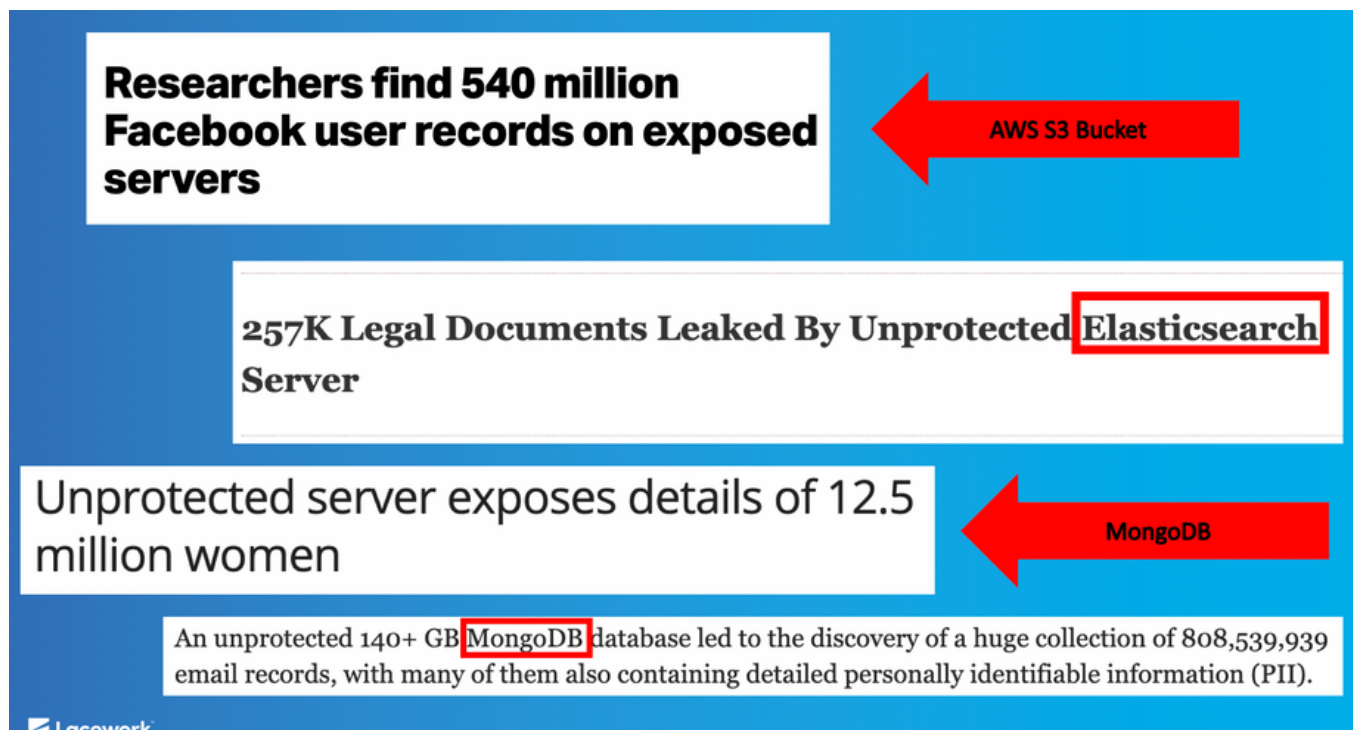
One of the biggest cryptojacking incidents, exposed by Checkpoint researchers in February 2018, exploited the known CVE-2017-1000353 vulnerability in the Jenkins Java deserialization implementation. By taking advantage of Jenkins' large-scale adoption as an automation platform with an estimated 1 million users, the hackers could accumulate 10,800 Monero (roughly 3 Million USD) with this malicious mining operation. Research by Sysdig uncovered cryptojacking attacks moving beyond EC2 exploits to container-specific and Kubernetes-specific exploits. In early 2018, an open, preconfigured Kubernetes instance located on honeypot servers was infected with malicious Docker containers for mining Monero.

Mitigations

- Deep process-level visibility to detect unexpected high CPU usage at all layers of the cloud: servers, containers, and orchestration platforms.
- Implement billing alerts.
- Monitor connections to popular pools.
- Harden host servers with regular patches and updates.

4) DATA LEAKS

Data leaks or data breach involves unauthorized transfers of classified data from end-devices, data centers, etc. to some external domain. Attackers often leverage the leaked data to launch attacks on a large scale. Ponemon Institute's 2018 "Cost of a Data Breach Study"[6] found the total cost of a data breach to exceed \$3.8 million (global average), a 6.4 percent increase from the previous year's average. Misconfiguration has been a leading reason for data leaks. In the cloud environments, threat actors (hackers and disgruntled insiders) leverage a lapse in configuration to expose confidential data, typically from unsecured databases like MongoDB, Elastic-search, Redis, or open cloud provider buckets.



(Figure Source: <https://www.slideshare.net/Lacework/lacework-top-10-cloud-security-threats/16>)

According to the same Ponemon study, the mean time to identify (MTTI) or to contain (MTTC) a data leak incident increased in the past year. The number of days to identify the data leak increased from an average of approximately 191 in 2017 to 197 days in 2018. The average days to contain the data breach increased from 66 to 69 days. Increased migration to the cloud and compliance

failures are cited as top reasons for this increase, coupled with the growth in the usage of IoT devices and mobile platforms.

Mitigations

- Enforce authentication for databases.
- Encrypt sensitive data at rest.
- Increase vigilance in the cloud using solutions that provide deep visibility into internet facing configuration.
- Use security tools to seamlessly audit and alert for open storage buckets.

5) DATA EXFILTRATION



Data exfiltration, a specialized form of a data breach, involves the unauthorized transfer of confidential data from the network to a domain that the threat actor controls. Increased network connectivity due to mobile, IoT, and migration to the cloud is contributing to a steady rise in data exfiltration incidents in enterprises around the globe. Data exfiltration is typically associated with advanced persistent threats (APT) and sophisticated, well-funded, nation-state threat actors.

In data exfiltration incidents, the attacker leverages misconfigurations, accidental exposure, spear phishing, etc., to gain unauthorized access to the network.

In March 2019, a data exfiltration incident for Toyota affected up to 3.1 million customers. The breach came from unauthorized access and affected Toyota Tokyo Sales Holding Inc. and possibly three other independent dealers in Japan. A month prior to APT32 (unconfirmed attribution) spear-phishing was launched against several multinational car companies. Some experts connect

this attack to the fact that Vietnam reportedly is trying to develop its domestic car industry.

Mitigations

- Requires fully mature security posture to protect against the determined and sophisticated threat actor.
- Business must understand where their most valuable information is and how to monitor and protect it.
- In cloud environments involving ephemeral containers and servers, host-based intrusion detection systems (IDS) are more effective than network-based IDS.

6) EVOLVING MALWARE



Any malicious software (RATs, Trojans, backdoors, downloaders, ransomware, etc.) designed to damage a computer, server, client, or computer network is a malware. There is a growing trend where known malware are re-appearing as evolved variants. Adversaries are leveraging social engineering and irregular patching habits to dispatch malware with unprecedented levels of sophistry and impact. As the number and variety of malware types and families grow, the threat landscape in the cloud gets chaotic, which undermines an organizations' efforts to stabilize their security posture.

In September 2018, researchers reported the Xbash malware family that targets Linux and Microsoft Windows servers. Xbash was attributed to the Iron Group, a threat actor group publicly linked to other ransomware campaigns in the past including Remote Control System (RCS). Xbash has botnet, ransomware, and coinmining capabilities.

Much like WannaCry or Petya/NotPetya, Xbash can self-propagate like a worm

in an organization's network by targeting weak passwords and unpatched vulnerabilities. It contains functionality that, when activated, can accelerate its self-propagation rate. Xbash's ransomware capabilities can destroy data in Linux databases. To date, there is no evidence of any Xbash functionality that could make recovery possible even after a ransom payment.

Mitigations

- Enforce strong, non-default passwords as part of IAM.
- Maintain strong patching hygiene to stay up-to-date on security updates.
- Enable endpoint security, e.g. Host-based IDS, file integrity monitoring and detection.
- In on-premises VMs and in public cloud IaaS, most workloads run a single application. Application control and whitelisting to permit/deny executables attempting to run on a server can effectively secure the runtime environment.
- Use an effective backup and restoration process.

7) CLOUD SERVER AND ORCHESTRATION COMPROMISE



Workloads running in a single cloud, multi-cloud, or hybrid cloud environments present various degrees of the attack surface. Containers and Kubernetes allow new workloads to be deployed in minutes. However, it also implies that the attack surface can change more quickly than you can change the security. Docker and Kubernetes are primarily container orchestration tools with limited container security capabilities. For example, they don't monitor container content or pod-to-pod communications for malicious behavior. In 2018, in a mega-security event, a leading automaker company's

unprotected Kubernetes console was compromised for cryptomining. In one specific Kubernetes pod, access credentials were exposed to the company's AWS environment. To conceal their identity and remain unnoticed, the attackers used servers placed behind a content delivery network to run their scripts. The attackers gained access to sensitive telemetry data stored in an Amazon S3 (Amazon Simple Storage Service) bucket. Current trends around orchestration platform vulnerability indicate an increase in the frequency of this type of attack.

In the event of a cloud server compromise, the attacker gains access to some or all of the resources on a given server. The source of the compromise can come from insider threats, exploits/malware, misconfigurations, and cloud service provider account compromise.

Mitigations

- Require complete security posture encompassing cloud service provider security, DevOps security, and run-time security.
- Monitor public cloud environments for risky configurations.
- Use IDS tools capable of inspecting server and workload activities at runtime. AI capabilities to detect behavioral abnormalities provides dynamic threat defense/control in the cloud.

8) REMOTE CODE EXECUTION

Remote Code Execution (RCE) is a very common infection vector in the public cloud and has frequent occurrence with so many technology stacks and new CVEs released every week. RCE is a vulnerability where threat actors (including disgruntled insiders) can execute code remotely in a cloud server.

Lacework Labs recently came across an RCE attack in a server exposed to the open Internet. The server was running an older version of Redis (2.8.4) on Ubuntu 14.04 without enabled authentication for incoming connections. The Redis server was quickly exploited by LUA vulnerability CVE-2015-4335. The exploit contained payload to download install script, which then downloaded backdoor, a miner that gathered information about certain processes to report back, killed competitive miners, and modified host files and crontabs. This is a good example of the extent to which RCE can expose servers.

Without run-time monitoring, an RCE attack can remain undetected for an extended period of time. Account masquerading/impersonating users, deleting account and activity trails are some of the common and complex techniques attackers are employing to hide their trail. Besides, due to increasingly shorter life spans of workloads, necessary information may disappear by the time an alert is flagged. This makes threat detection and incident response increasingly challenging and costly.

Mitigations

- Use behavioral baselining to detect malicious code execution early.
- Patch early and often.
- Control network access to services.
- Have incident response plans in place for 0-days (there will always be new exploits).
- Reduce the size of the attack surface

9) CONTAINER ESCAPE VULNERABILITY


As containers climb in popularity, Container Escape Vulnerability is an increasingly serious concern. Containers are less of a sandbox than VMs, where containerized applications share host resources, and an escape can lead

to attacks on other containers. A vulnerability that allows escaping from a sandbox or container can result in access to the host operating system or hypervisor.

CVE-2019-5736 is the first major container escape of its kind where the execution of malicious containers allows for container escape and access to the host filesystem. Root user in a container or specially crafted container could overwrite runc binary (used in most container platforms, most notably Docker) with the new binary of their choosing.

RUNC CONTAINER ESCAPE VULNERABILITY

- CVE-2019-5736: Execution of malicious containers allows for container escape and access to host filesystem
- First major container escape of its kind
- Root user in container or specially crafted container could overwrite `runc` binary with new binary of their choosing
- `Runc` used in most container platforms, most notably Docker



All Your **Linux Containers** Are Belong to Us

`runc`

kubernetes docker

Lacework

(Figure Source: <https://www.slideshare.net/Lacework/lacework-top-10-cloud-security-threats/>)

A lot of attacks could be avoided if the software has the latest patches. This remains an unresolved problem as security teams struggle with keeping software up to date, especially given all the open-source and software resembling a lego building instead of core development. Running vulnerability scanners in a dynamic runtime environment in the cloud also presents a challenge.

Mitigations

- Adhere to an effective vulnerability management discipline for the cloud. 0-days are difficult to detect.
- Prepare for rapid response to update container platforms and operating system to keep pace with new vulnerabilities.
- Follow container best practices (privileged container policy, read-only root filesystem, etc.) to minimize the chance of successful escape.

10) RANSOMWARE

Over the last five years, ransomware has significantly gained in prominence as an indomitable challenge for enterprise assets hosted in the public cloud. Ransomware is essentially a malware that encrypts files and asks for payment to decrypt those files. However, some ransomware doesn't unlock files even after a ransom payment.

Since August 2018, the ransomware Ryuk has targeted large enterprises for a high-ransom return. Hermes is commodity ransomware used by multiple threat actors.

Mitigations

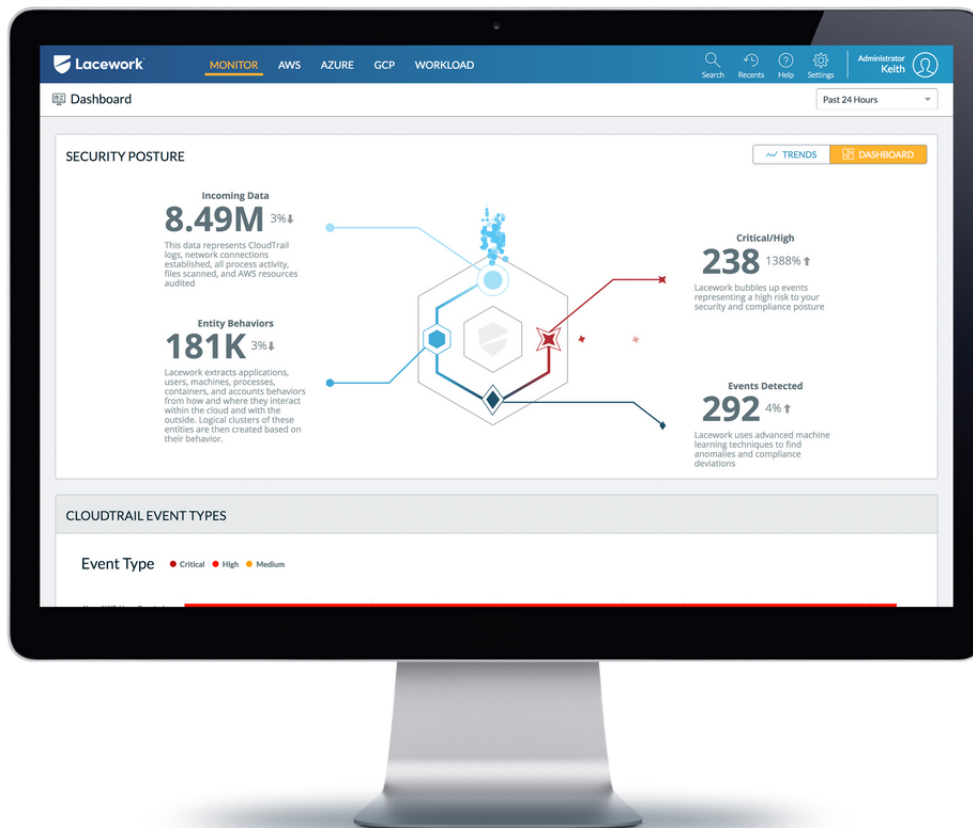
- Maintain a good security posture, especially in the cloud.
- Keep application up-to-date, enforce strong passwords, endpoint security, network monitoring, and threat intelligence.
- Use disaster recovery and backups.
- Secure runtime environment with process-level visibility and behavioral baselining to detect malicious activity early enough.

CONCLUSION

The trends discussed in this report suggest an increasingly dynamic threat landscape in the cloud. To cope with new security challenges, security teams are forced to update their security posture and strategies because traditional security tools and approaches are not well suited for the challenges of dynamic, virtual and distributed cloud environments.

The shift to cloud-native application development using container-based application architectures, microservices-based applications, and adoption of serverless PaaS will require new security capabilities both in development and at runtime. Cloud-native apps require solutions designed to address the protection requirements of cloud-based systems.

By 2022, 60% of server workloads will use application control in lieu of antivirus, which is an increase from 35% at YE18. Careful attention to the mitigation plan outlined in this article, coupled with appropriate security tools to automate aspects of cloud security, can empower your organization to effectively identify and address threats. Through continuous awareness, issues can be addressed upon discovery, with greater odds of thwarting attacks and strengthening the overall security of your cloud environment.



Get actionable recommendations on how to improve your security and compliance posture for your AWS, Azure, GCP, and private cloud environments.

Streamline security for AWS, Azure, and GCP. Gain unmatched visibility, ensure compliance, and enable actionable threat intelligence.

FREE ASSESSMENT