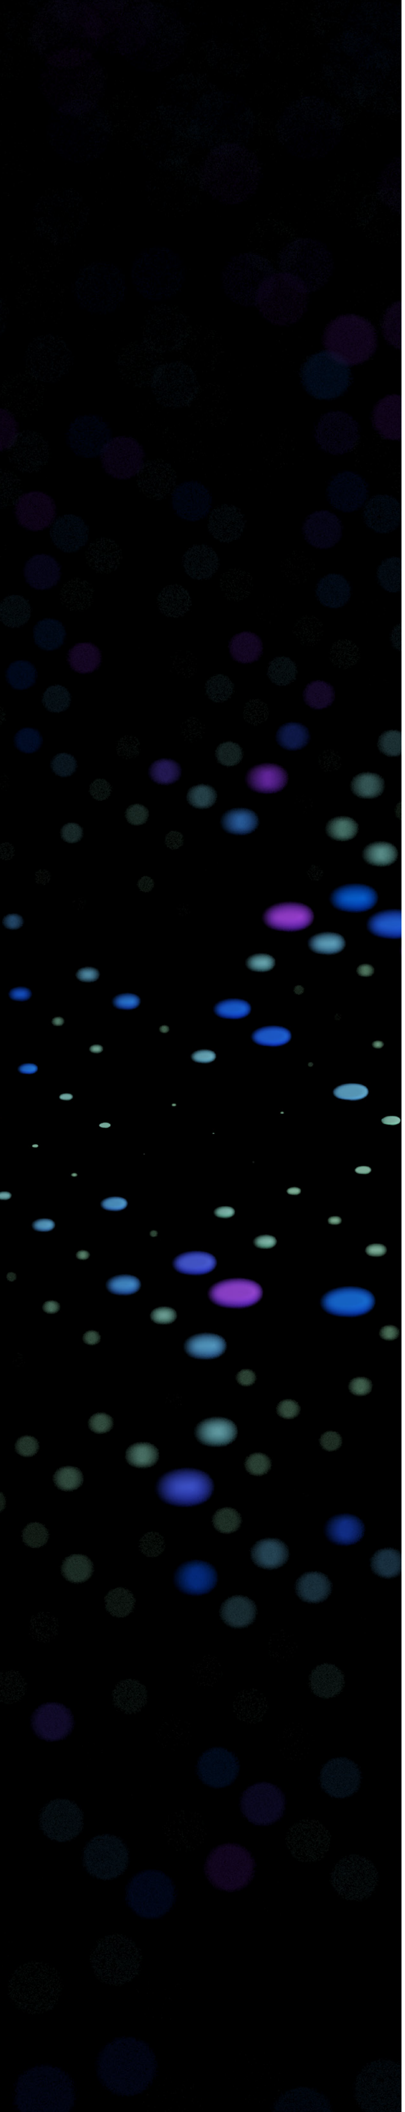

X-Force

Cloud Threat Landscape Report 2020

IBM Security X-Force® Incident Response
and Intelligence Services (IRIS)

Special Intelligence Report Q2 2020

IBM Security

Table of contents

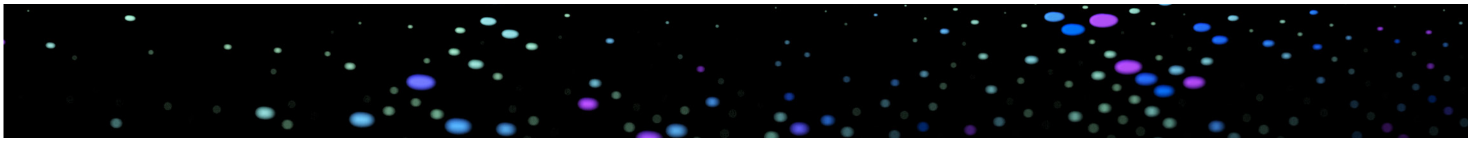

Introduction

With organizations increasingly using cloud environments for scaling and accelerating operations, understanding the unique cyber threat landscape of this field is critical. Leveraging IBM’s global presence and cloud incident response capabilities, we have conducted an in-depth analysis of cloud-related cybersecurity incidents our team has responded to over the past year in order to discern the top threats in this arena. This paper will dive into what IBM Security sees on the front lines of defending the cloud, including:

- Who is targeting cloud systems
- How we’re seeing threat actors access cloud environments
- What threat actors are doing once they have gained access to a cloud environment
- Common shortcomings we observe in cloud security
- Recommendations for improving your organization’s cloud security posture

Key findings

Financial gain is the most common motivation of threat actors targeting cloud environments, based on IBM Security Incident Response data collected since 2019.[1]

45%

Bruteforcing and exploitation of cloud applications are the two most common infection vectors, accounting for 45% of the cases examined in this report.

Data theft—such as appropriation of personally identifiable information (PII)—is the favored activity of cybercriminals once they penetrate a cloud environment.

>1 billion

Misconfiguration of cloud environments led to over one billion lost records in 2019.

Ransomware is the most commonly deployed malware in infiltrated cloud environments, accounting for three times as many cases as cryptomining and botnet malware, which follow in second and third place respectively.

Key findings

Leveraging cloud platforms for use as malicious infrastructure is often a favorite ploy of sophisticated threat actors, enabling them to ramp up operations with a single compromise. This has the added appeal of allowing them to minimize their own costs at the expense of their targets and appear to originate from otherwise legitimate sources.

Losses of >$50,000/hr

Depending on the organization impacted and the type of applications run in the cloud, infiltration can bring this swift and hefty price tag.

Redeploy assets, don’t reimage them: Organizations that redeploy assets vs. reimage affected cloud environments are more capable of performing effective forensic investigations, which may prevent subsequent harm to the organization.

Defense-in-depth is a necessity: Malware developers, aware of increasing cloud adoption, have begun making malware that disables many common cloud security products, leaving many companies unknowingly vulnerable.

Cloud computing: significant security benefits, but risks must be addressed

Cloud computing, the delivery of on-demand computing resources over the internet on a pay-per-use basis, provides a host of security benefits to organizations looking to expand their computing capability across the enterprise. While most enterprises today are approximately 20 percent into their transition to the cloud, as companies continue modernizing their core IT infrastructure and moving mission-critical data and apps to the cloud, organizations must also adapt to the unique cybersecurity challenges and opportunities posed by these new hybrid, multi-cloud environments.

There are key terms that need clarification when discussing cloud security, as these terms often come up while organizations work to secure their cloud-based assets.

Public cloud
A cloud environment hosted external to an organization.

Private cloud
A cloud environment hosted and maintained within an organization for their private use.

Hybrid cloud
Having clouds that are integrated and holistically managed — both on- and off-premises and across different clouds.

Containerization
Containers are an executable unit of software in which application code is packaged, along with its libraries and dependencies, in common ways so that it can be run anywhere, whether it be on desktop, traditional IT, or the cloud.

Corporate IT environments are getting increasingly complex as businesses incorporate cloud elements and need to manage their hybrid multicloud infrastructure in a way that is simple, consistent and integrated. But the architecture is not where it ends. Security threats inherent to the cloud's hyper-connected nature must be assessed and addressed as data and operations are moved to the cloud. Understanding the cloud security threat landscape can help organizations better protect themselves and prepare for potential security events.

Thinking about risk

As a result of increased cloud integration across enterprises, cloud-to-cloud interactions allow potential infections to spread across an enterprise even faster than on-premise attacks would and the vast amounts of data running in clouds can also increase the amount of data that threat actors could potentially steal. With data breach costs tightly linked with the number of records breached, cloud security incidents can be costly. Another aspect that can quickly add up costs is unauthorized access to cloud assets. Depending on the type of operation being breached, that access can easily generate losses of over $50,000 within less than one hour.

Unauthorized access to cloud assets can lead to

losses of >$50,000 in less than an hour.

Clouds are a new era for many organizations and securing them, addressing data privacy concerns, regulation and compliance, all demand a new and adapted approach with best practices that can differ from managing classic IT security.

Who is targeting cloud systems?

IBM Security has observed multiple types of groups targeting cloud systems through our IBM X-Force Incident Response and Intelligence Services (IRIS) incident response team, and working with our colleagues at DarkOwl, we've found even more interested parties on the deep web and dark web. Per our finding, financially motivated criminals are the top threat, but nation state actors are also a persistent risk.

Financially motivated criminals are the top threat group targeting cloud systems, but nation state actors are also a persistent risk.

Criminal cloud targeting

The most frequently observed threat actor category targeting cloud environments was cybercriminals, according to X-Force IRIS incident response data. Financially motivated threat actors were rarely associated with a specific group or organization, but rather simply trying to gain access to cloud assets in opportunistic attacks, oftentimes automating their attempts to cast a wide net. Further research done in collaboration with DarkOwl, a darknet data intelligence company, reveals a thriving market for cloud-targeted services and accounts on underground forums and markets.

Key offerings in the underground ecosystem include cloud service accounts at reduced prices. For example, one Russian-language forum featured vendors offering an unlimited monthly service for a large public cloud service that was cheaper than the public sale price. In another case, a different group of actors was observed selling access keys to another large public cloud service for as little as $15 each. These types of services can allow threat actors to use cloud infrastructure for their malicious activities under accounts that would otherwise appear legitimate.

Criminals also offer assistance in penetrating specific cloud-hosted accounts. On a platform that connects English and Russian speaking hackers, a user posted a link to a bypass utility to access personal accounts on a large public cloud platform. Another offering detailed the steps for exploiting a public cloud service's ID and included a script the hacker was selling on the forum. Another market offering observed in dark and deep web venues is the sale of cloud computing resources. Cloud provider access is sold to those wanting to host phishing sites, for example. In one case we observed a hacking forum member express their plans to host phishing pages on a public cloud environment that mimicked a product offered by the cloud provider.

Similar to legitimate online learning courses, there are multiple tutorials on how to use cloud computing credits or stolen credit cards to open accounts on popular cloud services, to then use them for malicious purposes.

Nation state attacks

Nation state threat actors have an established record of targeting cloud environments for financial gain and espionage purposes. Threat groups often target cloud environments extensively in an effort to access connected third-party systems.

As more sensitive data moves into cloud environments, we expect espionage-focused threat groups to continue targeting cloud environments to achieve their strategic intelligence objectives.

How threat actors compromise cloud environments

Like any network, cloud environments can be targeted in a variety of ways, some of which are unique to cloud but many of which affect systems across the board.

The IBM Security threat intelligence team leveraged X-Force IRIS's incident response data to find the most common ways we see threat actors targeting cloud environments. In many cases, threat actors use more than one of the following mechanisms to gain, maintain, and spread access, as well as escalate their access permissions inside the same cloud and on potentially connected clouds.

Exploitation of cloud applications

The most common infection vector we observed in cloud environments from January 2019 to May 2020 was remote exploitation of cloud applications, accounting for 45 percent of the cloud-related cybersecurity events examined. While agile and scalable, cloud-based applications can open up a broad avenue for threat actors. Since these applications are also often necessary for business operations, they are a key target for threat activity.

45%

Remote exploitation of cloud environments was the most common infection vector observed, accounting for 45% of cloud-related cybersecurity events we examined.

Over the past two years, IBM Security has responded to multiple incidents where vulnerable applications were present in the environment but remained undetected. Sometimes that is due to lack of security maturity in the cloud, and often it can also be due to “shadow IT,” where an employee goes outside approved channels to use cloud capabilities, opts to use a vulnerable cloud application, and thereby puts the entire environment at risk.

Addressing remote vulnerabilities in cloud environments has been challenging, in part due to the lack of public cataloging of discovered issues. Until 2020, vulnerabilities in cloud products were outside the scope of traditional CVEs, meaning that vulnerabilities related to cloud infrastructure were rarely publicly disclosed or recorded over time. As a result, cloud environments may have been more vulnerable than previously realized, harboring a variety of unaddressed issues.

Misconfiguration exploitation

Based on data from IBM’s 2020 X-Force Threat Intelligence Index, threat actors took advantage of misconfigured cloud servers to siphon over 1 billion records from compromised environments in 2019. Misconfiguration of cloud environments and subsequent data leaks remains one of the greatest sources of record loss across the board and can quickly allow a threat actor to access and steal sensitive information from organizations where that sort of oversight may affect assets in the cloud.

>1 billion records stolen

Threat actors took advantage of misconfigured cloud servers to siphon over 1 billion records from compromised cloud environments in 2019

IBM X-Force Threat Intelligence Index

Cross-cloud compromise

Threat actors can sometimes breach cloud environments by infecting one cloud environment, then use the trusted connection to move laterally to other clouds and infect additional environments.

This cross-cloud compromise can be especially insidious, as cloud environments, especially large public clouds, often have high volumes of communication and can make this type of infection more difficult to detect.

In one X-Force IRIS incident response, threat actors were detected when larger than normal amounts of data were observed being transferred between clouds that were designated for different geographies. This type of attack allowed the threat actors to quickly move between large data repositories and cause harm throughout the target's enterprise while evading many detection mechanisms, hiding their activity in the overall operational activity.

Swimming upstream

In another type of vector, we have observed threat actors attempting to gain privileged access to cloud repositories by getting onto the underlying hardware.

This "swimming upstream" technique requires threat actors to gain initial access to the cloud environment, then access the underlying host, then down to the management system to move between client environments. Swimming upstream can sometimes mask threat actor activity as legitimate administration, since administration often needs to move data between instances as well, so distinguishing between the two can be complex. This technique was on display when the severe Perfect 10.0 vulnerability was disclosed in 2020. That flaw allowed for a threat actor to break hardware-based isolation in a cloud environment, making them able to intercept code, manipulate programs, and otherwise affect other users' activities hosted on the same hardware. That flaw was since fixed.

How threat actors are using the cloud to cause harm

While there are many theoretical techniques to cause harm once in a cloud environment, IBM Security has observed threat actors using a variety of traditional attack tactics to profit from this new technology's increased capabilities.

Ransomware, data theft, and cryptomining lead the way in harming organizations while leveraging clouds. However, hosting malware and scam sites on cloud environments, or leveraging access to worm into other clouds, has enabled actors to cause broader damage outside the organization's own scope, compounding the potential risk to connected parties as well.

Ransomware

In X-Force IRIS incident response cases we analyzed from 2019 and 2020, ransomware was by far the most common type of malware deployed in the cloud, representing three times as many incidents as any other malware.

Ransomware was used 3x

as much as any other malware deployed in the cloud, based on IBM incident response cases studied.

Unlike ransomware attacks on traditional networked endpoints, ransomware in the cloud can have a more disruptive impact and cause greater data loss. That can be due to the wider range of operations supported by cloud environments, their potential impact on critical applications, and the sheer amount of data moving through clouds every day.

In a ransomware incident to which X-Force IRIS responded, the infection resulted from a gap in cloud administration responsibilities between the infrastructure provider and their client. That gap potentially allowed the compromise to go undetected longer, increasing the cost to the organization, and highlighting the criticality of defining cloud security roles for both the provider and the client.

Data theft

Cloud environments host large troves of information, and this data can be stolen by threat actors and sold on underground marketplaces. From incidents that IBM X-Force IRIS has handled, we've seen that the types of data stolen can vary. For example, we observed threat actors stealing sensitive PII, including credit card numbers, from a breach of cloud data. In a different incident, we observed threat actors potentially stealing client-related emails from a compromised cloud.

Data theft was the second most common threat activity IBM X-Force observed in breached cloud environments in 2019.

The type of data stolen can vary based on threat actor motivations and sophistication, but in cloud environments, the amount of data available can be much greater, making the potential impact of a breach that much more damaging to the organization.

Cryptomining

When it comes to illicit cryptomining threats, especially those leveraging scaled infrastructure like clouds, IBM X-Force IRIS has seen multiple instances of threat actors using cloud environments to mine for cryptocurrency over the last year.

In one case in 2019, X-Force IRIS responded to an incident where a cloud server was infected with a cryptominer that subsequently attempted to spread laterally to connected machines. This sort of compromise can have a few impacts: for on-prem cloud environments, the organization can incur increased electricity costs and faster degradation of hardware components as well as performance impact that can be rather meaningful in industries like the financial sector.

With external/public clouds, the organization could face increased charges due to higher data usage or slower response times as processing power is reduced. In all cases, cryptomining is a drain on organizational resources and can hinder business operations.

Hosting malware or malicious sites

Threat actors can use infected cloud environments to host malware that can subsequently spread onto other environments. For example, in late 2019, criminals were hosting a payment card skimmer on a cloud platform, which was subsequently downloaded onto targeted machines. In another case, threat actors hosted over 200 tech support scam sites on one cloud instance, directing users to these sites and using the cloud environment to lend their attack an air of legitimacy.

Hosting malware or malicious sites on cloud environments can help threat actors avoid network blocks by appearing as callouts to legitimate infrastructure. Also, the use of cloud hosting can provide threat actors with a layer of abstraction, making campaign activity more difficult to track. If an organization is unknowingly hosting malware on its cloud, especially for extended periods of time, this can lead to both direct and reputational harm as the organization may lose data itself and also be accused of allowing this activity to persist.

DNS compromise

IBM X-Force IRIS observed multiple incidents in which threat actors compromised cloud-hosted DNS services to redirect employees to different sites.

Playing on the DNS cache poisoning concept, this insidious attack method leverages existing cloud access to cause further harm to an organization and can be difficult for users to detect. This type of compromise can redirect users to sites that attempt to exploit their browsers and drop a malicious payload onto endpoint machines, or it could redirect them to phishing sites to steal their network credentials. In some cases, redirection to advertisements, or click-fraud, is set up to line the pockets of criminals.

While DNS compromise is not a new attack, as organizations continue to shift these services to an external cloud provider, threat actors can find new avenues for potential compromise, and the impact could expand throughout the organization.

Lateral spread

Threat actors have used a variety of methods to expand from an initial infection into other parts of the cloud environment or on to endpoint boxes that access cloud resources. In 2019, IBM X-Force IRIS responded to an incident in which malware deployed to a cloud environment attempted to spread to other machines via SSH bruteforcing, affecting external local machines accessing the cloud.

Also in 2019, the Exim worm, a Linux worm spreading via Exim servers, was reportedly spreading infection automatically via remote exploitation of CVE 2019-10149. The worm would take over a server with the goal of dropping a cryptojacking malware onto it.

This type of lateral spread can compound the impact of a cloud infection and bring it into an organization’s internal network space.

Cloud-native malware and its evolution

Many cloud-based systems run the same operating systems and applications as their on-premise counterparts, and as a result, much of the malware found operating in cloud environments is the same as that found outside of the cloud. However, there are instances of malware which are specifically designed to either target or make use of cloud systems.

These malware variants fall into three groups:

— Malware that uses clouds to scale

— Malware adapting to cloud environments

— Malware using cloud environments for operational infrastructure

Malware that uses clouds to scale

By targeting cloud applications or platforms specifically, malware operators can quickly ramp up operations and potentially gain significant profits with a single compromise. One example of a malware family adapting to the new cloud reality is DemonBot, a Linux-based bot which was reported in October 2018. DemonBot targets cloud servers running Hadoop and infects them via a vulnerability in Hadoop's resource management tool.

X-Force IRIS investigated an instance of DemonBot found in a cloud environment which was detected by the affected organization due to a significant increase in billing and resource usage. In that case, the DemonBot binaries' primary function was to launch distributed denial of service (DDoS) attacks as part of a botnet. With the addition of cloud-targeting capability, malware operators are able to amplify these attacks using cloud resources.

Another malware family that has found use in organizations' cloud adoption is Graboid, a cryptomining worm identified in October 2019. This malware targeted and compromised unsecured Docker hosts, onto which it downloaded malicious Docker containers. The malicious containers performed cryptomining and also spread the malware to other hosts.

In June 2019, researchers reported on an attack campaign which also targeted misconfigured Docker hosts in order to infiltrate them. In this case, it was a Linux botnet malware known as AESDDoS that exploited API misconfigurations to deploy the malware on vulnerable containers. This malware was then able to receive commands from a command control server and launch a variety of DDoS attacks.

Malware adapting to cloud environments

During the last two years, researchers at Intezer have observed a significant increase in the number of cyber-attacks targeting Linux servers, mostly in cloud environments. The Linux operating system accounts for nearly 90 percent of all cloud servers.

One example of a threat actor that uses malware to target clouds is the Chinese-affiliated Pacha Group. This group has been targeting cloud-based infrastructures with new, previously undetected malware variants of Linux. GreedyAntd, which shares significant amounts of code with previous variants. Another malware targeting Linux-based file storage systems (NAS servers) in cloud environments is the QNAPCrypt ransomware. This sort of threat can affect a very large user base and affect damage on troves of data hosted in the cloud.

As cloud environments continue to increase in popularity, Linux-focused malware is likely to continue to grow.

Malware using cloud environments for operational infrastructure

Just as organizations scale up operations, malware distributors, especially those linked with organized crime and nation state attacks, can opt to scale their operations as well. X-Force IRIS has observed threat actors leveraging cloud environments for their malware operations in a variety of ways.

One investigation of RokRat, a remote access tool (RAT) observed to target victims in South Korea and attributed to ITG10 (AKA APT37, Scarcruft), used legitimate, commercial cloud storage services for hosting payloads and C2 communications. This use of cloud services for infrastructure can be difficult to detect, as organizations' networks can have significant cloud communication as part of normal operations and whitelist legitimate providers. By hosting malware in that sort of a cloud environment, detection of malicious payload downloads can also be more challenging.

Much like RokRat, Karae is another backdoor attributed to ITG10, which is also known to use cloud storage providers for C2 communications. In Karae samples analyzed by X-Force IRIS, a legitimate cloud storage service provider was used by the threat actors to host the malware, with account credentials hardcoded into the malware's binaries. Karae gathers information about the victim's system and writes it to a file, which it then uploads to the cloud. It also attempts to download and execute additional binaries from this service.

This tactic overall presents a detection problem because it can blend in with normal and legitimate user activity, and it's also the reason that malware operators of varying motivations opt to use it.

Recommendations to enhance your cloud security

Responding to a security incident in a cloud environment requires some special considerations outside of the usual method of incident response. From X-Force IRIS's extensive experience in this field, we've compiled some key lessons we've learned while preparing for, and responding to, cloud incidents.

Preparing a safer cloud environment

Begin with the end in mind

Before considering moving workloads or data to the cloud, develop a plan as to its purpose. Build security controls into the conception process and take into the account the criticality and sensitivity of operations running in the cloud. Consider using a partner that provides comprehensive security services to help you gain visibility and control of all aspects of your hybrid cloud security as you develop your program.

Use proactive simulation

Simulate both expected and unexpected security events within your cloud environment to understand the effectiveness of your preparation. This preparation can provide an opportunity to exercise internal playbooks and standard operating procedures. By focusing on testing and improving both technical and operational response skills, organizations can move quickly to remedy issues before damage is done or expands. Furthermore, these exercises can be augmented with tactical information, such as indicators of compromise (IOCs), to support a threat intelligence-driven response scenario.

Prevent policy “dead spots”

With external clouds, the responsibility for protecting cloud environments often falls on both the organization and the cloud hosting provider. Cloud hosting is rarely a set-and-forget service, and requires security both on the part of cloud service providers and the organization using the service. Delineating the role for each party at the time of negotiating the contract can help define responsibility, controls, monitoring, and potential liability before an incident. This can help prevent incidents resulting from policy gaps, as well as allow for more efficient detection and incident response.

Apply best practices to cloud security

Cloud security does have its own approaches, but they also resemble the security of other networks in some respects. Organizations therefore need to apply security best practices to their cloud environments, as clouds are not immune to compromise. To mitigate the threat of unauthorized access, implementing multifactor authentication can help prevent infiltration using stolen credentials.

Privileged account management (PAM) is another important concept in keeping clouds better protected. Restrict accounts to least-required privileges to minimize harm from account compromise and consider using the zero-trust model. By implementing these practices in cloud environments, organizations can mitigate the risks of incidents or reduce the impact of potential security events.

Monitor and log

Cloud environments need to be monitored for a variety of matters. From cloud-sprawl to third party access and unexpected outages, proper monitoring can help detect malware and the first sign of an attack. Cloud users should maintain robust logging of events in cloud environment for forensic investigation of malicious activity. The responsibility of monitoring and logging of cloud events should be decided by the organization and the cloud hosting provider prior to starting services to prevent policy dead spots.

Use threat intelligence to monitor threats

Threat actors continue to evolve and augment their existing arsenal of tactics, techniques, and procedures with new capabilities specifically to target cloud environments. As these capabilities continue to develop, organizations should leverage threat intelligence to monitor changes in targeting and implement effective defenses.

Responding to cloud security incidents effectively

Don't reimage, redeploy instead

When organizations terminate their cloud instances, they lose potentially valuable forensic artifacts. Instead of immediately destroying this data, isolating affected systems and standing up known clean images could allow forensic investigators to analyze the infected instance and potentially find additional leads to understand what went wrong and how to prevent it in the future.

In the case of an incident, build a workstation in the cloud for your investigators to perform the work on, then create an image from the compromised servers and collect volatile memory data. Also, doing root cause forensic analysis can prevent IT administrators from redeploying a tainted base image.

Remember bandwidth costs

One challenge organizations face nowadays is that downloading a large server image after an incident may be cost prohibitive due to the high bandwidth costs associated with moving data out of the cloud environment. Having relevant policies or requirements in place prior to an incident can help reduce the cost of download, especially when an infection could potentially harm other instances on the same cloud.

Have the right investigative tools

Cloud security requires the right tools to conduct a thorough investigation if something goes wrong. Many common incident response and forensic tools are only effective in a local environment or an on-premise hosted servers and are not equipped for a cloud environment. However, preparing the right toolkit can enable effective cloud investigations.

Moreover, organizations should include cloud assets in their overall incident response plans and test their cloud security incident response at a tactical level to ensure that the tools they have are capable of working across all the cloud environments they use.

Automate incident response

Implementing effective security automation in cloud environments can improve detection and response capabilities, rather than relying on manually reacting to events. For example, by following the Infrastructure as Code (IaC) approach and using tools such as CloudFormation, declarative approach, and the serverless, event-driven Lambda services, a compromised organization can efficiently rebuild the environment from a predefined template. This methodology can lead to accelerated recovery during ransomware or destructive cyber-attacks against that environment.

Conclusion

For those who contract their infrastructure from a vendor, cloud security is a joint effort by both the provider and the user of cloud services. Organizations that operate in the cloud must be aware of the threats to cloud environments in order to properly secure their data and services.

IBM X-Force research indicates that threat actors are also keenly aware of organizations shifting to cloud infrastructure and are evolving accordingly. With the cost of cloud breaches only continuing to grow, organizations need to take steps to protect their cloud-based assets. By being proactive and taking the recommended steps, organizations can better protect themselves as they move into a cloud-based world.

About IBM X-Force

IBM X-Force studies and monitors the latest threat trends, advising customers and the general public about emerging and critical threats, and delivering security content to help protect IBM customers. From infrastructure, data and application protection to cloud and managed security services, IBM Security Services has the expertise to help safeguard your critical assets. IBM Security protects some of the most sophisticated networks in the world and employs some of the best minds in the business.

Learn about IBM X-Force IRIS

Endnotes

1. Methodology Caveat: The statistics cited in this paper used a subset of X-Force IRIS's incident response reports between June 2018 and March 2020; this limitation was required due to a variety of reasons including privacy concerns. As a result, the statistics contained herein, while reflective of broader trends observed in cloud security during this timeframe, may be affected by a degree of collection bias.

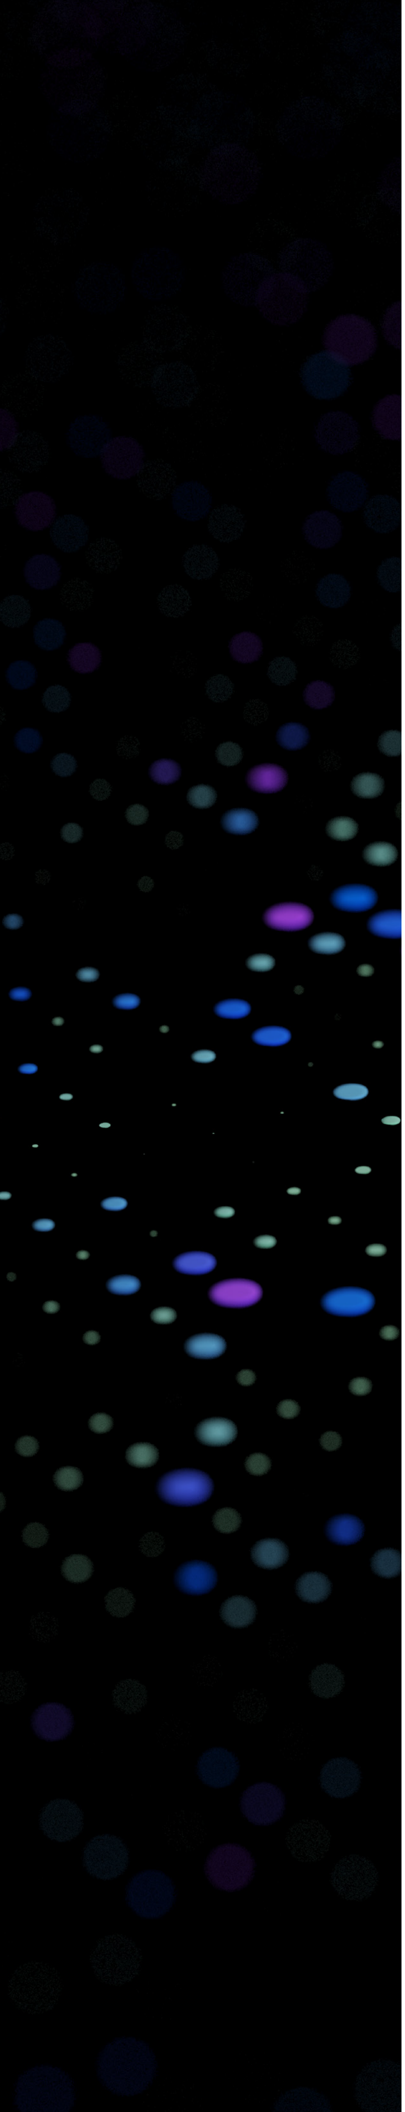

IBM Security

Contributed research

Intezer
DarkOwl

IBM Security
New Orchard Rd
Armonk, NY 10504

Produced in the United States of America
May 2020

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates. THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.