



Threat detection and response in the cloud

TABLE OF CONTENTS

Uncharted territory: Detecting threats in the cloud.....	3
Attack lifecycle in the cloud	3
Top cloud security threats	4
Analysis of a real cloud attack.....	5
The Cloud Hopper attack lifecycle	5
Shared responsibility model.....	6
Key takeaways	7
Managing access	7
Detect and respond	7
Security operations.....	7

Uncharted territory: Detecting threats in the cloud

Cloud environments change fundamental assumptions in how to perform threat detection and response.

The highly dynamic inventory of cloud workloads means systems come and go in seconds. A heavy focus on automation amplifies the potential for human error in system configurations. Shared responsibility with the cloud service provider (CSP) creates potential threat detection gaps in the attack lifecycle.

Everything in the cloud is moving to an API data access method, and traditional approaches to monitoring traffic flow no longer apply.

In addition to challenges in threat detection and response, the pace of innovation in the cloud leaves businesses constantly behind. Increasing business competition means organizations focus more on shipping features first and outsourcing non-core capabilities business models – often at the expense of information security.

An explosion of cloud services means the concept of a perimeter is gone and using perimeter controls becomes futile. A growth of new infrastructure and deployment tooling results in new environments with new security models and attack surfaces. And finally, the existing shortage in security expertise becomes amplified with all the newly released features and services.

Most critically, the introduction of multiple access and management capabilities creates variability that adds significant risk to cloud deployments. It is difficult to manage, track, and audit administrative actions when those users can access cloud resources from inside or outside the corporate environment.

Traditionally, accessing a server required authentication to the organization perimeter and monitoring could be implemented inside the private network to track and monitor administrative access.


Attack lifecycle in the cloud

Attackers have two avenues of attack to compromise cloud resources. The first is through traditional means, which involves accessing systems inside the enterprise network perimeter, followed by reconnaissance and privilege escalation to an administrative account that has access to cloud resources.

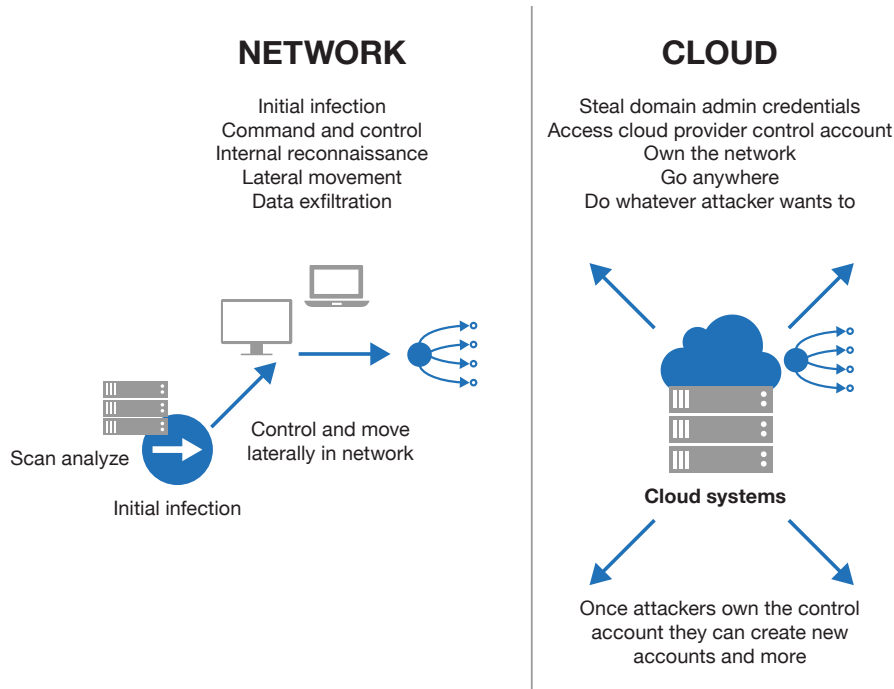
The second involves bypassing all the above by simply compromising credentials from an administrator account that has remote administrative capabilities or has CSP administrative access.

This variability in administrative access models means the attack surface changes with new security threats via unregulated access to endpoints used for managing cloud services. Unmanaged devices used for developing and managing infrastructure exposes organizations to threat vectors like web browsing and email.

When the main administrative account is compromised, the attacker does not need to escalate privileges or maintain access to the enterprise network because the main administrative account can do all that and more. How does the organization ensure proper monitoring of misuse of CSP administrative privileges?



CYBERATTACK LIFECYCLE



Organizations need to review how the system administration and ownership of the cloud account is handled.

1. How many people are managing the main account?
2. How are passwords and authentication performed?
3. Who is reviewing the security of this important account?

Who is at fault if there is a security problem? The CSP or the cloud tenant organization? Initially it seems to be dependent on the problem, but some CSPs want to push that responsibility to the tenant organization.

Most importantly, how does an organization monitor for the existence and misuse of administrative credentials? A lack of visibility to back-end CSP management infrastructure means cloud tenant organizations need to identify misuse of CSP access within their own environments when used as a means of intrusion.

Top cloud security threats

In 2017, the Cloud Security Alliance (CSA) conducted a survey to compile professional opinions about what it believed at the time to be the most pressing security issues in cloud computing.

Of the 12 identified concerns, five were related to managing credentials and methods of compromising those credentials to gain access to cloud environments for malicious intent. Those five, in order of severity per survey results, are:

1. Insufficient identity, credential and access management – Lack of scalable identity access management systems, failure to use multifactor authentication, weak passwords, and a lack of ongoing automated rotation of cryptographic keys, passwords and certificates.
2. Insecure interfaces and APIs – From authentication and access control to encryption and activity monitoring, these interfaces must be designed to protect against both accidental and malicious attempts to circumvent policy.
3. Account hijacking – Attackers can eavesdrop on user activities and transactions, manipulate data, return falsified information and redirect your clients to illegitimate sites.
4. Malicious insiders – A current or former employee, contractor or other business partner who has or had authorized access to an organization's network, systems or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity or availability of the organization's information or information systems.
5. Insufficient due diligence – Not performing due diligence exposes a company to a myriad of commercial, financial, technical, legal and compliance risks that jeopardize its success.

Analysis of a real cloud attack

The APT10 group has been credited for a tactical campaign known as Operation Cloud Hopper, a global series of sustained attacks against CSPs and their customers. These attacks aimed to gain access to sensitive intellectual and customer data.

US-CERT noted that a defining characteristic of Operation Cloud Hopper was that upon gaining access to a CSP, the attackers used the cloud infrastructure to hop from one cloud tenant to another, gaining access to sensitive data in a wide range of government and industrial entities in healthcare, manufacturing, finance and biotech in at least a dozen countries.

The Cloud Hopper attack lifecycle

In Operation Cloud Hopper, attackers initially used phishing emails to compromise accounts with access to CSP administrative credentials. This is the most common method of infection for any attack and is still the easiest way of getting initial access to a network. The attacker would leverage malware designed to collect the necessary credentials to pivot directly into the CSP and client managed infrastructure.

Once access is attained on the management infrastructure, PowerShell could be used inside client managed infrastructure for command-line scripting to perform reconnaissance and gather information used for lateral movement to get access to additional systems.

The attackers continued to leverage compromised credentials to cross security boundaries, effectively using cloud service providers as a step to gain access to corporate data of multiple organizations.

To ensure persistent connectivity to the cloud infrastructure in the event an administrative account no longer worked, the attackers installed remote access trojans for command and control to sites spoofing legitimate domains.


These were open source, off-the-shelf malware used in many attacks like Poison Ivy and PlugX. Many of the systems compromised with remote access were non-mission critical, which could be used to continue lateral movement and avoid detection by system administrators.

The final stage of Operation Cloud Hopper was data exfiltration of intellectual property. Data was collated, compressed and exfiltrated from the CSP infrastructure to the infrastructure controlled by the attackers.

As CSPs take on responsibilities from tenants in the managed infrastructures, the amount of control and visibility those cloud tenants maintain diminishes. APT10 took advantage of this diminished visibility and leveraged credentials and systems that had access to both CSP and enterprise infrastructures.

Because cloud tenants do not have visibility or control in the CSP infrastructure itself, it is a formidable challenge to monitor and detect attackers who access one system then quickly pivot within the CSP infrastructure to access another system.

It is important to note that the complexity of hybrid environments that involve CSPs and on-premise systems makes it difficult to adequately address problems like stolen credentials or lateral movement by attackers from a cloud tenant to a CSP and then to a second cloud tenant. One careless and inattentive cloud tenant can increase the risk for other cloud tenants who exercise greater diligence.



Shared responsibility model

Ensuring threat detection and response capabilities in cloud environments starts with a basic understanding of the shared responsibility model and the impact that model has on security management and monitoring capabilities.

The security of cloud services is a partnership and a shared responsibility between cloud tenants and the CSP. The CSP is responsible for the cloud platform and the physical security of its data centers.

Tenants own their cloud data and identities, the responsibility for protecting them, the security of on-premises resources, and the security of cloud components over which they have control. CSPs deliver security controls and capabilities to help protect data and applications, and the degree of tenant responsibility for security is based on the type of cloud service.

The level and balance of control by CSPs and cloud tenants depends on the computing model used. The model below provided by Microsoft for Azure illustrates the level of shared responsibility based on a cloud platform.

On-premises deployments involve data centers that leverage a virtualized infrastructure owned by the enterprise. In this model, an enterprise is responsible for the entire security stack, from physical devices to data.

An infrastructure-as-a-service (IaaS) virtual data center model replicates existing internal data centers. In this instance, physical segregation of hardware is not possible and requires hypervisor-level capabilities to create security zones and for remote access.

When choosing between managing the infrastructure in a private or public cloud, most organizations find themselves with a hybrid cloud, a combination of the private and public cloud with shared resources and distribution components. Usually, the critical back-end infrastructure is private and the access and distribution is public.


Security and compliance concerns are first-order priorities for virtualized data center and cloud deployments. Security requirements for virtualized data centers and clouds include the ability to monitor virtualized environments while maintaining the highest levels of VM host capacity and performance. Techniques include hypervisor-based stateful firewall, network detection and virtualization-specific endpoint protection.

In a platform-as-a-service (PaaS) model, applications are installed and managed on existing outsourced platforms. A server can be provided for exclusive access or a server is shared between multiple applications.

Confidential information can be exposed to other users or to the service provider because no control is provided over existing hardware. Controls must be applied to the data within the applications and databases using encryption and external key management designed for virtual environments.

With software-as-a-service (SaaS), third-party applications like Salesforce are utilized to provide a specific service. Data is stored on the application providers' back-end using access controls they provide.

Enterprise applications now support integration with Active Directory using ADFS and SAML for communication. Controls must be provided for authentication and access management as well as monitoring to ensure the enterprise retains control over how these applications are used.



The Microsoft shared-responsibility model

Key takeaways

In the APT10 Operation Cloud Hopper attack, the method of initial intrusion was cloud specific, but the attack behaviors within those cloud environments were the same behaviors found in private cloud and physical data centers.

This is because all attacks must follow a certain attack lifecycle to succeed, especially when the goal is data exfiltration. Preventing a compromise is increasingly difficult but detecting the behaviors that occur – from command and control to data exfiltration – are not. More importantly, when an attack is carried out in hours rather than days, the time to detect becomes critically more important.

A key takeaway from the shared responsibility model is that regardless of the data center model deployed – infrastructure, platform or software as a service – the enterprise organization is always responsible for data, endpoints, accounts, and access management.

Managing access

While CSPs need to ensure their own access management and controls that limit access to cloud tenant environments, tenants themselves must assume this can be compromised and focus on learning the who, what, when and where of access management.

Properly assigning user access rights helps by reducing instances of shared credentials so cloud tenants can focus on how those credentials are used. Resource access policies can also reduce opportunities for movement between the CSP infrastructure and cloud tenants.

Detect and respond

When it comes to cloud and on-premises monitoring, it is necessary to monitor both as well as determine how to correlate data and context from both into actionable information for security analysts.

Monitoring cloud-deployed resources by cloud tenants is essential to increase the ability to detect lateral movement from the CSP infrastructure to tenant environments and vice versa.

Coordinating with the CSP – as well as CSP coordination with cloud tenants – can provide a powerful combination of information that can increase the likelihood of detecting the post-compromise activities.

More importantly, visibility into attacker behaviors is dependent on the implementation of proper tools that can leverage cloud-specific data.

Security operations

Knowing and managing the infrastructure as a part of due diligence should help to identify systems and operations that are compromised by malware implants like those used in Operation Cloud Hopper.

Changes to production systems can be difficult to detect. But when visibility is available in the cloud infrastructure, it is much easier to detect attacker behaviors in compromised systems and services that are clearly operating outside of expected specifications.

Ideally, security operations teams will have solid information about expectations for that infrastructure, so deviations from normal activity are more likely to identify malware and its activity.



Email info@vectra.ai Phone +1 408-326-2020
vectra.ai