XM Cyber
See All Ways™

CSA cloud security alliance℠

The Necessity Of Attack Path Management For The Hybrid Cloud

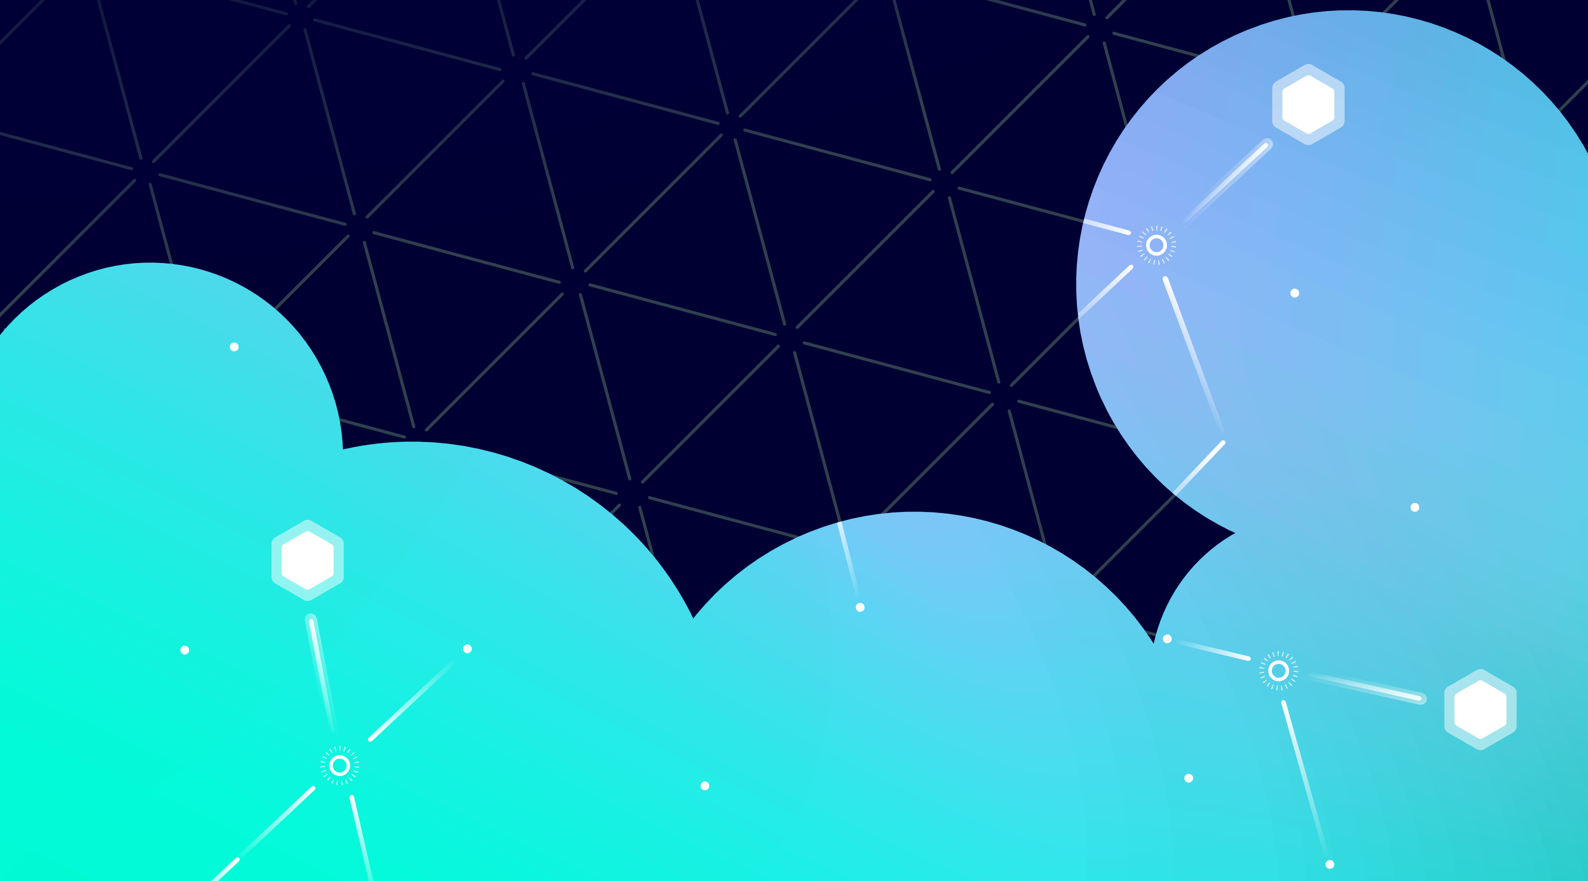

Table of Contents

Introduction

Summary

Introduction

Lifting the Fog, Understanding how to Secure the Hybrid Cloud

Accelerated Digital Transformation

Now it's no secret businesses have ramped up and driven the adoption of the cloud faster than any period previously. One of the key players in driving the mission of businesses all around the world to digitize is COVID-19. The necessity of operations to keep pace with the high demand of access has driven companies to accelerate their plans from a matter of years to months. In addition, remote working is now a common business practice which further complicates the issues of organizations protecting their most critical assets. This is due in part to the fact that our perimeter has dramatically expanded, and with a silver lining, kind of disappeared at the same time. No longer is protecting everything on the inside enough - companies are lifting and shifting workloads and critical assets from on premise, to the cloud and back again to match the needs of the users and services that require them. This has been made even more evident by the recent log4j / log4shell vulnerability disclosure that raises the need not only to look into the asset distribution but prioritize work on which one of them is more prone to attacks.

New security gaps are constantly being created due to new ways of working in a hybrid environment. Cyber attackers take advantage of this change to obtain the initial foothold and breach an organization leveraging misconfigurations, overly permissive identities, vulnerabilities, and human errors. But it's clear why everyone is moving to the cloud.
It offers increased agility, improved ability to collaborate, there's a minimal IT infrastructure to manage, evergreen services (particularly SaaS), predictable pricing, and it has the potential for increased automation. Research is showing that 90% of organizations are actually going to be using multi-cloud or have some type of hybrid strategy by 2022. Chances are you are already using the cloud and if not are on your way to it.

Whilst COVID-19 has driven some common cross-industry approaches to adoption of remote working tools like video-conferencing and collaboration, there remain differences in adoption patterns of other cloud services across industry verticals and technology capabilities. These differences are often driven by regulatory concerns, with the financial services industry being the canonical example of a heavily regulated environment. Financial services regulators across the globe are often risk-averse in the area of availability with systemic risk a key concern; regulatory requirements regarding resilience and exit planning are often seen to be forcing institutions to adopt not just multi-cloud approaches but full workload portability. You will quite often see major financial services organizations building complex IT stacks that allow containerized applications to be moveable across cloud environments, albeit at the cost of abstraction from the underlying cloud platforms (and native capabilities) and the management and maintenance headaches associated with technologies such as Kubernetes. Manufacturing industries have, in general, been quite slow to move to new technologies, however the long-foretold revolution of Industry 4.0 is beginning to become real and such businesses are now investigating the adoption of cloud services as part of the wider strategy of IT/OT consolidation. Perhaps surprisingly, Governments the world over have been enthusiastic adopters of cloud services, with initiatives such as Fed Ramp in the United States helping to improve security assurance across the wider cloud environment.

In this paper, we talk through the importance of understanding the attack paths a motivated attacker can use to try and compromise your on-prem, multi-cloud, and hybrid environments. Understanding the attack path enables organizations to identify potential choke points, monitoring requirements and architectural weaknesses that can be addressed in order to improve their overall security posture.

Top 6 Challenges of Securing the Cloud

Alongside offering new IT capabilities to consumers, or at least new ways of delivering those capabilities, the cloud also opens up new possibilities for attackers to exploit:

1. Unknown hybrid attack surface risk: When your environment expands, your attack surface expands with it. Every time a new cloud environment is added, whether as infrastructure as a service from either AWS, Google Cloud, Microsoft Azure – or software as a service that transmits data to or from, you have to manage risk within each of those environments and in the transitional spaces between them. Visibility and control across cloud infrastructures are the keys to enabling superior application security and reliable connectivity from client to server, regardless of location.

2. A highly dynamic environment: Hybrid cloud has the potential to offer more resource options via use of public cloud providers vs. an organization’s physical data center. This makes it easier to provision, deploy and scale resources to meet demand spikes. The connected nature of cloud environments makes the attack surface even wider. The line begins to blur when stitching together various PaaS capabilities to optimize application delivery. Between that and commonly used APIs that can be potentially insecure or poorly managed credentials, the threat continues to grow without a clear understanding of the security boundaries between various platforms.

3. Prioritizing large volumes of vulnerability data: If it seems as if vulnerability risk management is increasingly more challenging, it’s probably for good reason: the sheer number of vulnerabilities to manage can overwhelm even a strong security team. With COVID-19 disruptions and an ongoing cybersecurity skills gap complicating matters even further, it’s critical that organizations identify the right support tools to help get the job done.

4. No visibility of exploitable attack paths: In an ever-changing elastic environment, it is crucial to prioritize high impact risks and know what to fix first. Security and engineering teams must first tackle those vulnerabilities that are open to attackers, compromise critical assets or allow lateral movement within the systems / networks.

5. A lack of context of critical assets: It’s more important to understand the risk to an asset versus the risk on an asset. Understanding the difference means providing a direct ability to prioritize remediation that will make the biggest impact on risk. A prerequisite to understanding the context of your critical assets is understanding where those assets are located; a non-trivial task when organizational units are able to spin up resources with little or no central control or visibility.

6. Attackers are now aggressively targeting cloud environments: API attacks are now a threat vector many companies are facing. Due to the scale and number of configurations for AWS, Azure, and other cloud platforms, understanding the risk from the configurations and changes to the cloud provider security tooling is imperative. Threat actors are aware of the security challenges presented by the cloud and take advantage of just that. As organizations have shifted to incorporate remote work and more disconnected, hybrid multi-cloud environments, a zero trust strategy has the potential to help protect data and resources by making them accessible only on a limited basis and in the right context.

Other common cloud security challenges include:

- Malware leveraged by threat actors
- Insider threats, both malicious and unintentional
- Third-party risk – An organization does not control the infrastructure or applications used by their cloud services provider, whilst larger organizations also have managed services providers, e.g. application management or managed security providers, working on their behalf
- Weak tenant separation within PaaS services
- API risk – Cloud application integrations must be protected from the threat of actors attempting to intercept or redirect transactions

The Threat is Real

It doesn't matter if you're a small organization or large enterprise, 95% of cloud breaches occur due to human errors such as configuration mistakes* and 67% of breaches included the use of credential theft and misconfiguration**. Cloud providers, and the wide variety of managed services providers operating within those clouds, are tempting targets for advanced threat actors due to the number of high-profile enterprises making use of those services. There is increasing visibility of threat actor activity in the managed services provider space, with the activities of APT10/ Stone Panda being widely publicized by western national security agencies back in 2018 representing a significant turning point in transparency.

The Cloud Security Alliance has produced a number of reports related to cloud threats over the years, with the latest iteration of this work, titled, *Top Threats to Cloud Computing: Egregious Eleven.* The threats listed in this report include:

- Data Breaches
- Misconfiguration and Inadequate Change Control
- Lack of Cloud Security Architecture and Strategy
- Insufficient Identity, Credential, Access and Key Management
- Account Hijacking
- Insider Threat
- Insecure Interfaces and APIs
- Weak Control Plane
- Metastructure and Applistructure Failures
- Limited Cloud Usage Visibility
- Abuse and Nefarious Use of Cloud Services

Whilst we often talk about technical security threats, it is important to consider the implications of a lack of cloud security architecture and strategy. Cloud adoptions often fail due to a lack of understanding of organizational ownership, responsibility and accountability and the impacts of these deficits on the derivation of common security patterns. Strong enterprise governance is a pre-requisite to successful, secure, cloud adoption and operation.

Cloud environments can be ultra-sensitive to misconfigurations and escalation of privilege problems. Once the initial foothold is achieved, it can become possible to execute lateral movement to breach critical assets on-prem or in the cloud. Top 3 methods for Threat Actors to compromise the Cloud*:

- **Password spraying –** a hit and hope approach towards re-using compromised passwords across multiple systems and services, easily automated and can scale to target many organizations
- **Software vulnerability –** allow attackers to gain the initial foothold into an organization's environment
- **Pivoting from on-prem environments into cloud environments** almost 25% of incident responses occurred via this vector to deepen the hold on the organization and compromise even more resources

95%

Of cloud breaches occur due to human errors such as configuration mistakes*

67%

Of breaches included the use of Credential theft and miscofiguration**

*Source: 2021 IBM Security X-Force Cloud Threat Landscape Report
**Source: 2021 Verizon Data Breach Investigations Report

The Cloud Shared Responsibility Model

All cloud service providers adhere to a shared security responsibility model (whether formally documented or not), which means your security team is responsible for security as you move applications, data, containers, and workloads to the cloud, while the provider takes some responsibility, but not all. Defining the line between your responsibilities and those of the cloud providers you use is imperative for reducing the risk of introducing vulnerabilities into your public, hybrid, and multi-cloud environments.

Layer	Infrastructure-as-a-Service (IaaS)	Infrastructure-as-a-Service (PaaS)	Infrastructure-as-a-Service (SaaS)
Data	🟩	🟩	🟩
Application	🟩	🟩	🟦
Operation System	🟩	🟦	🟦
Virtualization	🟦	🟦	🟦
Servers	🟦	🟦	🟦
Storage	🟦	🟦	🟦
Network	🟦	🟦	🟦
Physical	🟦	🟦	🟦

🟩 Customer 🟦 Cloud provider

Cloud service providers are responsible for the security of its cloud infrastructure, while users are responsible for everything that they put into the cloud, including data, and for configuring their firewalls, servers and deployed code.

The defining lines for responsibility between provider and consumer are not always obvious. A great example of how blurred lines of responsibilities can be problematic was the security breach of a financial services giant's AWS environment*. In this incident, the consumer was responsible for the configuration of a hosted WAF and unfortunately this WAF was misconfigured, allowing an attacker to conduct a Server-Side Request Forgery (SSRF) attack. SSRF attacks allow a malicious actor to force a server-side application to make requests on their behalf. In this case, the attacker made use of the SSRF vulnerability in the consumer-managed environment to make a request for the AWS-managed Instance Meta-Data Service (IMDS). The attacker was able to use IMDS to retrieve the AWS access credentials of the virtual server hosting the vulnerable application. Unfortunately, these access credentials allowed a degree of lateral movement, including access to a variety of S3 buckets hosting sensitive data. This incident is a great example of an attack path including elements of both consumer and provider technology. AWS subsequently released IMDSv2 to prevent this particular attack path from succeeding again in the future. Organizations need to be aware of the accessible attack surface in their cloud provider management infrastructure in addition to more traditional concerns relating to application, operating system, and network.

How am I to secure the Cloud?

When it comes to securing the cloud, knowing is half the battle. Let's examine the key components our security teams need to focus on when applying their security measures in the cloud.

Resources

No matter how you use the cloud, or multi-cloud providers, there is one thing in common - all resources are valuable (and chargeable, exploited systems can prove expensive if used for crypto mining!). When using one of the big 3 cloud providers AWS, Google Cloud, or Azure, there are different resources to protect that each have their own configurations and best practices. Types of resources include databases, DynamoDB, EC2, Beanstalk, ElastiCache, RDS Cluster, RDS Instance, Redshift, S3 Buckets. Virtual Machines & Containers. In enterprise networks it could be Linux, Windows or MacOS hosts. All of these different types of resources can offer potential routes for an attacker.

Misconfigurations & Identities

Misconfigurations & Identities are the Achilles heel of the Cloud. Each provider has differing, but usually complex, configurations to grant access and authorization to services and resources. Determining the exact roles and appropriate level of permissions for each user can be difficult and time consuming. The example compromise of a financial services giant described earlier demonstrates the perils of allowing unnecessary access to a set of cloud access credentials.

73%

Of incidents involve cloud critical assets**

**Source: 2021 Verizon Data Breach Investigations Report

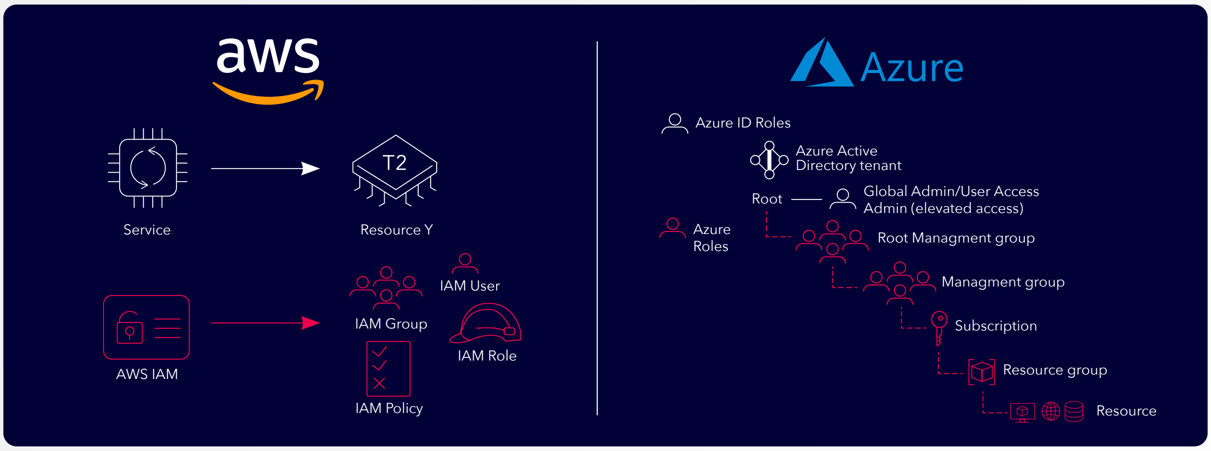

An example of high-level complexities and differences in Amazon AWS and Microsoft Azure visualized

Permissions

Permissions within cloud providers are granular and complex, often inhibiting least privileges approaches when developers or operators may (unfortunately) seek shortcuts – or perhaps be subject to time constraints that make adopting such good practices difficult. For example, when creating a custom policy with permissions on an EC2 instance, it may be time consuming to allow permission for each user that potentially needs access, but much quicker to just allow permission for a whole group. The trick is to review all permissions frequently, and automatically where possible – cloud providers now provide tooling to enable consumers to identify the permissions they are allocating and the resource access that those permissions allow.

The Cloud Security Alliance offers a vast array of security guidance and research to help organizations make use of cloud services in a risk-managed manner. The core guidance document offered by the CSA ("*Security Guidance for Critical Areas of Focus in Cloud Computing v4.0*") offers advice and guidance across a wide range of areas, whilst also informing the Cloud Controls Matrix (CCM) used by many organizations across the globe as a compliance baseline. CSA research is often produced as output from working groups made up of industry volunteers and corporate member representatives, with current *working groups* considering a wide range of issues such as zero trust, securing containerized workloads, the Internet of Things and quantum-safe security.

An ever-expanding attack surface

Businesses do not always clearly define a strategy in their migration to the hybrid cloud world. Organizational decisions can be made to allow individual units to adopt their own migration strategies but sometimes business units make their own arbitrary decisions to source cloud resources without input from IT. Sometimes the lack of a single strategy is down to wider business events, like if an organization with one cloud vendor acquires or merges with another organization using a different cloud vendor. This unplanned large scale of cloud environment complexity and attack surface volume affects security across the entire enterprise IT resources, including:

Your assets: In order to effectively secure your assets you need to know where they are – these assets could be virtual servers, they could be data, they could be functions or any other category of technology asset that supports the enterprise. The recent issues with Log4J have demonstrated the importance of understanding your assets in order to be able to prioritize patching them. The other aspect that the Log4J vulnerability has highlighted is the importance of Software Bill of Materials (SBOMs) in order to understand the components used to build an application service. Adoption of devsecops approaches and tooling can help with both SBOMs and asset inventories if built appropriately. A failure to understand where your assets are located can lead to significant regulatory concerns, particularly if you are unaware of where you may be storing or processing personal data within the scope of regimes such as the EU General Data Protection Regulation.

Your network security: The integrity and security of your data depends on appropriately robust network technologies. The number one priority in your network is connectivity with maximum uptime and zero downtime, across all environments. This is even more crucial if you're using a multi-cloud environment.

The security of your platform: As your infrastructure expands and new innovations are adopted, so too, does your attack surface. The number one priority will be visibility to enable properly configured security policies and effective remediation of risks. Adoption of newer application architectures such as serverless also require a change in security mindset and patterns as it becomes infeasible to rely on traditional agent-based security tooling in the context of ephemeral functions.

Application security: Instead of being limited to which cloud providers can be safely integrated with your critical applications, enhanced application layer security controls allow you to choose and use multiple cloud platforms based on matching performance with your priorities.

The Most Common Attack Techniques

The hacker's goal to succeed is simple. First it is the initial foothold, do heavy reconnaissance by propagating throughout the network, then selecting their target(s) and breaching, either directly or indirectly.

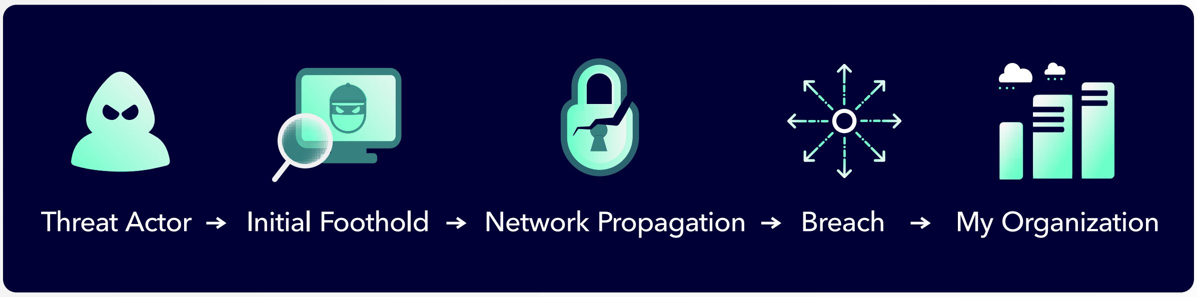

Threat Actor → Initial Foothold → Network Propagation → Breach → My Organization

The initial foothold

Attacker Perspective: Gain/expand access, easy to accomplish, a simple phishing email is sometimes all it takes
Defender Perspective: difficult to control, tough to prevent

The breach

Attacker Perspective: easy to execute on target, get more mileage using advance techniques like double extortion, wreak havoc in the network compromising the organization
Defender Perspective: too late, damage is done. Pay the consequences

The network propagation

Attacker Perspective: Lateral movement, Credential harvesting, Vulnerability exploitation.
Defender Perspective: Longest and slowest phase for the attacker. This allows an organization to take advantage and get the highest return on investment. Focus your security efforts during the network propagation phase to disrupt attacks in the making before they launch.

One of the more commonly used sources of information relating to attack techniques is the MITRE ATT&CK®framework. The MITRE ATT&CK®framework was recently updated to include some attack techniques specific to cloud environments:

https://attack.mitre.org/matrices/enterprise/cloud/
https://attack.mitre.org/matrices/enterprise/cloud/iaas/

The XM Cyber Attack Path Management platform is aligned to the MITRE ATT&CK®framework. Let's look at a few techniques that XM Cyber finds to be most commonly used by attackers and the steps the hacker needs to take below.

On-premise to the Cloud and back again
Attack paths can become very complex in hybrid network architectures. A compromised on-premise desktop offers a low sophistication attack path to compromise Active Directory via Azure. The XM Finding: Intune Administrator user was able to compromise Active Directory

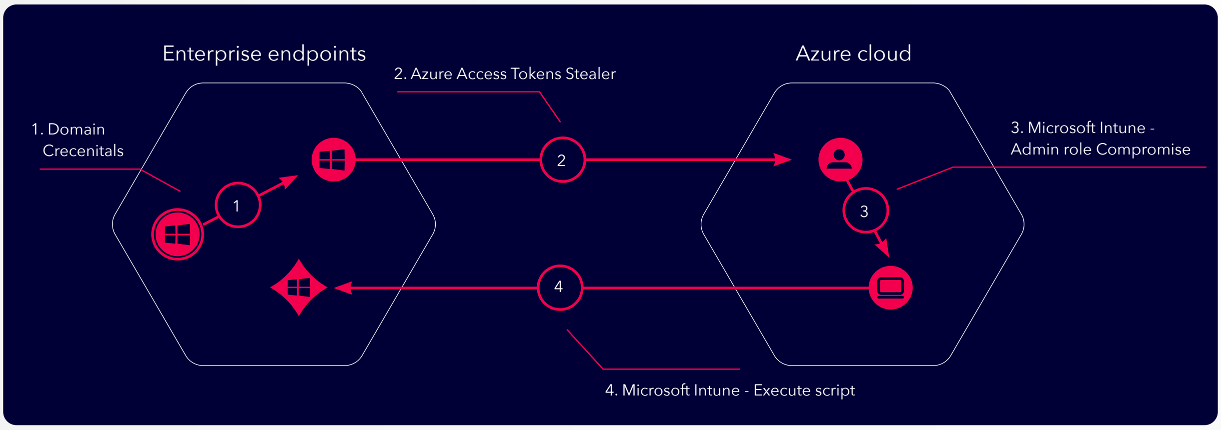
Enterprise endpoints
Azure cloud
1. Domain Crecenitals
2. Azure Access Tokens Stealer
3. Microsoft Intune - Admin role Compromise
4. Microsoft Intune - Execute script
1
2
3
4

Step 1: the initial breach point is via compromise of a Windows machine

Step 2: the attacker steals domain credentials from the breach point

Step 3: the hacker takes the access token from the compromised endpoint and uses it to authenticate to the Azure tenant

Step 4: the compromised access token has Intune privileges and allows attacker to execute commands back on the on-prem critical asset machine(s)

Risk from External Identities and 3rd Party applications to the Production Environment
The Federation of identities greatly increases the attack surface and widens risk.

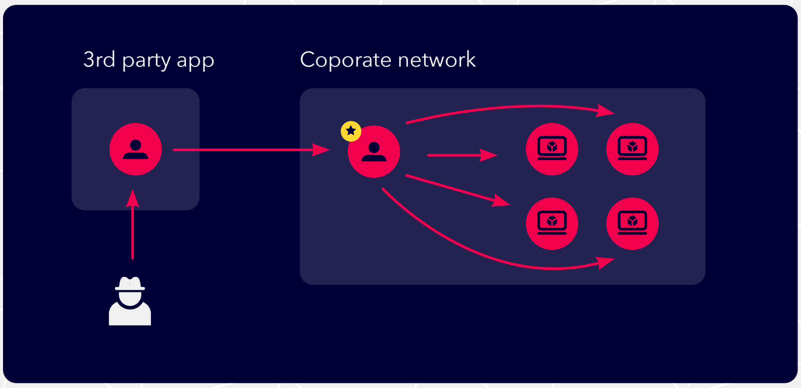

Step 1: In the first example, an attacker compromises the 3rd party application and is able to re-use the privileges of the 3rd party app to access services within the targeted organization

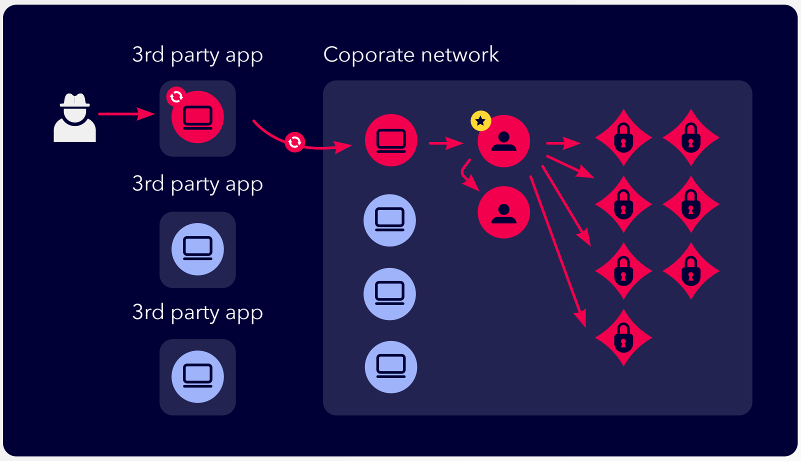

Step 2: In the second example, an attacker is able to compromise the trusted identities managed by the 3rd party app and is able to misuse these federated identities to gain access to services and data within the corporate network

Escalating and Leveraging Privileges within the Cloud

Escalating privileges in the cloud can be a straightforward task for threat actors. Just one developer role can compromise all machines in an AWS environment depending upon how that role is scoped and secured.

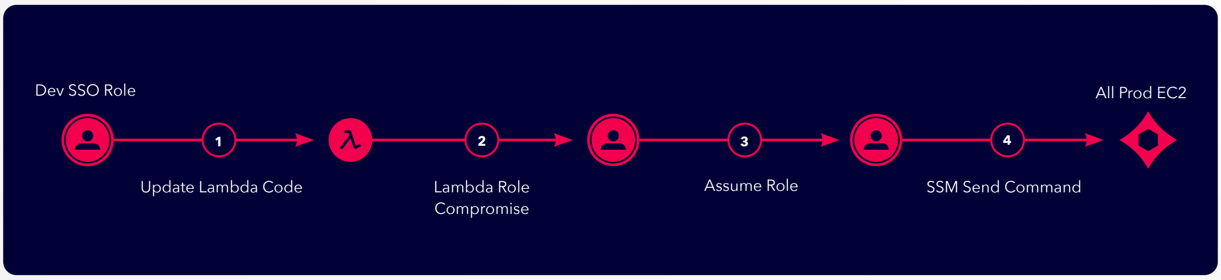

Step 1: an SSO developer role is overly scoped and under protected, allowing an attacker to assume the role and update the lambda code – SSO developer role is accessible from everyone in the environment

Step 2: The lambda function has a specific role attached to it which allows the attacker to generate access tokens of the role and move laterally throughout the environment

Step 3: The same role has permissions to assume another privileged role inside the same environment

Step 4: The privileged role can execute commands on all the EC2 production instances

Best Practices to Improve Your Cloud Security Posture

Prioritize Risks in the Cloud and Reduce Your Attack Surface

The most influential change that will help improve your hybrid cloud security posture is using the attacker's perspective to see the attack before it happens by mapping all possible attack paths and getting a clear view of the security posture across your hybrid cloud ecosystem. Focusing in on the key intersections where multiple attack paths converge to exploit a critical asset, offers more actionable intelligence then receiving a simple vulnerability alert about a single component with a high CVSS score assessed using the Common Vulnerability Scoring System (CVSS). Without the insights of attack paths threat actors take, and how they can compromise your critical assets, it's difficult to retain a high security posture and keep an upper hand against your adversaries.

The five best security practices of Hybrid Cloud users

1
The need for governance, strategy and architecture

2
Understanding of your environment and the critical assets in them

3
Continuous vulnerability monitoring both for threat intelligence and vulnerability assessment with prioritization of high volume of vulnerabilities

4
Continuous and safe attack path modeling to discover high impactattack paths to critical assets

5
Cost-effective remediation with low effort, targeting the high-value elements of an environment that would be key to an attacker's successful compromise

Tradition vs Innovation

There are 3 principles that proactive hybrid cloud security relies on: visibility, analysis, and action. Prior to the emergence of attack path management tooling, there was not an efficient way to identify and then break the critical points in the attack chain. In order to do that you need a clear view into the entirety of your environment from the eyes of an attacker. It is not enough that you are just monitoring the threats and alerts; it's about understanding the context of these vulnerabilities within your environment and the attack paths that these vulnerabilities offer to an attacker looking to breach your critical assets. This is achieved through a deep analysis of the environment and only then can we define the steps needed to eradicate, or at least mitigate, the risk to our organizations.

	WITHOUT Attack Path Management	WITH Attack Path Management
Data Application	Manually define Environment (VM, devices, DB, critical assets, etc)	Easily defined and mapped with fast deployment (VM, devices, DB, critical assets, etc)
Operation System Servers	Constantly and manually monitor environment according to least-priviliges, without any knowledge of best practices	Continuosly and automatically monitor environment against latest attack techniques (aligned with MITRE)
Storage Physical	First- hand knoweledge of risk remediation	Prioritized and guided remediation steps for highest risks.

From the perspective of the UK Chapter of the Cloud Security Alliance, we see attack path management as a useful addition to the armoury of cloud security professionals. Hybrid cloud environments can be complex, made up of a wider variety of SaaS, PaaS, IaaS and FaaS services running code developed either in-house, by software vendors or by managed services providers. This complexity is reflected in novel attack surfaces at a variety of boundaries including the traditional Internet boundary as well as boundaries more specific to cloud services such as the boundary between tenants in multi-tenant solutions and between the cloud consumer and cloud provider in any cloud solution. The compromise at the financial services organization discussed earlier in this paper is a great example of an attack path traversing the responsibility boundary between consumer and provider with elements of both consumer and provider services being misused to allow lateral movement across the consumer's cloud environment. Identifying potential attack paths allows defenders to identify potential choke points or other critical elements of a potential route to compromise where efforts can be made to either remove vulnerabilities or else ensure that suitable compensating controls are implemented.

Examples of the research and guidance that is made available by the Cloud Security Alliance were provided earlier. The full set of artifacts that the CSA has published can be found here:

https://cloudsecurityalliance.org/research/artifacts/

Outside of the CSA, a variety of standards and security agencies across the globe have published guidance and compliance baselines relating to cloud security including NIST and the International Organization for Standardization (ISO).

Summary

Where should you invest your security effort to disrupt threat actors' attack paths that may compromise critical assets?

The pandemic has accelerated moves towards zero trust and hybrid cloud. This has implications for enterprise security architectures in areas from remote access, device security and identity through to micro segmentation across hybrid environments. Whatever the architecture or ecosystem, attackers will still be out there looking to make a quick buck, whether through ransomware, crypto mining, or perhaps old-fashioned extortion.

No matter how you run your organization you can leverage attack path insights to get a clear view into your hybrid cloud security posture. Focusing security efforts during the network propagation phase to disrupt attacks in the making before they launch will yield the greatest return on investments. By combining how attackers can exploit security gaps like misconfigurations and vulnerabilities in relation to your critical assets you can disrupt the opportunity for lateral movement across the network and pinpoint the exact changes needed to quickly eliminate the risk of compromise.

To learn more about attack path management, visit: **https://www.xmcyber.com**

About the UK Chapter of the Cloud Security Alliance

The Cloud Security Alliance is a global not-for-profit organisation dedicated to defining and raising awareness of best practices to help ensure a secure cloud computing environment. The UK Chapter of the Cloud Security Alliance is a voluntary organisation, established in August 2012. The purpose of the chapter is to take the best of the Global Cloud Security Alliance guidance and make it actionable for a UK audience, encouraging the growth of local cloud security talent. We are a vendor-agnostic organisation, but we do work with vendors to produce content of interest to our members.

To learn more about the UK Chapter of the Cloud Security Alliance,
visit: **https://www.cloudsecurityalliance.org.uk/**

About XM Cyber

XM Cyber is a leading hybrid cloud security company that's changing the way innovative organizations approach cyber risk. Our attack path management platform continuously uncovers hidden attack paths to your critical assets across cloud and on-prem environments, so you can cut them off at key junctures and eradicate risk with a fraction of the effort. This approach is a complete game-changer, which is why some of the world's largest, most complex organizations choose XM Cyber to help eradicate risk. Founded by top executives from the Israeli cyber intelligence community, XM Cyber has offices in North America, Europe, and Israel.

Tel-Aviv:	+972-3-978-6668
New-York:	+1-866-598-6170
London:	+44-203-322-3031
Munich:	+49-163-6288041
Paris:	+33-1-70-61-32-76

xmcyber.com

XM Cyber