CYBER
THREAT
ANALYSIS

Recorded Future®

By Insikt Group®

August 4, 2025

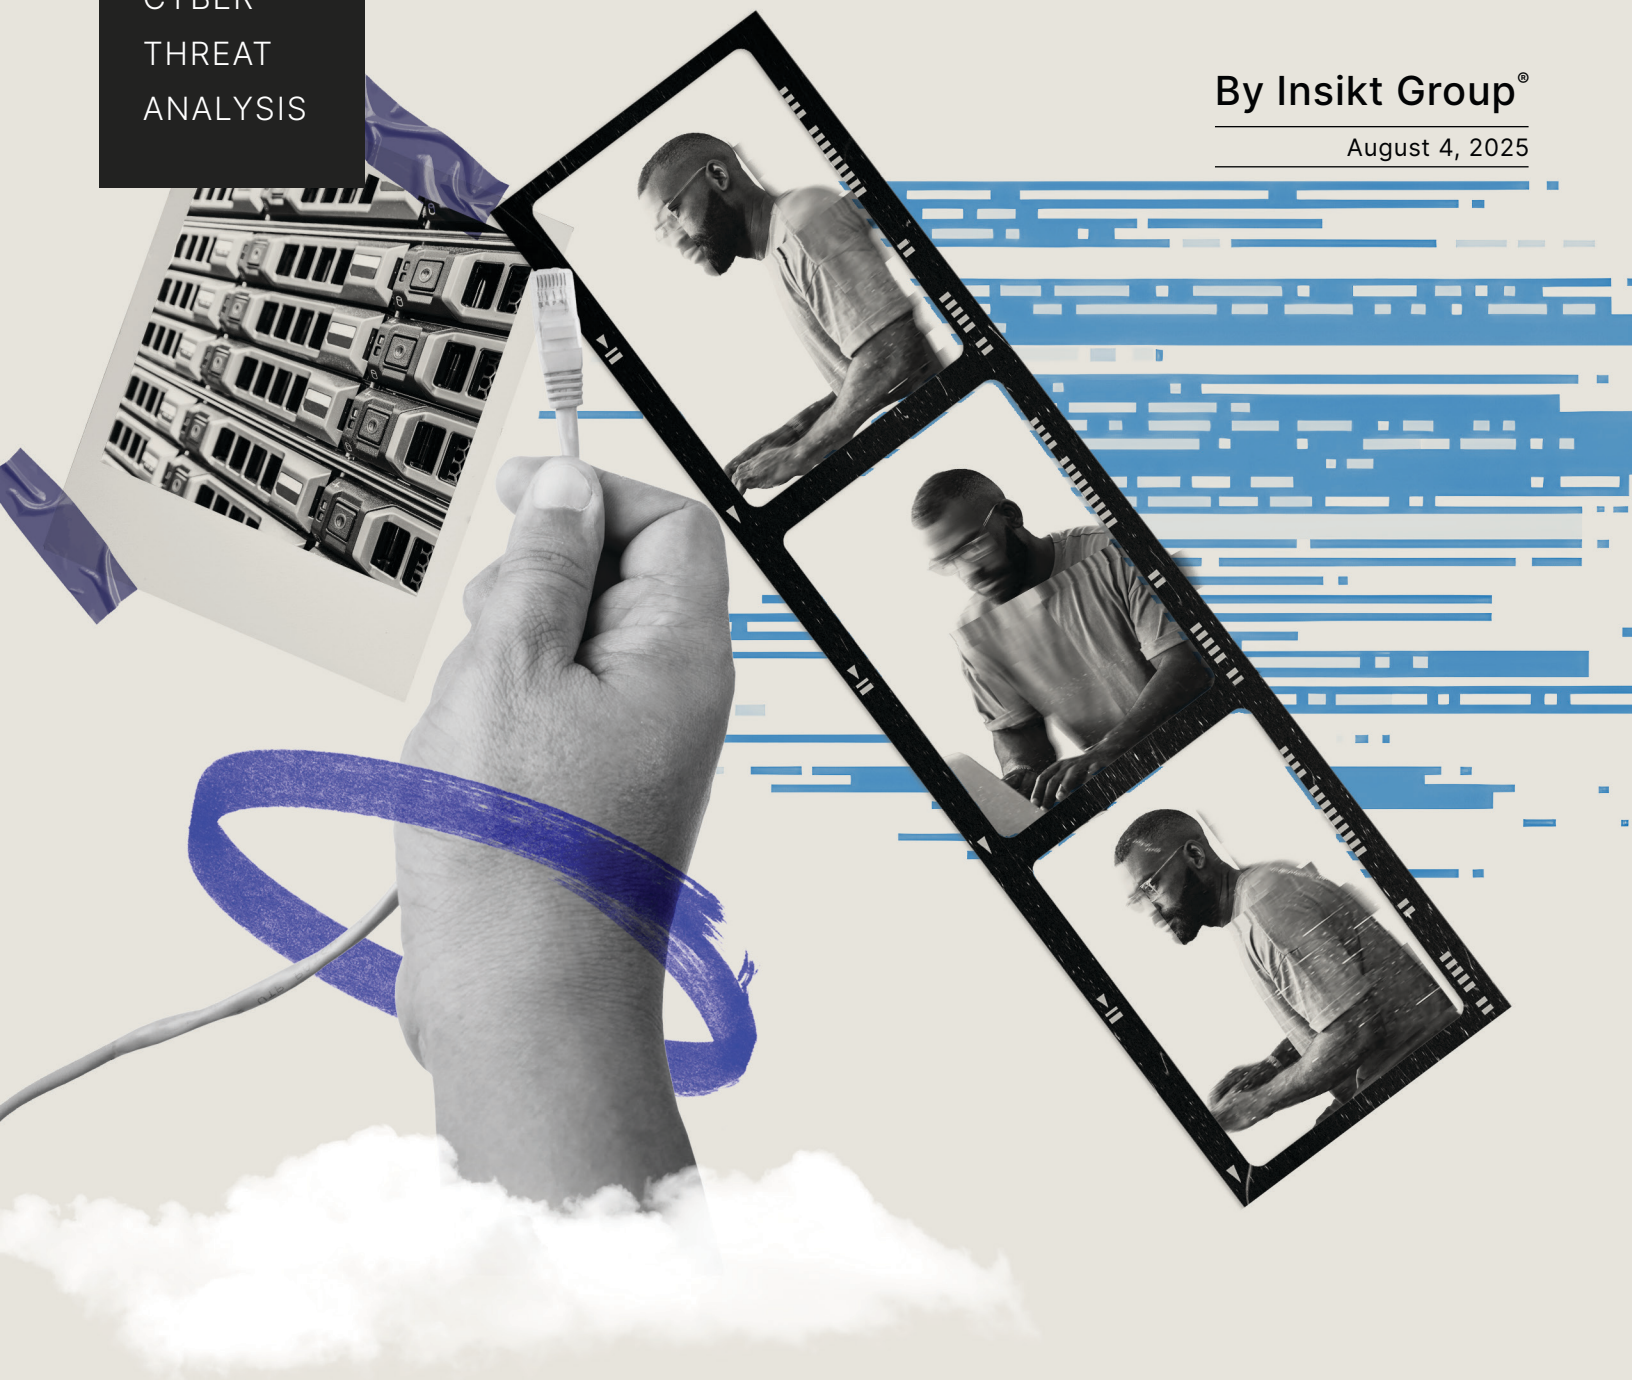

Cloud Threat Hunting and Defense Landscape

As organizations increasingly adopt cloud infrastructure, they encounter novel and unique security challenges that threat actors are actively exploiting

Threat actors targeting cloud environments rely mainly on exploiting misconfigurations and employing coercion tactics for initial access.

Vulnerability and misconfiguration scanning campaigns, alongside initial access brokers, represent the primary means by which threat actors obtain cloud credentials.

Executive Summary

In a review of recently observed attack methods, Insikt Group identified five attack vectors that currently pose the greatest potential threat to cloud environments. Three of these attack methods, vulnerability exploitation, endpoint misconfiguration, and credential abuse leading to account takeover, can grant threat actors initial access. In certain circumstances, these three attack methods can also be employed following initial access to gain increased permissions within a cloud environment, modify the cloud environment, and allow lateral movement, either to additional cloud environments, traditional on-premise environments, or user devices. The two remaining attack methods, cloud abuse and cloud ransomware, demonstrate impact actions threat actors can perform within a cloud environment.

Hunting for each of these threats often requires the implementation of robust logging within cloud environments to ensure that data such as network communications, user access, and cloud service usage metrics can be readily accessed and scrutinized for aberrations. Log data assists in both proactive discovery of suspicious activity originating at the edge of cloud environments, such as in instances where misconfiguration and vulnerability scanning occur, and in identifying instances where cloud accounts and resources are abused for malicious purposes.

To mitigate threats from impacting cloud environments, proper configuration of the environment is paramount, both at the edge of the cloud environment, including the methods by which users and services interact with the environment, and within the environment itself. Cloud environments that are configured appropriately minimize the risk of initial access and can significantly limit the malicious actions a threat actor is capable of performing post-initial access. Additionally, the most common cloud platforms provide native services focused on security for cloud environments, such as web application firewalls (WAF), identity and access management (IAM) services, secrets storage and management suites, and secure data connectors for hybridized cloud environments, that allow cloud architects to mitigate the threats discussed in this report with relative ease.

Key Findings

- Most initial compromises start with exposed or misconfigured cloud endpoints, with attackers using open-source scanners to identify misconfigured endpoints.
- Stolen or weak credentials, often gathered from initial access brokers (IABs) and previous malicious actions performed by the attacker, remain the fastest path to full-tenant cloud takeover.
- Threat actors increasingly abuse legitimate SaaS and IaaS resources, shifting costs to the owners of victimized environments and abusing resources to complicate the detection of follow-on malicious actions, such as phishing campaigns.
- Ransomware groups have adopted cloud-native tactics, encrypting S3 and Azure storage directly and disabling backups to maximize leverage.
- Hybrid infrastructure lets attackers pivot seamlessly between on-premise and multi-cloud environments, so visibility and controls must extend beyond the cloud environment to the devices and services that access it.

Introduction

During the past decade, a steady shift from traditional on-premise IT infrastructure to cloud-based infrastructure and hybrid cloud infrastructure has taken place. According to PwC's 2023 Cloud Business Survey, 39% of private respondents stated that the entirety of their operations had been moved to cloud environments. Cloud computing has become a trusted and integral part of many corporations' day-to-day operations. Since the time of PwC's reporting, cloud computing as an industry has only grown with no signs of slowing.

The breadth of cloud products and the depth of services provided by cloud environments continue to grow daily. In a joint study conducted by Amazon and Telecom Advisory Services, cloud adoption accounted for a total of $1 trillion in the global gross domestic product, with a projected increase to $12 trillion between 2024 and 2030. This estimate indicates that traditional computing environments will continue to migrate to cloud environments rapidly in the coming years. That demand for cloud computing resources will continue to increase for the foreseeable future.

The success of cloud computing can be squarely attributed to the benefits that adopters are provided. When properly configured, cloud environments allow their adopters to shift costs associated with traditional on-premise environments, create high-availability to remote assets, and eliminate development overhead by gaining access to managed services. As cloud providers continue to offer additional services and products that make similar offerings for traditional environments less effective from cost and operational perspectives, cloud adoption will only continue to grow in the future.

Background

Cloud technologies, platforms, and services are increasingly implemented into corporate structures, providing all of the benefits of traditional on-premise environments while reducing costs associated with an on-premise environment in nearly every conceivable way. This relationship was demonstrated in PwC's "2024 Cloud and AI Business Survey," which reported that, out of a survey of 1,000 companies that implemented cloud technologies, 74% of the surveyed companies that have optimized their cloud environments reported increased profitability, and 65% of the same respondents reported increased cost savings. While these benefits are highly appealing to corporations, cloud environments pose unique risks and security challenges, challenges that require a fresh approach to cybersecurity to mitigate properly.

The advancement of cloud environments has also increased the number of network-accessible endpoints that an organization must monitor and defend. In instances where large enterprise entities have fully migrated their operations to cloud environments, the endpoints required to facilitate user access, deploy web applications, support data transfer, and provide many other kinds of access on a day-to-day basis add up quickly and create a diverse boundary that is constantly interacting with the broader internet. The technologies that interface with and are embedded within this boundary pose unique risks and security challenges. Looking inward, similar issues persist, with cloud defenders requiring a fresh understanding of how cloud environments can be effectively architected to provide the benefits of a cloud environment without allowing undue access to sensitive information and control over mission-critical assets hosted in these environments.

As Insikt Group discusses in this report, threat actors have become increasingly aware of the security challenges cloud defenders must address, as well as the opportunities that cloud technologies, environments, and services afford them. The overwhelming amount of data, applications, systems, and other assets hosted on cloud environments, coupled with the task of defending these assets, provides threat actors with novel opportunities to compromise information, abuse environment resources, and profit from illicit activities in ways previously unattainable in on-premise environments. Additionally, threat actors have begun to understand the usefulness of cloud resources as part of an attack chain, realizing they are afforded all of the same benefits of legitimate cloud users, with the added benefits of anonymity and reduced detection capabilities in a way that is unobtainable with traditional infrastructure.

Understanding the threat posed by these adversaries, this report was created to shed light on the most impactful and emerging tactics, techniques, and procedures (TTPs) displayed by threat actors that target and abuse cloud environments. In doing so, it aims to provide an understanding of how threat actors are impacting and abusing cloud environments at a granular level, as well as how to mitigate these threats and hunt for indicators of compromise associated with them so that cloud defenders are better able to identify and respond when necessary.

Methodology

This report identified five main threats to cloud environments, each of which are explored their respective sections:

- Cloud Abuse
- Exploitation
- Endpoint Misconfiguration
- Cloud Ransomware
- Credential Abuse and Account Takeover

Each section includes radar charts that measure the following attributes associated with a given threat. These determinations were derived by Insikt Group investigating instances where this threat vector was observed to answer the following questions:

- **Cost of Impact:** How much would this threat cost a victim in terms of monetary, reputational, and operational losses? In the radar chart, the higher the number, the higher the cost the victim can expect to incur monetarily, reputationally, operationally, or otherwise.
- **Commonality:** How often is this threat vector observed in attack chains against cloud environments in the wild? In the radar chart, the higher the number, the more likely a cloud defender is to observe this behavior in their own environment.
- **Evolution Potential**: What is the potential for threat actors to further "evolve" this attack vector in terms of new tools, attack methods, and TTPs that can be employed to achieve this threat vector? In the radar chart, the higher the number, the more likely it is threat actors will be able to perform actions demonstrating this threat in ways previously unobserved, thus complicating detection of the behavior.
- **Effort to Perform:** What are the technical and monetary costs associated with performing this threat vector? In the radar chart, the higher the number, the greater the barrier for an attacker to demonstrate this threat against a cloud environment, generally in terms of monetary cost or technical capability.

Threats To Cloud Environments

Cloud Abuse

Key Takeaways

- Attackers registered their own cloud infrastructure to host malicious content and exfiltrate stolen data to their own cloud environments.
- Uses for compromised cloud environments varied heavily and were determined by the responsible threat actor's goal or proficiency.

Figure 1 illustrates and compares attributes associated with cloud abuse. A description of each attribute can be found in the **Methodology** section of this report.

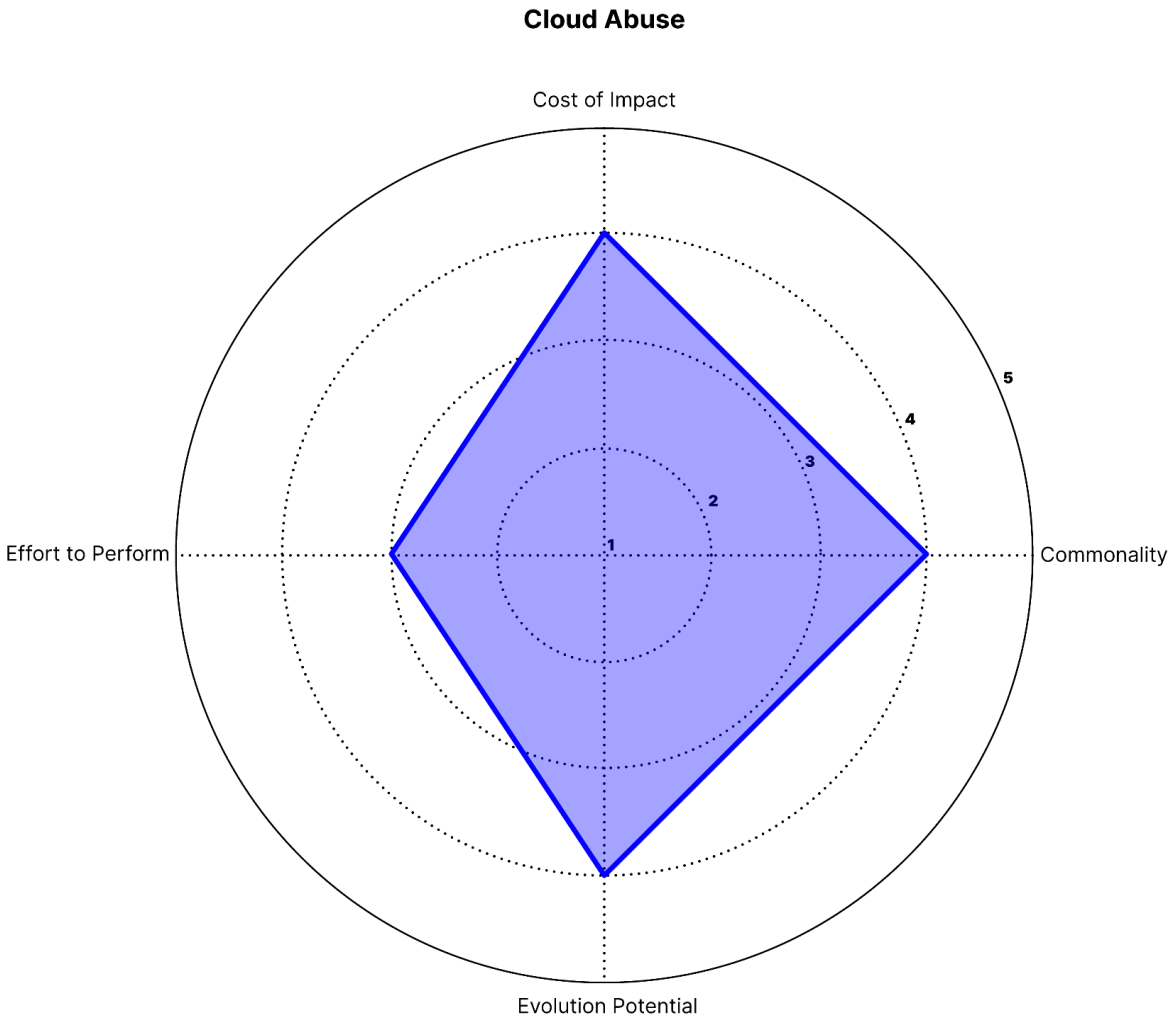

Figure 1: *Radar chart illustrating cloud abuse as a threat vector (Source: Recorded Future)*

Cost of Impact: 4 (High)

Attacks where threat actors abuse victim cloud environments are highly costly, whereas instances where threat actors register and abuse legitimate services are comparatively less costly. In both instances, threat actors are able to masquerade as legitimate entities, leading to reputational losses for the abused environment and owner. Instances where threat actors abuse compromised victim cloud infrastructure often result in increased costs to the owner of the cloud environment.

Commonality: 4 (High)

Abuse of legitimate cloud infrastructure registered by a threat actor is very common, whereas abuse of compromised victim cloud infrastructure is comparatively less common. Many observed attacks against cloud infrastructure include threat actors attempting to gain control of cloud services for follow-on actions at some point, indicating that this type of threat remains common with respect to other cloud threats.

Evolution Potential: 4 (High)

Threat actors have demonstrated that there are a plethora of ways cloud abuse can be achieved and then leveraged to perform malicious actions within the past year. Additionally, novel techniques such as "LLMjacking," where threat actors sell access to compromised, cloud-based LLM models, indicate that threat actors are continuously considering how to monetize the abuse of cloud services, forecasting an increase of cloud service abuse in the future.

Effort to Perform: 3 (Moderate)

Both the abuse of legitimately registered cloud infrastructure and compromised victim cloud infrastructure pose moderate difficulties to threat actors. In the former threat type, attackers must determine how to register for larger cloud platforms anonymously and conduct malicious actions without being detected, all while paying for the environment. In the latter threat type, threat actors are only able to abuse victim cloud infrastructure after adequately compromising cloud services and systems that are necessary for them to achieve their overarching goals.

Threat Summary

The term cloud abuse refers to two overarching behaviors threat actors have displayed when targeting cloud environments:

- Abuse of legitimate cloud infrastructure obtained by a threat actor to perform malicious activities
- Abuse of legitimate cloud infrastructure owned by a victim a threat actor compromises to perform malicious activities

In both instances, threat actors abuse legitimate cloud infrastructure for nefarious purposes; however, the behaviors demonstrated by threat actors in each of these scenarios differ significantly. In the former example, threat actors will mainly abuse these resources to appear as part of legitimate traffic and

remain anonymous; this behavior is often used to carry out phishing campaigns, host malicious content, and act as part of the threat actor's command-and-control (C2) infrastructure. In the latter example, threat actors may still abuse the cloud environment to masquerade as a legitimate entity, but they may also hijack the environment's resources, shifting costs to the environment's owner. In such an instance, additional actions such as cryptojacking and a more recent technique, LLMjacking, may occur and result in inflated monetary costs.

Outlook

Threat actors will almost certainly continue obtaining their own cloud infrastructure for several reasons:

- Threat actors are afforded all of the same benefits legitimate cloud users are provided, in addition to anonymizing factors that aid in malicious actions (see **Figure 2**).
- It is relatively easy to obtain cloud infrastructure from CSPs without extensive scrutiny from the provider, allowing attackers to create extensive cloud environments without suspicion.
- Abuse of legitimately registered cloud environments is often identified reactively following malicious actions originating from the environment, indicating that CSPs do not have a reliable method of detecting cloud abuse prior to victim compromise.
- Threat actors are easily able to pivot from one cloud provider to another and are able to mask their identities while performing malicious actions when abusing cloud resources.

Figure 2: *Benefits of cloud adoption, legitimate user versus threat actor (Source: Insikt Group)*

Threat actors will also continue abusing compromised cloud infrastructure because:

- Compromised cloud infrastructure opens the possibility of third-party risk, possibly allowing attackers to masquerade as the victim when targeting and compromising additional targets.
- Attackers can abuse compromised cloud environments to shift the cost of otherwise costly operations that cloud environments can provide onto the victim, such as cryptomining and requests to cloud LLM services.

Mitigations and Detections

Figure 3 demonstrates a hypothetical attack chain where both abuse of a legitimate cloud account and abuse of a victim cloud account occur. Throughout this visual, Insikt Group has identified parts of the

attack chain where defenders can most efficiently hunt for and mitigate behaviors associated with cloud abuse.

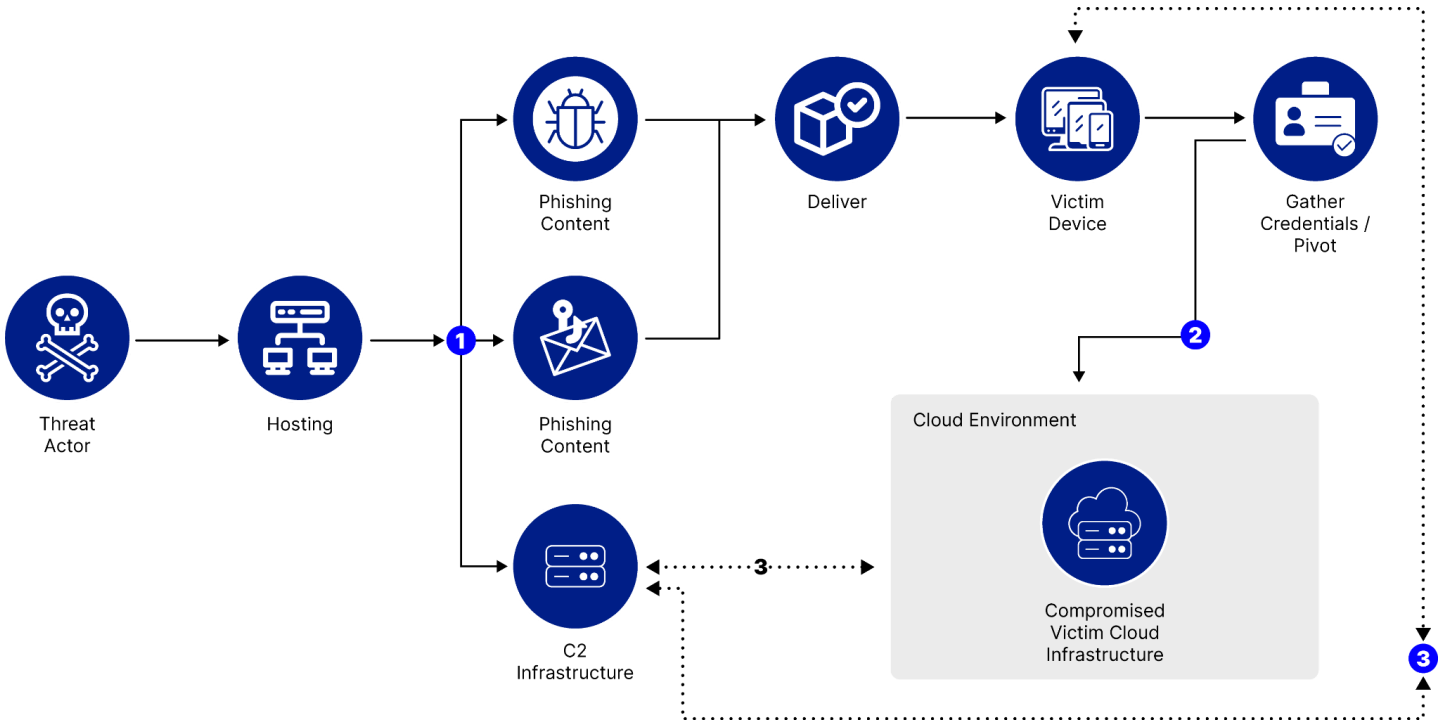

Figure 3: *Visual representation of potential cloud abuse attack vectors (Source: Recorded Future)*

① Threat Actor-Controlled Cloud Infrastructure

Threat actors will host malicious content and infrastructure using cloud platforms. This activity can be mitigated using the following techniques:

- Monitor for phishing emails that could redirect to the malicious content hosted on cloud architecture
- Flag user browsing activity; threat actors will often implement phishing links that redirect the victim to malicious phishing webpages
- Use threat intelligence to assist in identifying malicious websites and other infrastructure that a user may unknowingly access during their usual browsing

② Pivoting from User or On-Premise Hosts to Cloud Infrastructure

Attackers have been observed compromising traditional IT infrastructure to gain access to cloud infrastructure. This activity can be hunted using the following techniques:

- Search cloud access logs for failed authentication requests originating from an external source that ultimately succeed in authenticating.

- Search cloud API logs for aberrant behavior, such as an account attempting to make API calls with improper privileges or an account making a high number of calls within a short period of time.
- Search cloud access logs for aberrant login behaviors, such as authentication attempts at uncommon times of the day or login attempts from unknown or suspicious IP addresses or geographic locations.

Notably, Insikt Group observed multiple instances (discussed in detail later in this section) where threat actors implemented man-in-the-browser (MitB) or token replay attacks where threat actors captured legitimate authentication materials, session tokens, and multi-factor authentication (MFA) codes. These methods allowed attackers to bypass common authentication standards, so additional mitigation efforts should be implemented to discourage and limit the impact of unauthorized cloud access.

Additional mitigations include the following:

- Rotate cloud access credentials upon suspecting that a user's workstation has been compromised.
- Implement additional authentication measures aside from MFA, such as passwordless options, like passkeys, and physical authentication technologies, like physical security keys, to grant access to the cloud environment.
- Require regular reauthentication to the cloud environment, and ensure that access tokens are regenerated on a regular basis.
- Implement the principle of least privilege to minimize possible attacker actions within a cloud environment following a compromise.
- Create user baselines to identify aberrant behavior stemming from a compromised user account following a compromise.

③ Cloud-Hosted C2 Infrastructure

Threat actors have been observed abusing cloud infrastructure and services to provide functionality similar to traditional C2 infrastructure, most often as an exfiltration endpoint. This activity can be hunted in the following ways:

- Monitor logs associated with cloud services that provide file storage and transfer, such as AWS S3, and Azure Storage, as well as Google Cloud Storage for unexpected or high-volume, outbound data transfers.
- Monitor cloud network endpoints for high-volume outbound communications or a pattern of inbound/outbound communications to an unrecognized source that follow a time-based pattern.
- Hunt for newly created artifacts and applications within your own cloud environment that either provide a point of egress or that are configured with overly permissive roles.

Additionally, this behavior can be mitigated in the following ways:

- Ensure technologies such as access control lists (ACL), web application firewalls (WAF), cloud firewalls, and other cloud-based network controls are implemented and configured to ensure that both inbound and outbound access is only allowed for known and necessary functions.
- Restrict the creation of new applications and assets within a cloud environment by implementing identity and access management (IAM) policies and roles for both user and service accounts.

Examples In the Wild

Insikt Group curated a list of events published within the past year that demonstrate the threats posed by cloud abuse. These events are discussed below.

Docker and Kubernetes Instances Abused in High-Volume Cryptojacking Campaign

On September 23, 2024, Datadog Security Labs published an article discussing a cryptojacking campaign exploiting Docker and Kubernetes environments to mine cryptocurrency at scale. Threat actors gained initial access by scanning for exposed Docker API endpoints and deploying a malicious Alpine container that mounts the host filesystem. The attack progressed by abusing Docker Swarm's orchestration features for C2, allowing the adversaries to join compromised hosts into their Swarm cluster.

Attackers targeted Kubernetes clusters and containers through the kubelet API exploitation, enabling the deployment of additional malware and the execution of unauthorized workloads. The attackers then used scanning tools such as `masscan` and `zgrab` to identify additional candidates for lateral movement, including additional Kubernetes containers and Secure Shell (SSH) servers. The campaign used Docker Hub to distribute malicious container images under accounts such as `nmlmweb3`.

The attackers achieved persistence through dynamic linker hijacking, SSH key manipulation, and malware injection into compromised hosts. Additionally, attackers manipulated Docker Swarm clusters by forcing hosts to leave existing configurations and join a threat actor-controlled Swarm, effectively converting infected nodes into a botnet. The attack infrastructure included the `solscan[.]live` domain for payload distribution. Attackers also implemented obfuscation techniques such as process hiding and custom binaries to ensure stealth during the campaign's duration.

Peach Sandstorm: Abuse of Azure Infrastructure for C2 Operations

On August 28, 2024, Microsoft reported that the Iranian state-sponsored threat group Peach Sandstorm, also tracked as APT33, was observed abusing Azure infrastructure to conduct intelligence-gathering campaigns. The campaigns occurred between April and July 2024 and targeted satellite, communications, oil and gas, and government sectors in the United States and United Arab Emirates. As part of this campaign, Peach Sandstorm deployed Tickler, a multi-stage backdoor, and abused legitimate Azure infrastructure for C2 operations. The group also conducted password spray attacks against educational sector organizations to identify valid credentials, which they used to create and manage malicious Azure subscriptions.

Peach Sandstorm abused Azure resources by creating Azure tenants using fraudulently registered Microsoft Outlook accounts and abusing compromised user accounts associated with the education sector. After creating these Azure tenants, Peach Sandstorm established Azure for Students subscriptions within the tenants. These attacker-controlled Azure environments were then used to deploy Azure-based C2 nodes, facilitating the operation of the Tickler backdoor. Deployment of Azure based environments allowed Peach Sandstorm to blend C2 activity with legitimate cloud traffic, thereby minimizing activity detection.

The Azure App Service was the primary service abused for hosting C2 servers. Peach Sandstorm registered a series of malicious subdomains under *azurewebsites[.]net*, which served as C2 endpoints for the Tickler malware. These domains were used to relay commands, receive stolen data, and distribute secondary payloads.

FLUXROOT and PINEAPPLE: Google Cloud Infrastructure Abused to Host Credential Phishing Pages

Google's Threat Horizons H2 2024 report published the discovery that, since 2023, the financially motivated threat groups FLUXROOT and PINEAPPLE exploited Google Cloud infrastructure to conduct credential phishing campaigns.

FLUXROOT used Google Cloud serverless infrastructure, including Cloud Run and Cloud Functions, to host credential harvesting pages targeting users of the Latin America payment platform, Mercado Pago. The group registered Google Cloud container URLs and tested their detection rates on VirusTotal before deployment to refine their stealth tactics. Google later identified and suspended FLUXROOT-associated projects while updating Safe Browsing detections.

PINEAPPLE deployed phishing sites and malicious redirects using Google Cloud domains, such as *cloudfunctions[.]net* and *run[.]app*, embedding these URLs in phishing emails impersonating Brazil's Receita Federal. The group exploited email security weaknesses by bypassing Sender Policy Framework (SPF) checks and manipulating SMTP Return-Path fields to evade detection. When Google blocked PINEAPPLE's serverless infrastructure, the group pivoted to Google Compute Engine (GCE) instances with static public IP addresses, distributing phishing payloads via `.zip`, `.lnk`, and `.html` files. PINEAPPLE later expanded its operations to Azure and Tencent Cloud, demonstrating a multi-cloud abuse strategy to evade Google's countermeasures.

Credentials Stolen in Laravel Exploitation Campaign Used in "LLMjacking" Campaign

On May 6, 2024, the Sysdig Threat Research Team (TRT) published a writeup discussing an attack campaign using stolen cloud credentials to exploit cloud-hosted LLM services, a tactic referred to as "LLMjacking." In this instance, attackers exploited versions of Laravel, an open source web framework, that were vulnerable to CVE-2021-3129, allowing the attackers to exfiltrate cloud credentials and systematically probe multiple AI services, including Anthropic, OpenAI, and AWS Bedrock.

Sysdig alleged the threat actors used an OpenAI (OAI) reverse proxy to facilitate unauthorized LLM access for financial gain while shifting the cost burden to compromised cloud account holders, which

Sysdig estimated could result in costs exceeding $46,000 daily. The attackers employed evasive techniques such as issuing API calls with deliberately invalid parameters to confirm service access without raising alarms. Additionally, they queried cloud logging configurations to determine whether monitoring was enabled, ensuring their activity remained undetected. This campaign highlights the increasing sophistication of cloud credential abuse, particularly in AI, where attackers capitalize on the high operational costs of LLM services for illicit profit.

Nearly 20% of All Public Docker Hub Instances Used to Host Malicious Content

On April 30, 2024, JFrog Security researchers published findings that nearly 20% of public repositories on Docker Hub, totaling nearly three million repositories, were being used to distribute malware and phishing scams. Two of the most significant campaigns, dubbed "Downloader" and "eBook Phishing," leveraged Docker Hub's infrastructure to coerce users into downloading malicious payloads.

The Downloader campaign, which was active in 2021 and resurfaced in 2023, primarily used automatically generated repositories that promoted pirated content, game cheats, and cracked software. These repositories contained links to malicious sites that mimicked legitimate URL shorteners, dynamically redirecting users to malware-hosting domains such as *failhostingpolp[.]ru*. The downloaded files were trojanized executables that connected to C2 servers, collecting system information and installing additional payloads while using obfuscation techniques like XOR encoding and dynamic URL redirection to avoid detection.

The eBook Phishing campaign, active primarily in 2021, targeted users searching for free eBooks by deploying nearly a million repositories filled with SEO-optimized descriptions and links to phishing pages such as *rd[.]lesac[.]ru*. Victims were directed to fake download portals that prompted them to enter credit card details under the pretense of unlocking content, effectively stealing payment information and enrolling them in fraudulent subscription services. Attackers ensured a consistent presence by continuously creating new repositories to replace the ones taken down.
JFrog researchers identified these campaigns by analyzing anomalies in Docker Hub's repository creation patterns and linking them to clusters of automated account activity. Following JFrog's disclosure, Docker removed the malicious repositories, but the abuse highlights the growing trend of attackers exploiting open-source registries to facilitate large-scale cybercrime.

Exploitation

Key Takeaways

- Threat Actors continually leveraged vulnerabilities found in cloud services to exploit victims.
- Credentials obtained by threat actors were consistently used to gain and maintain access to victim cloud environments.

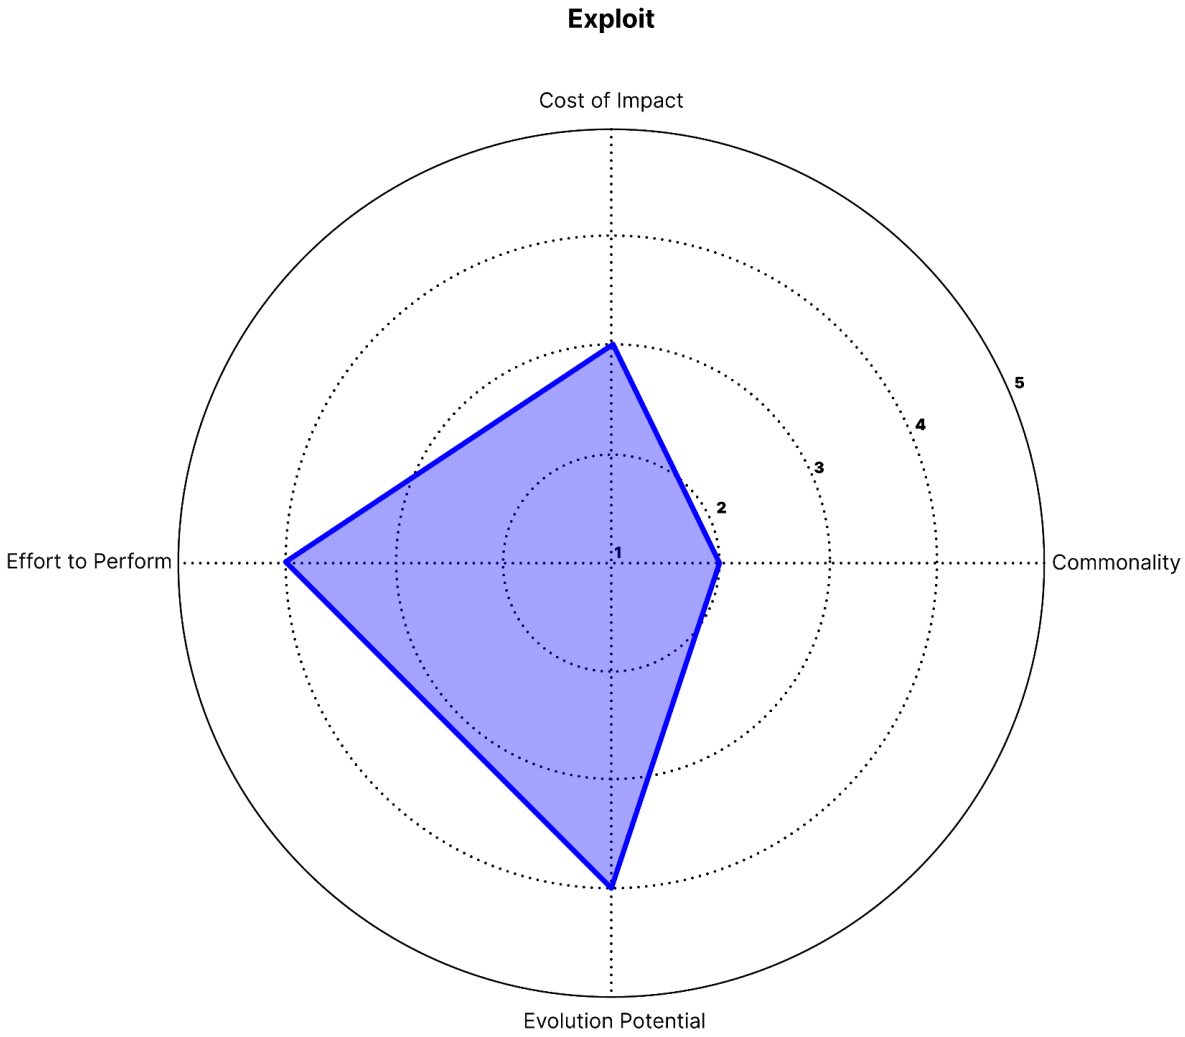

Figure 4: *Radar chart illustrating exploitation as a threat vector (Source: Recorded Future)*

Cost of Impact: 3 (Moderate)

Successful exploitation of cloud infrastructure or the technologies embedded in it may result in multiple benefits to a threat actor, but may not directly translate to victim cost. Based on the vulnerabilities discussed in this report, multiple additional steps or chained exploitation is required in addition to a singular exploit to make an impact on the victim environment. Despite this, vulnerability exploitation is difficult to detect in some instances and may result in threat actors having more time to increase their level of impact within an environment.

Commonality: 2 (Low)

Due to the overhead and managed nature of larger cloud environments, vulnerabilities in the underlying infrastructure are often identified and mitigated before the vulnerabilities are publicly disclosed. This restricted timeline aids in mitigating threat actor exploitation of these vulnerabilities.

Evolution Potential: 4 (High)

The ongoing possibility of undisclosed vulnerabilities means cloud environments will continually face the risk of zero-day vulnerability exploitation.

Effort to Perform: 4 (High)

The discovery of a vulnerability requires a deep technical understanding of the targeted cloud environment, including the infrastructure and services in which a vulnerability may reside, as well as an understanding and background in exploit development. Only sophisticated threat actors are capable of demonstrating this threat.

Threat Summary

Threat actors are increasingly exploiting cloud services and infrastructure, both as a means to conduct malicious activity and as targets for gaining access to enterprise and end-user systems. Despite the growing adoption of cloud technologies, they remain vulnerable to exploitation. As with any evolving technology, new vulnerabilities will inevitably be discovered, some of which are weaponized by adversaries before being patched or even disclosed. Threat actors often depend on chaining multiple vulnerabilities, rather than relying on a single exploit, to breach cloud environments, highlighting the complexity and persistence of modern cloud-targeted attack strategies.

Over the course of the past year, Insikt Group found multiple events where threat actors targeting cloud services used a unique series of exploits to carry out their attack. A sample of these events are detailed in the **Examples in the Wild** section. These incidents underscore that attackers are not merely opportunistic but are tailoring multifaceted exploitation strategies to defeat cloud defenses.

In Insikt Group's research of exploits against cloud technology, several events were found that highlighted threat actors leveraging vulnerabilities against one or more cloud services and infrastructure to exploit a system. In many of the events, Insikt Group observed that it was not a singular

vulnerability or exploit that led to the compromise, but rather a series of exploits tailored by the threat actors to attack each system. The evolving threat landscape demands continuous vigilance and adaptive security practices to detect and respond to multi-stage exploit campaigns in the cloud.

Outlook

Exploitation of cloud environments and cloud services will almost certainly continue and escalate as entities continue to rely on and grow their usage of cloud services. The readily available nature of cloud environments provides ease of use for end users, but the increased attack surface and vulnerabilities from using extra, third-party software and surfaces also increase the number of exploitation opportunities for threat actors for both initial and sustained access.

As was seen in the previous section, attackers can easily access cloud resources and infrastructure to carry out attacks. Similarly, attackers can easily study these same resources for vulnerabilities that will enable their operations against cloud adopters and end users. Cloud providers are not responsible for the security of the deployments, only the infrastructure they are hosted on. As a result, in these easily deployable and standard configurations, secure implementation and configurations of the environments is crucial, as will be discussed in subsequent sections, to avoid gaps and configuration errors an attacker can exploit.

Mitigations and Detections

Figure 5 demonstrates a hypothetical attack chain where both abuse of a legitimate cloud account and abuse of a victim cloud account occur. Throughout this visual, Insikt Group has identified parts of the attack chain where defenders can most efficiently hunt for and mitigate behaviors associated with cloud abuse.

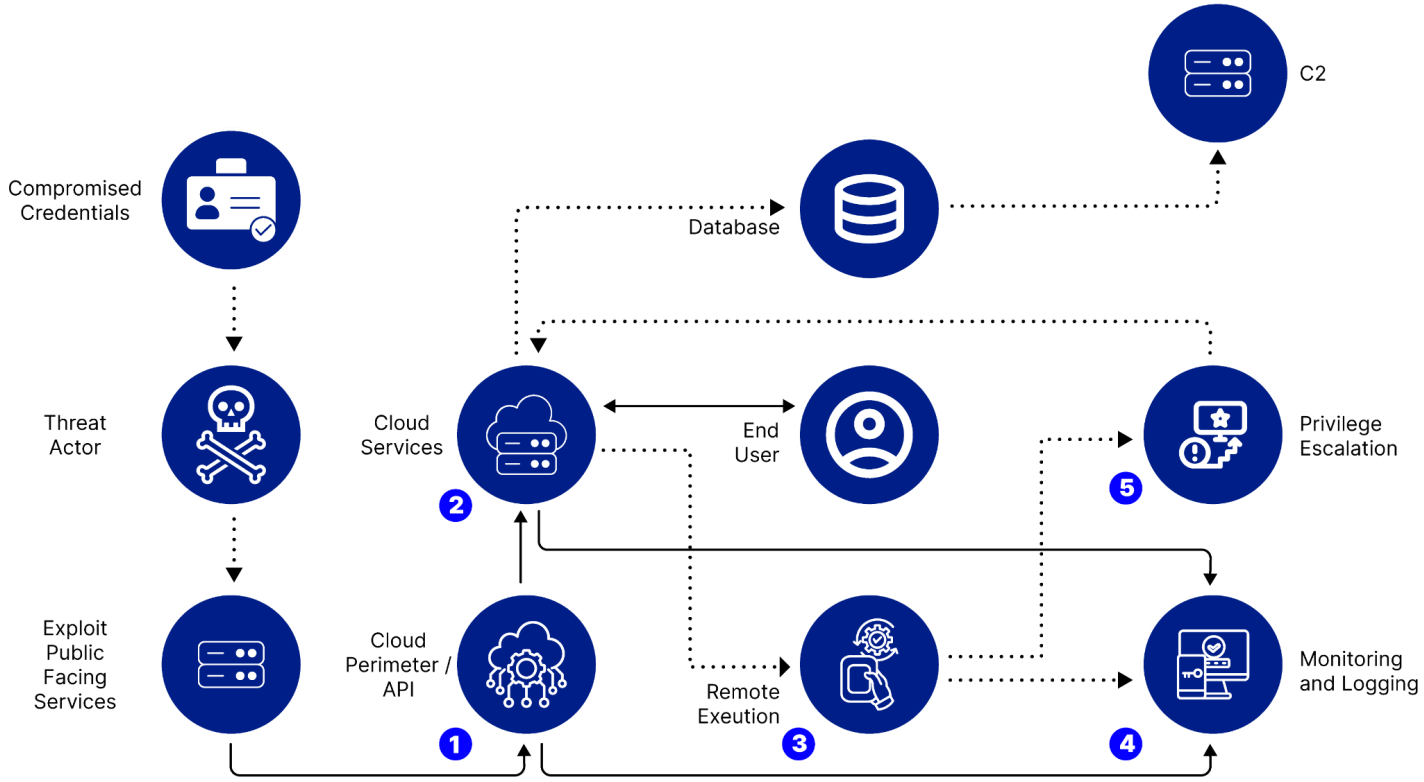

Figure 5: *Visual representation of potential exploitation attack vectors (Source: Recorded Future)*

① Credential-Based Access from Unknown or Rotating IPs

Credential abuse continues to be one of the most exploited vectors for initial access across cloud-focused threat campaigns. Threat actors such as Storm-0940 and Storm-0501 use brute-force and password spray attacks, often leveraging distributed proxy networks or compromised small office/home office (SOHO) routers to mask their origin. These techniques allow attackers to obtain valid credentials, access cloud services undetected, and escalate privileges. Anomalous login attempts frequently precede successful intrusions, highlighting the importance of correlating authentication patterns across identity providers.

Mitigation Recommendations:

- Enforce MFA or passkeys for all cloud accounts.
- Use behavioral baselining to detect anomalous login patterns.
- Correlate failed and successful login attempts by IP address or user agent.
- Block logins from known malicious IP address ranges.
- Employ identity protection and sign-in risk policies.

② Exploit Attempts Targeting Public-Facing Cloud Services and Appliances

Threat actors regularly exploit unpatched vulnerabilities in internet-exposed cloud infrastructure, chaining multiple flaws to achieve initial access and control. Campaigns exploiting Citrix NetScaler, Ivanti CSA, and Oracle WebLogic show how weaknesses in edge systems become primary entry points. These exploits often lead to further actions including remote code execution, credential harvesting, web shell implantation, and lateral movement. Monitoring for suspicious HTTP behavior and scanning patterns is critical for early detection.

Mitigation Recommendations:

- Patch and monitor all public-facing services.
- Restrict external access to critical cloud admin interfaces.
- Deploy virtual patching via WAFs and reverse proxies.
- Detect scanning behavior and exploit patterns via web logs and intrusion detection and prevention systems (IDS/IPS).
- Harden cloud edge appliances.
- Isolate cloud-facing appliances where feasible.

③ PowerShell Use

Post-compromise activity often involves heavily obfuscated PowerShell scripts designed to download payloads, disable defenses, or execute malware. Threat actors such as Water Sigbin and Storm-0501 rely on script-based execution to remain fileless, evade detection, and maintain stealth during lateral movement or data collection. Detection efforts must focus on identifying unusual script execution patterns and encoded payloads indicative of obfuscated activity.

Mitigation Recommendations:

- Restrict script execution to trusted users using AppLocker or Windows Defender Application Control (WDAC).
- Enable script block logging and AMSI in PowerShell.
- Detect and alert on Base64-encoded PowerShell commands.
- Use endpoint detection and response (EDR) tools to hunt for obfuscated or encoded scripts.
- Limit use of interpreters to secured contexts.

④ Unauthorized Remote Access Tools or Anomalous Binary Execution

Threat actors maintain persistence and bypass detection by deploying remote management tools or repurposed binaries. Storm-0501 and other adversaries often use AnyDesk, NinjaOne, and renamed executables to mimic legitimate system activity. Reflective dynamic-link library (DLL) loading, masquerading, and task scheduling are used to evade process monitoring and maintain covert access. These techniques demand strict controls on software installation and binary execution behavior.

Mitigation Recommendations:

- Monitor process execution for known remote monitoring and management (RMM) tools and unusual binary paths.
- Block unsigned or renamed binaries with behavior-based EDR. Similarly implement application allowlisting to block unauthorized binaries.
- Restrict installation of RMM software via group policy object (GPO) or app control.
- Monitor scheduled task creation involving suspicious binaries.
- Audit startup locations and autorun keys.

⑤ Privilege Escalation through Identity Federation or Credential Manipulation

Advanced campaigns increasingly rely on identity abuse in cloud environments, particularly within Microsoft Entra ID. Threat actors escalate privileges by modifying federated domain settings, forging Security Assertion Markup Language (SAML) tokens, or exploiting weakly protected synchronization accounts. These actions allow persistent access and enable threat actors to bypass MFA, access high-privilege resources, and impersonate legitimate users. This technique was observed in multiple campaigns, including those by Storm-0501 and Scattered Spider, underscoring the critical nature of identity-based persistence and escalation in hybrid cloud environments.

Mitigation Recommendations:

- Limit modifications to identity federation settings.
- Enforce conditional access, MFA, or passkeys for all cloud administrative roles.
- Mandate usage of secure password managers.
- Detect SAML tokens issued without MFA using identity logs.
- Alert on the creation of new federated domains or backdoor accounts.
- Use Microsoft Defender for Identity to detect unauthorized privilege elevation or token abuse.

Examples in the Wild

Insikt Group curated a list of events published within the past year that demonstrate the threats posed by exploitation in cloud environments. These events are discussed below.

Citrix NetScaler Instances Targeted in Brute-Force Attacks Exploiting Zero-days

On December 13, 2024, Cyber Security News published a report detailing a surge in brute-force attacks targeting Citrix NetScaler devices. These attacks were carried out by threat actors using outdated NetScaler systems to target and exploit CVE-2024-8534 and CVE-2024-8535. The attackers targeted organizations across various sectors, notably critical infrastructure, as confirmed by reports from the German Federal Office for Information Security (BSI) and other international partners.

CVE-2024-8534 is identified as a memory safety vulnerability leading to potential memory corruption and denial of service, while CVE-2024-8535 allows authenticated users to access unintended

functionalities due to a race condition flaw. Attackers masked their distributed brute-force attack by using disparate IP addresses to avoid detection, complicating defense measures for security teams.

Cyber Security News's report listed multiple IP addresses and IP ranges used in the attack campaign, including but not limited to *45.145.4[.]0/24*, *185.92.180[.]0/24*, *194.113.37[.]0/24*, and *212.87.223[.]3*. The campaign began shortly after the release of patches for the identified vulnerabilities, indicating the threat actors were taking advantage of organizations that had delayed patch deployment.. The US Cybersecurity and Infrastructure Security Agency (CISA) has also issued alerts regarding these vulnerabilities, highlighting the risk of full system compromise if exploited.

Chinese Threat Actor Storm-0940 Employs Credentials from CovertNetwork-1658 for Password Spray Attacks

On October 31, 2024, Microsoft published a report detailing the activities of Chinese state-sponsored threat actor Storm-0940. The report focuses on the group's use of credentials obtained from password spray attacks conducted through a covert network, identified as CovertNetwork-1658, to gain initial access to Azure resources. The covert network was composed of compromised SOHO routers, primarily manufactured by TP-Link. Storm-0904's operations using CovertNetwork-1658 primarily targeted organizations in North America and Europe, including think tanks, governmental and non-governmental organizations, law firms, and the defense industrial base. Storm-0940 has been active since at least 2021 and is associated with initial access via brute-force attacks, password spraying, and misuse or exploitation of network edge applications.

CovertNetwork-1658 is a collection of compromised SOHO routers leveraged to proxy malicious traffic, particularly password spray campaigns. Threat actors exploit vulnerabilities in these routers to gain remote code execution, though the vulnerabilities used in these attacks remains under investigation. Upon gaining access to a router, the threat actor downloads Telnet and xlogin backdoor binaries from a remote FTP server to start an access-controlled command shell on transmission control protocol (TCP) port 7777. A SOCKS5 server is then deployed on TCP port 11288, allowing for anonymized and distributed password spray operations. CovertNetwork-1658's infrastructure involves thousands of rotating IP addresses, with approximately 8,000 compromised devices active at any time and 20% participating in active password spraying.

Storm-0940 rapidly used credentials harvested from CovertNetwork-1658, sometimes on the same day they were acquired, indicating a tightly coupled operational relationship. After gaining initial access, Storm-0940 conducted post-compromise activities, including lateral movement through scanning and credential dumping, targeting network devices to install proxy tools and remote access trojans (RATs), and attempting data exfiltration. Attackers used the following user agent strings during the attacks:

- `Mozilla/5.0 (Windows NT 10.0; WOW64; Trident/7.0; rv:11.0) like Gecko`
- `Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/80.0.3987.149 Safari/537.36.`

Storm-0501 Expands Ransomware Attacks to Hybrid Cloud Environments

On September 26, 2024, Microsoft Threat Intelligence published an analysis detailing a hybrid cloud environment compromise campaign carried out by Storm-0501. The campaign relied on the compromise of Microsoft Entra Connect Sync to pivot from on-premise networks to Microsoft Entra ID, formerly Azure Active Directory (AD). The financially motivated threat group extracted credentials of the on-premise and cloud service accounts involved in identity synchronization, achieving persistent backdoor access to Microsoft Entra ID tenants.

The attack chain began with initial access through compromised credentials or exploitation of public-facing vulnerabilities such as CVE-2022-47966 (Zoho ManageEngine), CVE-2023-4966 (Citrix NetScaler), and ColdFusion flaws (CVE-2023-29300 or CVE-2023-38203). After lateral movement and domain compromise facilitated via Impacket and Cobalt Strike, Storm-0501 identified and accessed Entra Connect Sync servers. The group retrieved plaintext credentials and DPAPI keys stored on disk, allowing authentication to Microsoft Entra as privileged users.

Using these credentials, Storm-0501 established persistent cloud access by using AADInternals to register a federated domain, enabling SAML token forgery for user impersonation. The group set `ImmutableId` properties to impersonate any synced user and bypass MFA, logging in to applications like Office 365 as Global Administrators. In environments with MFA disabled or misconfigured, cloud session hijacking was used against on-premise accounts synced to the cloud. Password resets or credential theft enabled further cloud access escalation.

Post-compromise, the threat actor retained cloud access by creating a federation trust and used forged tokens to bypass security policies. These actions allowed persistent control over Microsoft Entra environments independent of on-premise infrastructure.

Chained Vulnerabilities Exploited in Ivanti Cloud Service Applications for Initial Access

On January 22, 2025, CISA and the Federal Bureau of Investigation (FBI) released a joint advisory detailing exploitation of vulnerabilities in Ivanti Cloud Service Appliances (CSA). The advisory documents threat actors chaining four vulnerabilities, CVE-2024-8963, CVE-2024-8190, CVE-2024-9379, and CVE-2024-9380, to gain initial access, execute remote code, obtain credentials, implant web shells, and in some cases, achieve lateral movement within victim environments. The primary victims are organizations using Ivanti CSA versions 4.6 patch 518 and below and CSA version 5.0.1 and below, with exploitation confirmed in September 2024.

The first exploit chain combined CVE-2024-8963, an administrative bypass vulnerability, with CVE-2024-8190 and CVE-2024-9380, both remote code execution vulnerabilities. Threat actors began by exploiting CVE-2024-8963 to bypass administrative restrictions, using `GET` and `POST` requests targeting the `datetime.php` endpoint to acquire CSRF tokens and manipulate the system's timezone configuration. Exploiting these vulnerabilities allowed attackers to execute Base64-encoded Python scripts, which harvested encrypted admin credentials from the database. Decryption was performed offline or via a PHP file executable located within the `/tmp` directory. The threat actor used the regex `php\w{6}` to find and execute the PHP script. Subsequently, CVE-2024-9380 was exploited to implant

web shells using crafted `POST` requests and to establish reverse TCP C2 channels. `sudo` commands were also used to disable logging mechanisms, modify and remove web shells, and delete evidence of compromise.

The second exploit chain exploited CVE-2024-8963 alongside CVE-2024-9379, a Structured Query Language (SQL) injection vulnerability, to insert malicious SQL statements and create a web shell. This involved a `POST` request manipulating the lockout attempts input box to improperly handle a SQL injection, which inserted `bash` commands to build a web shell. Although threat actors repeated the injection process, there was no confirmation of a successful web shell deployment.

In one incident, lateral movement was observed, with the attackers accessing a Jenkins server. Tools such as Obelisk and Gogo were employed for reconnaissance and vulnerability scanning. Logs from the Jenkins server revealed bash history containing credentials for a PostgreSQL server. Attempts to access the VPN server using these credentials were unsuccessful.

Water Sigbin's Multi-Stage Infection Routine Delivered XMRig Cryptominer

On June 28, 2024, Trend Micro published a report detailing a campaign by Water Sigbin, also known as 8220 Gang, targeting Oracle WebLogic servers, commonly deployed in both cloud-hosted and enterprise environments, to deploy the XMRig cryptocurrency miner. The infection chain is composed of multiple distinct stages that emphasize stealth, persistence, and system resource hijacking within virtualized infrastructure.

The threat actors began their attack by exploiting CVE-2017-3506 and CVE-2023-21839 in vulnerable Oracle WebLogic instances to gain initial access. Following successful exploitation, a PowerShell script was executed to decode a Base64-encoded binary. This binary delivers a disguised executable named `wireguard2-3.exe`, mimicking the legitimate WireGuard VPN client.

The second stage of the attack began with the execution of `wireguard2-3.exe`, which decrypted and loaded `Zxpus.dll`. This DLL contained a second-stage AES-encrypted, GZip-compressed payload that was deserialized and injected via process hollowing into the legitimate .NET binary `cvtres.exe`. The technique leveraged reflective DLL loading to evade detection and ensure that the malicious code remained memory-resident.

In the third stage, the injected payload (`Tixrgtluffu.dll`, PureCrypter v6.0.7D) established persistence by creating a hidden scheduled task under `Microsoft\Windows\Name`, which executed `IsSynchronized.exe`, a replica of the loader placed in `AppData\Roaming\Name`. To further evade detection, the malware disabled Windows Defender protections by executing Base64-encoded PowerShell commands that excluded its executable path and associated processes from antivirus scanning:

- `Add-MpPreference -ExclusionPath …`
- `Add-MpPreference -ExclusionProcess …`

The malware performed system profiling using Windows Management Instrumentation (WMI) to collect details including processor ID, disk signatures, GPU model, username, and antivirus software. This information is hashed with MD5 and encrypted using TripleDES, with the key derived from the mutex `6cbe41284f6a992cc0534b`. The hashed information was then transmitted to a C2 located at *89.185.85[.]102:9091* or *god.sck-dns[.]cc*.

Lastly, in the final stage, the C2 returned a TripleDES-encrypted configuration containing parameters for XMRig. This configuration was stored under `HKEY_CURRENT_USER\SOFTWARE\<Victim ID>`. The final payload, `plugin3.dll`, was downloaded, decrypted, and injected into `AddinProcess.exe`, allowing the miner to operate under the guise of a legitimate system process. The infected host then connected to the mining pool at *217.182.205[.]238:8080* using the wallet address `ZEPHYR2xf9vMHptpxP6VY4hHwTe94b2L5SGyp9Czg57U8DwRT3RQvDd37eyKxoFJUYJvP5ivBbiFCAMyaKWUe9aPZzuNoDXYTtj2Z.c4k`.

Endpoint Misconfiguration

Key Takeaways

- Cloud endpoint misconfiguration is the most common way threat actors gain access to a cloud environment.
- Misconfiguration can stem from the misconfiguration of native cloud assets and technologies that are embedded at the edge of a cloud environment.
- Threat actors will often target misconfigurations in cloud environments opportunistically by performing broad misconfiguration scanning campaigns.

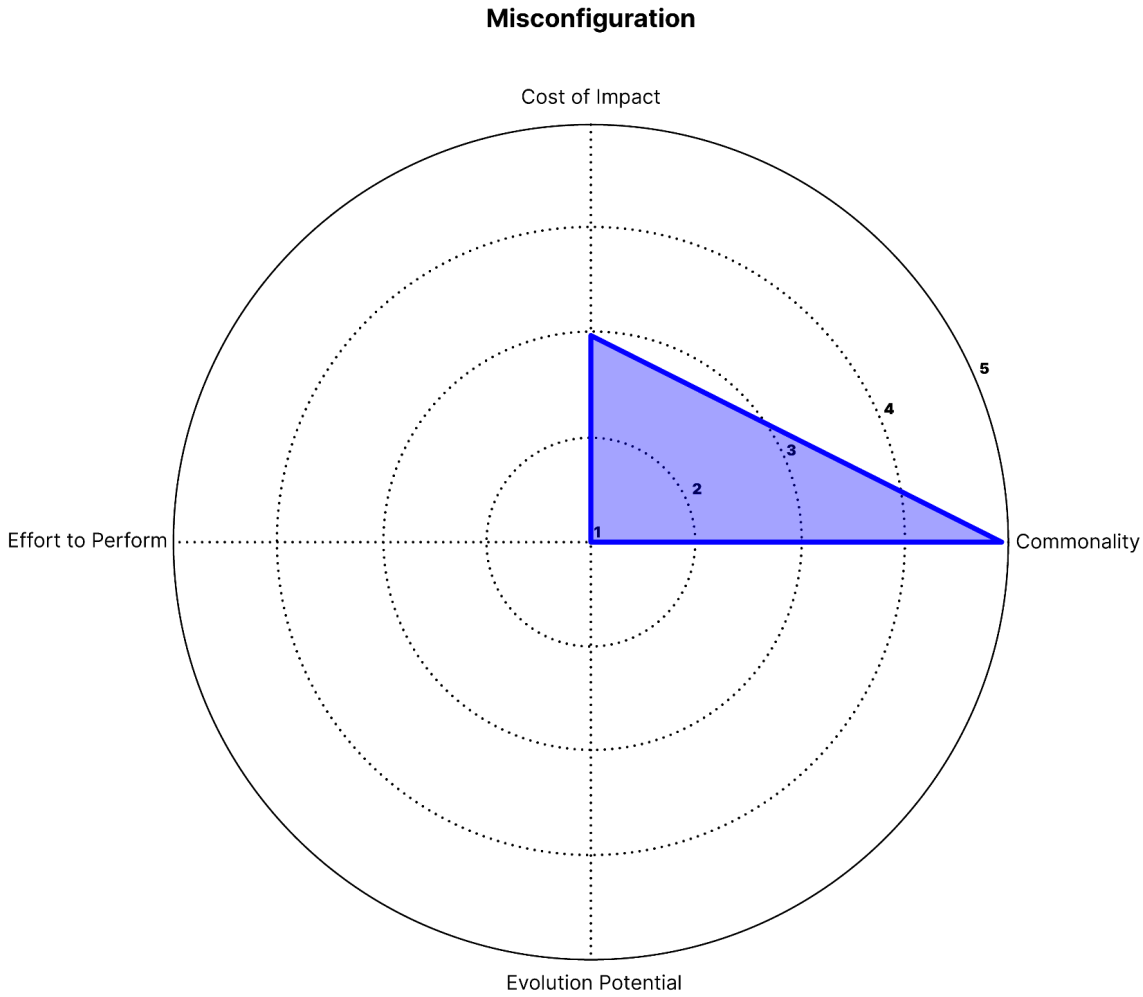

Figure 6: *Radar chart illustrating endpoint misconfiguration as a threat vector (Source: Recorded Future)*

Cost of Impact: 3 (Moderate)

Based on the events discussed in this section, misconfigurations in cloud environments are often exploited to gain initial access to cloud environments and services, usually resulting in data theft and exposure. This outcome may cost the cloud operator financially as well as in terms of reputation.

Commonality: 5 (Severe)

The majority of the events discussed throughout this report include some aspect where cloud endpoint misconfiguration was involved. Furthermore, cloud endpoint misconfigurations are often identified indiscriminately via mass scanning campaigns, allowing threat actors to opportunistically exploit this weakness in bulk.

Evolution Potential: 1 (Minimal)

Since the onus of proper endpoint configuration lies with the owner or operator of the cloud environment, there is little that threat actors can do to evolve their attack methodologies after identifying misconfigured cloud endpoints. The attack chain associated with this threat will always consist of threat actors identifying misconfigured endpoints and then accessing them.

Effort to Perform: 1 (Minimal)

There are many automated endpoint scanning repositories and additional tools for scanning publicly accessible infrastructure, such as Shodan and Censys, that a threat actor has access to. These allow for the discovery of misconfigured endpoints, eliminating the technical barrier to endpoint misconfiguration discovery. Once a threat actor can identify an endpoint misconfiguration, they are often easily able to access or verify the exposed endpoint; they can also do this by using automated tooling.

Threat Summary

Well-architected cloud environments aim to be highly accessible, allowing authorized users and services to freely access its data, services, and systems remotely. Often, this includes providing access to many users at the same time, potentially at a broad geographic level. As such, the cloud environments owned and operated by large corporate entities often require many endpoints situated at the edge of a cloud environment that facilitate both ingress and egress from the environment. However, the access these endpoints provide also carries the risk of unauthorized access should these endpoints be misconfigured.

The misconfiguration of cloud endpoints or technologies hosted at the edge of a cloud environment is a risk that threat actors commonly exploit when targeting cloud environments. By identifying misconfigured or vulnerable cloud endpoints, attackers can gain unauthorized access to a cloud environment and potentially the data and systems hosted within. Threat actors may attempt to identify endpoint misconfiguration in specific cloud environments; however, this threat is often characterized by

opportunistic threat actors that employ broad scanning campaigns or open-source infrastructure searching tools to identify misconfigured cloud endpoints.

Outlook

Threat actors will almost certainly continue to actively hunt for and abuse misconfigured cloud endpoints as an initial access vector during attacks on cloud environments.

As cloud platforms and environments continue to grow and become more widely adopted, the endpoints associated with these cloud environments can be expected to grow at a relative rate. This ever-increasing attack surface creates a compound security issue for cloud defenders and architects attempting to ensure that cloud endpoints are secure, only allowing data in or out based on the principle of least privilege, while also attempting to ensure that data is still accessible and cloud-user communications are feasible.

When attempting to address these concerns across potentially thousands of endpoints, and in some cases third-party technologies embedded within them, the likelihood of risk associated with endpoint misconfiguration increases. Eventually, even in well-architected cloud environments, it is likely that an organization will meet a threshold where the breadth of its cloud infrastructure exceeds the operational capabilities of its security and cloud architecture teams, increasing the risk of misconfiguration. This concept is represented in **Figure 7** below.

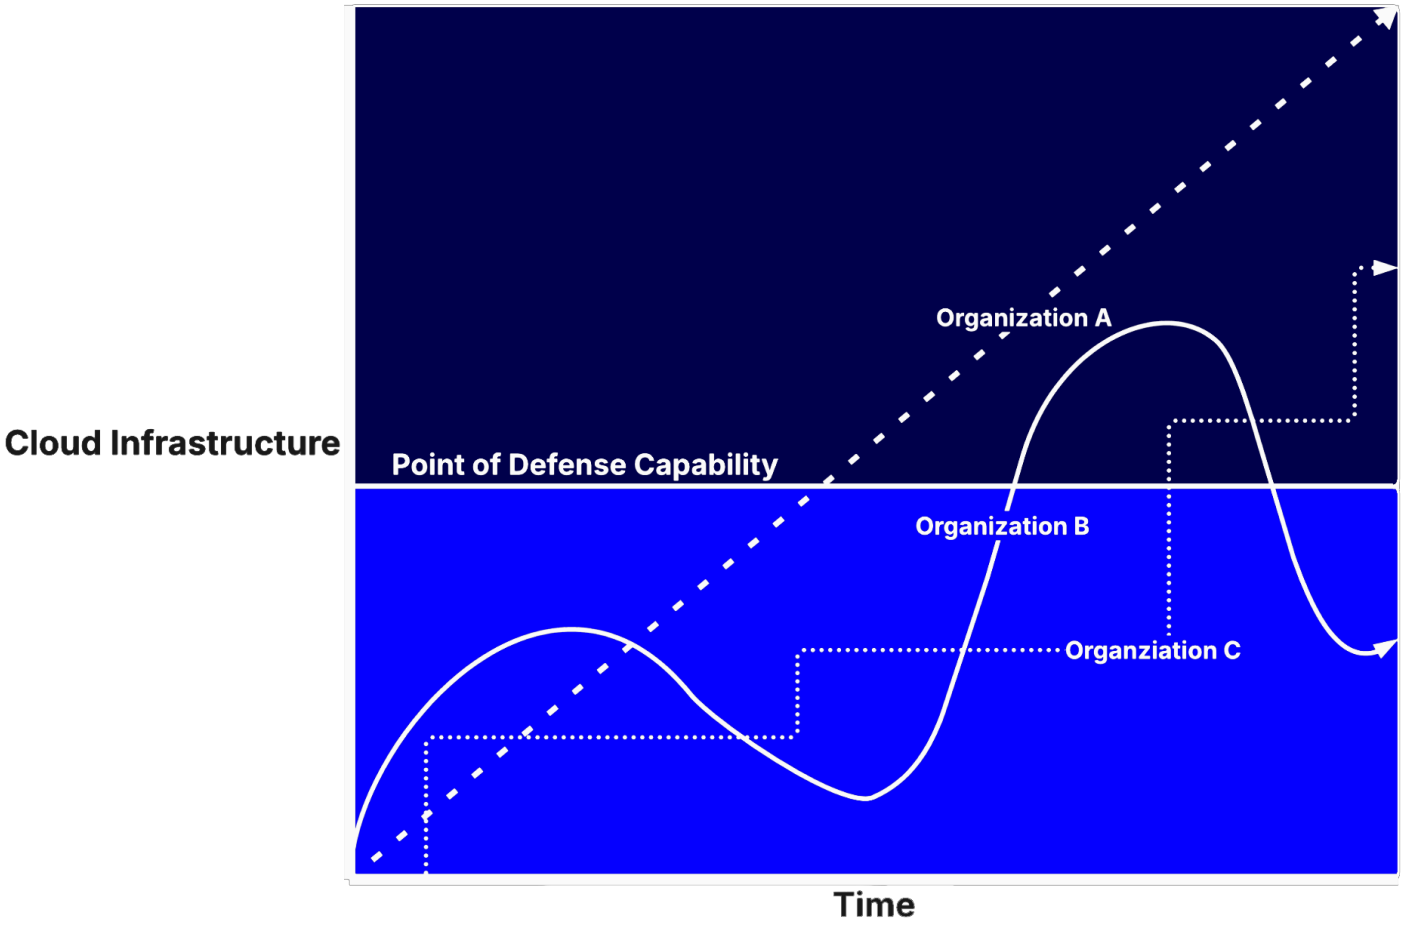

Figure 7: *Theoretical representations of increased risk associated with growing cloud infrastructure (Source: Recorded Future)*

Additionally, as evidenced by the examples presented later in this section and throughout this report, many attacks in the wild rely on endpoint misconfiguration as an initial access vector when targeting cloud environments. Misconfiguration has proven to be a reliable initial access vector for threat actors due to the lack of technical expertise needed to exploit this weakness and the aforementioned issue of defending against this weakness. A multitude of open source cloud misconfiguration scanners are also available from sources such as GitHub, which further lowers the bar in terms of technical literacy needed to identify and access cloud environments.

Mitigations and Detections

Figure 8 demonstrates a hypothetical attack chain where both abuse of a legitimate cloud account and abuse of a victim cloud account occur. Throughout this visual, Insikt Group has identified parts of the attack chain where defenders can most efficiently hunt for and mitigate behaviors associated with cloud abuse.

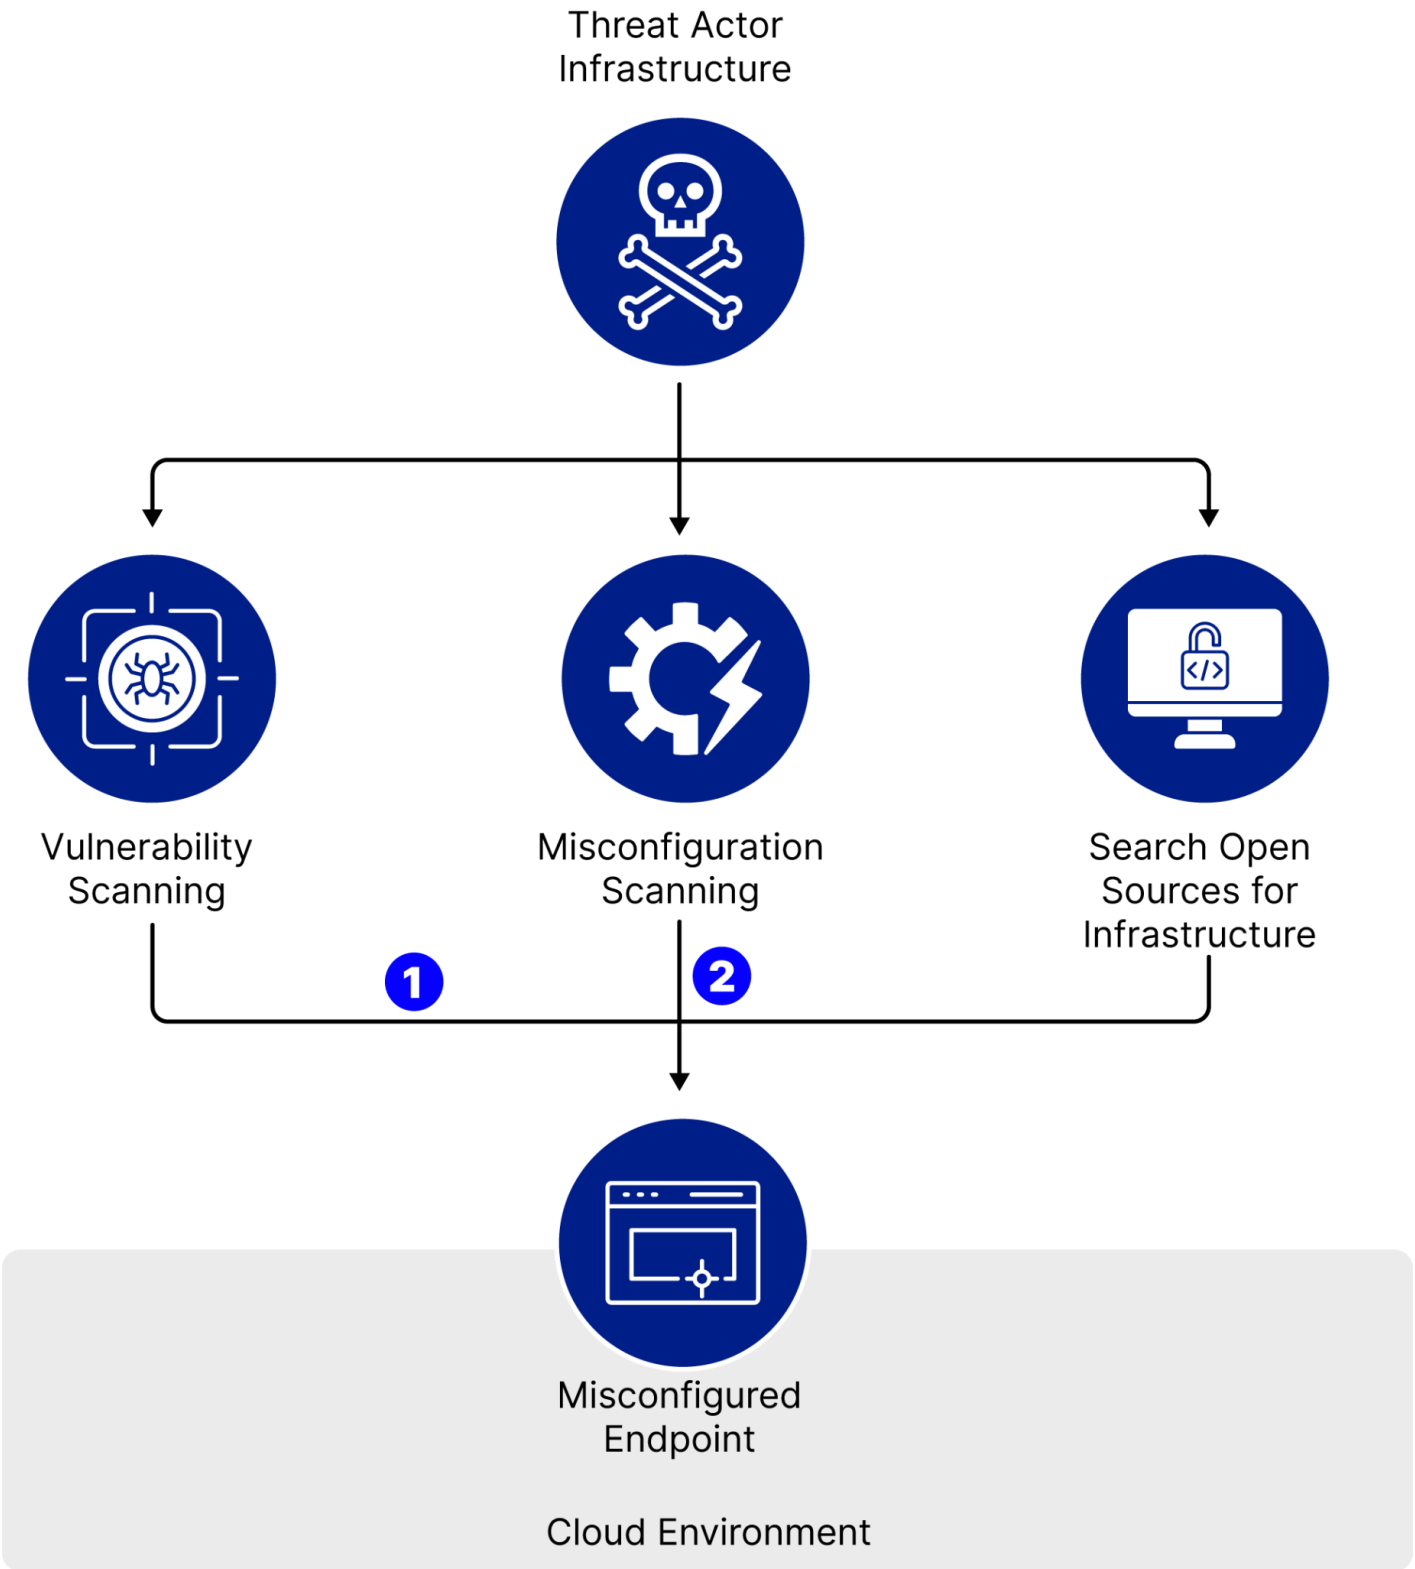

Figure 8: *Visual representation of potential misconfiguration attack vectors (Source: Recorded Future)*

① Vulnerable or Misconfigured Third-Party Applications Hosted at the Edge of Cloud Environments

Third-party software and services are often integrated into the edge of cloud environments to provide functionality that native cloud services or in-house development teams are unable to. While this software may provide operational benefits in these environments, they may require additional maintenance not handled directly by the native cloud environment and are therefore at increased risk of misconfiguration or containing vulnerabilities.

To mitigate the risks associated with this weakness, implement the following mitigations and policies:

- Institute a plan for regular penetration testing, specifically against cloud endpoints where third-party software and applications are hosted.
- Third-party authentication software, such as OAuth, Okta, and similar products, that are used for cloud authentication should be scrutinized to ensure they are properly configured. Threat actors will commonly attempt to abuse these services for initial access and persistence if they are not configured properly.
- Maintain an inventory of the third-party software and services deployed at the edge of cloud environments and compare this list against security bulletins issued by the products' developers to remain aware of high-severity vulnerabilities or other issues associated with these products (such as an uptick in targeting activity against a specific product).
- Whenever a new third-party technology is introduced at the edge of a cloud environment, determine a uniform policy for how this technology should be configured to ensure its deployment is uniform and adheres to the principle of least privilege, both in terms of who can access the service and what the service is capable of accessing or granting access to within the cloud environment.

② Misconfiguration Scanning Attempts

When threat actors attempt to identify misconfigurations at the edge of a cloud environment, this activity will almost always result in the generation of a network request that can be logged or filtered via network security products, native or otherwise, that sit at the edge of a cloud environment. The following general strategies can be used to identify attempts at misconfiguration scanning at the perimeter of a cloud environment:

- Hunt for an irregularly high volume of requests made to a cloud endpoint that originate from an unknown source.
- Hunt for multiple, identical requests made to a cloud endpoint that originate from multiple unknown or unexpected sources.

Additionally, as previously mentioned in this report, threat actors may abuse cloud infrastructure or other legitimate infrastructure to mask malicious behavior as legitimate network activity. In these instances, scrutinizing the requests associated with this network activity is the best option to determine whether this activity is malicious. Maintaining a knowledge base of common requests used when

scanning or attempting to access misconfigured endpoints will assist in the creation of detections for and the identification of malicious scanning activity from otherwise legitimate sources.

Examples in the Wild

Insikt Group curated a list of events published between H1 2024 and H1 2025 that demonstrate the threats posed by cloud endpoint misconfiguration. These events are discussed below.

EMERALDWHALE Targets Exposed Git Configuration Files in Cloud Credential Harvesting Campaign

On October 30, 2024, Sysdig reported on a campaign dubbed EMERALDWHALE, which exploited exposed Git configuration files to harvest over 15,000 cloud service credentials. This operation targeted AWS services, among others, to access sensitive data and leverage compromised accounts for malicious activities like phishing and spam. Sysdig alleged that the attackers used tools such as MZR V2 and Seyzo-v2 to automate the identification of exposed `.git` configuration files and extract credentials, with a specific focus on AWS access keys and secrets. (However, there is no evidence indicating these specific tools were used, aside from the threat actors demonstrating capabilities requiring them.)

Sysdig stated that the attack chain began with the automated scanning of IP ranges for exposed Git configuration files, with attackers likely using open-source tools such as HTTPX to perform the scanning. Once identified, the attackers retrieved sensitive data, including AWS IAM credentials, by cloning repositories and parsing for access keys. The attackers stored their findings in an Amazon S3 bucket, which itself belonged to a prior victim. Exposed Laravel `.env` files and other improperly configured services also provided a pathway for attackers to harvest cloud credentials, further broadening the attack's impact. Ultimately, AWS services like IAM, S3, and Simple Notification Service (SNS) were abused to execute reconnaissance, manage compromised credentials, and create resources for illicit activities.

Leaked Environment Variables Lead to Mass Cloud Extortion Campaign

On August 15, 2024, Unit 42 researchers disclosed a large-scale extortion campaign that targeted cloud misconfigurations in endpoint technologies to steal credentials from environment variable files (`.env`). Threat actors scanned over 230 million targets, identifying `.env` files that contained AWS IAM credentials, database access keys, and API secrets.

Attackers gained unauthorized access to AWS environments by scanning for publicly exposed `.env` files, which contained sensitive information such as cloud API credentials. Using these credentials, the attackers authenticated into the environments and deployed malicious Lambda functions. These functions executed bash scripts and accessed files previously staged in compromised S3 buckets, which contained potential targets or additional payloads. The attackers also performed various AWS API calls to manipulate cloud resources, including the creation of new cloud roles and assets.

To escalate privileges, the attackers created new IAM roles within the victim environments and attached the AdministratorAccess policy to gain elevated permissions. During credential access, they extracted

additional API keys and database credentials stored in publicly accessible `.env` files. For discovery, the attackers enumerated resources by calling AWS APIs to gather information on cloud assets, network configurations, and IAM permissions.

Exfiltration of sensitive data was conducted using the S3 Browser tool, allowing the attackers to transfer data from the victim's cloud environment. The final impact involved extortion tactics, where attackers threatened to release stolen sensitive information unless demands were met.

Cloud Ransomware

Key Takeaways

- Ransomware campaigns are leveraging sophisticated, multi-stage attack chains that combine social engineering, credential abuse, and cloud-native toolsets.
- Exploitation of third-party cloud management tools and backup software introduces systemic, cross-platform risk.
- Persistence and evasion techniques are cloud-optimized, with attackers forging SAML tokens, misusing Cross-Tenant Synchronization, and deploying unmanaged virtual machines for staging ransomware payloads.

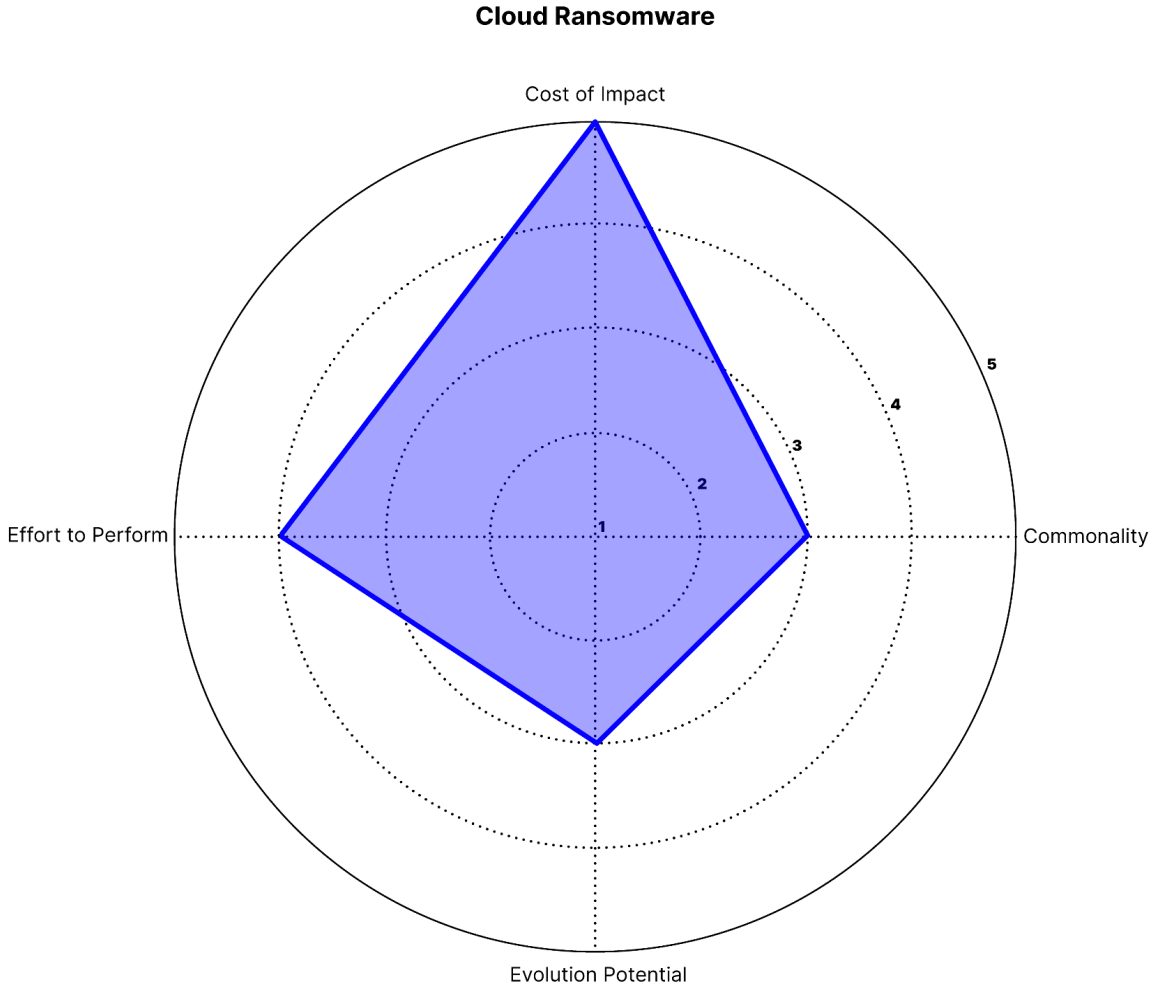

Figure 9: *Radar chart illustrating cloud ransomware as a threat vector (Source: Recorded Future)*

Cost of Impact: 5 (Severe)

Similar to ransomware attacks in traditional IT environments, successful cloud ransomware attacks often result in high costs to victims. These costs extend not only to monetary costs, which are the highest when compared to all other threats discussed in this report, but also to reputational and legal impacts due to name-and-shame and data exposure tactics commonly employed by ransomware threat actors who extort victims for payment.

Commonality: 3 (Moderate)

Based on the evidence found in this report, there is not a steady trend associated with the frequency of cloud attacks, nor are cloud ransomware attacks often marked by clusters of attacks occurring within a short period of time, which is usually indicative of a broader ransomware campaign. However, despite these irregularities, cloud ransomware attacks, or ransomware attacks in traditional IT environments, that occur as a result of cloud compromise appear to have become more common throughout the past year.

Evolution Potential: 3 (Moderate)

Cloud ransomware threat actors have demonstrated that there are multiple ways to encrypt or otherwise ransom data stored within cloud environments. The majority of these techniques, however, rely on the abuse of cloud native services or utilities intended for legitimate encryption. This differentiating feature between cloud ransomware attacks and ransomware attacks against traditional IT environments limits the techniques threat actors can employ during cloud ransomware attacks.

Effort to Perform: 4 (High)

As mentioned above, and based on the evidence found in this section, threat actors must rely on native cloud services and utilities to effectively perform a cloud ransomware attack and must therefore be familiar with the intricacies of the cloud platform to effectively encrypt this data.

Threat Summary

Ransomware remains one of the most persistent and rapidly evolving cyber threats, with a steady increase in reported incidents across diverse sectors. As organizations continue migrating critical infrastructure and data to cloud platforms, ransomware operators have adapted their techniques to target cloud environments. This shift has introduced new attack surfaces and operational complexities threat actors exploit, including cloud-specific vulnerabilities, credential compromise, phishing, and misconfigurations. Cloud-native services and shared responsibility models further complicate security postures, making cloud infrastructure an attractive and viable target for ransomware deployment.

The distributed and scalable nature of cloud systems enables threat actors to move laterally, maintain persistence, and carry out data exfiltration and encryption activities with increased efficiency.

Ransomware campaigns leveraging the cloud often involve multi-stage attack chains, with initial access frequently obtained through social engineering or exploitation of exposed services.

Outlook

Ransomware threat actors are expected to continue intensifying their focus on cloud infrastructure, leveraging cloud-native capabilities and administrative missteps to conduct scalable and disruptive operations.

The ongoing expansion of enterprise cloud adoption is driving a proportional increase in attack surfaces adversaries can exploit. Features such as cloud-native automation, identity federation, and scalable storage solutions offer operational efficiencies not only to defenders, but also to attackers. Threat actors are taking advantage of misconfigured permissions, overly broad access policies, and the misuse of legitimate services like AWS S3 and Microsoft Entra ID to execute ransomware attacks without relying on zero-day vulnerabilities. These native capabilities, when coupled with compromised credentials, allow attackers to encrypt or exfiltrate sensitive data while bypassing many traditional security controls.

The case studies outlined in this section highlight how ransomware operators are embedding themselves deeper into cloud environments by abusing built-in functionality such as custom key management and identity synchronization. Tactics like forging SAML tokens, initiating data encryption through client-controlled keys, and staging payloads on unmanaged virtual machines illustrate a growing sophistication in attack methods. These operations increasingly blur the line between misconfiguration and malicious exploitation, presenting new challenges for defenders as ransomware campaigns become more cloud-centric, automated, and persistent.

Mitigations and Detections

Figure 10 demonstrates a hypothetical attack chain where both abuse of a legitimate cloud account and abuse of a victim cloud account occur. Throughout this visual, Insikt Group has identified parts of the attack chain where defenders can most efficiently hunt for and mitigate behaviors associated with cloud abuse.

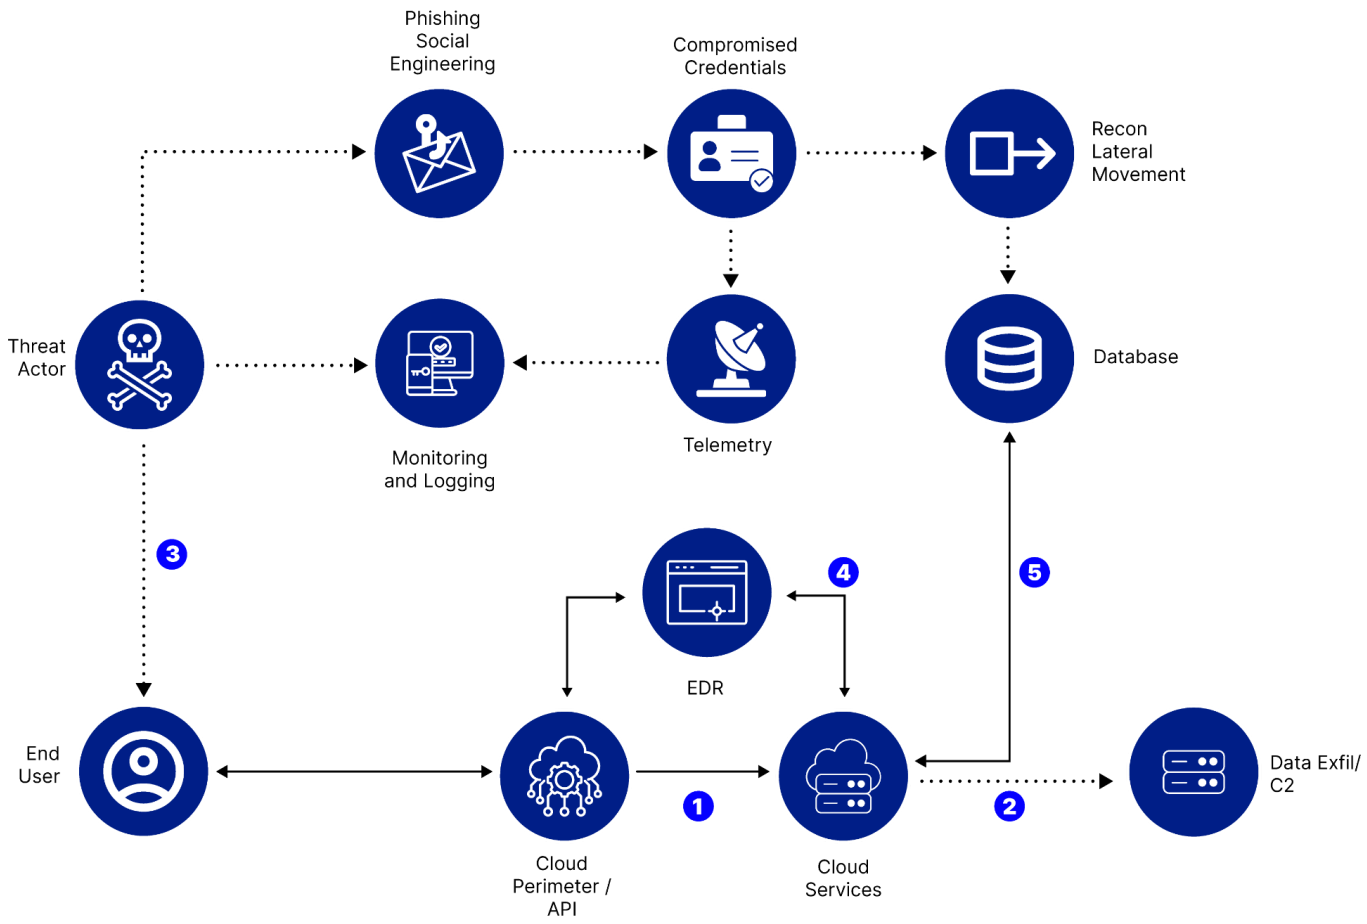

Figure 10: *Visual representation of potential cloud ransomware attack vectors (Source: Recorded Future)*

① Unusual Access Patterns

Compromised cloud credentials are a primary vector in ransomware operations, often obtained via phishing, leaked secrets, or abuse of overly permissive access policies. Attackers use valid credentials to authenticate into environments, enabling them to blend in with legitimate user activity while carrying out malicious operations undetected.

Mitigations:

- Enforce multi-factor authentication (MFA) or passkeys for all user accounts.
- Regularly rotate access keys and credentials.
- Use temporary credentials via IAM roles or identity federation instead of static keys.
- Implement IAM policies that follow the principle of least privilege.
- Alert on credential use from unfamiliar IPs or geolocations.
- Monitor AWS, Google GCP, and Azure environments for unusual access patterns, including the sudden use of inactive or newly created access keys, especially for privileged accounts.

② Cloud-Native Encryption Methods

Cloud-native encryption methods, such as AWS SSE-C or Azure disk encryption, are being repurposed by threat actors to lock critical assets. The use of custom keys not stored within the environment prevents recovery and enables attackers to bypass some traditional anti-ransomware protections.

Mitigations:

- Enable ransomware-specific protections in EDR/XDR platforms.
- Restrict permissions to invoke encryption-related API operations.
- Block use of SSE-C or similar unmanaged encryption keys via policy enforcement.
- Monitor CloudTrail logs for anomalous PutObject requests with encryption headers.
- Implement automated alerts for bulk file changes or extension renaming patterns.
- Alert on mass file renames, use of encryption APIs, or spikes in API calls involving custom encryption headers (for example, Server-Side Encryption with Customer-Provided Keys [SSE-C] in AWS).
- Monitor for entropy changes in file content and the execution of encryption scripts.

③ Phishing and Social Engineering

Phishing remains a foundational tactic to obtain credentials and bypass MFA protections. Threat actors are increasingly mimicking SSO portals and using social engineering via SMS and phone calls to trick users into revealing access information.

Mitigations:

- Deploy phishing-resistant MFA methods such as FIDO2 or number matching.
- Use tools such as the Recorded Future Intelligence Operations Platform to identify and monitor suspicious domain registrations.
- Scan inbound communications for links that mimic legitimate enterprise portals.
- Train employees to recognize phishing indicators and report suspicious activity.
- Enforce secure helpdesk protocols for identity verification during account resets.
- Detect typosquatted domains in emails and SMS messages using threat intelligence products, such as Recorded Future's Brand Intelligence module.
- Monitor for cloned SSO login portals and anomalous domain registrations mimicking enterprise login services.

④ Identity Federation Exploitation

Advanced adversaries are exploiting identity federation to maintain persistent access and elevate privileges. Through manipulation of SAML tokens and federated domains, attackers can impersonate legitimate users or admins without triggering normal security controls.

Mitigations:

- Restrict modification of federated identity and domain sync configurations.
- Require MFA for all SAML token issuance and federation changes.
- Continuously audit logs for changes in authentication provider settings.
- Alert on creation of SAML tokens linked to high-privilege roles.
- Monitor for federation with unknown or unauthorized identity providers.
- Monitor for unauthorized changes to SAML configurations or unexpected updates to federated identity domains.
- Log all token generation events, especially those issued without MFA.

⑤ Use of Cloud Automation Tools to Deploy Payloads

Ransomware operators are leveraging built-in cloud automation tools to deploy payloads without introducing new binaries or malware signatures. These techniques allow malicious script execution using approved APIs, making them harder to detect through conventional means.

Mitigations:

- Restrict use of automation features like RunCommand to authorized roles only.
- Enforce conditional access and Just-in-Time permissions for script execution.
- Audit all command invocations in cloud-native logging platforms.
- Apply behavioral analysis to detect anomalous use of scripting utilities.
- Integrate script execution with security approval and workflow systems.
- **Flag unauthorized or unexpected use of cloud scripting tools like Azure Run Command.**
- Watch for the use of shell scripts invoking ransomware payloads or remote download commands in execution logs.

Examples in the Wild

Insikt Group curated a list of events published within the past year that demonstrate the threats posed by cloud ransomware. These events are discussed below.

Abusing AWS Native Services: Ransomware Encrypting S3 Buckets with SSE-C

On January 13, 2025, the Halcyon RISE Team published a report detailing a novel ransomware campaign in which the threat actor "Codefinger" exploited Amazon S3's Server-Side Encryption with Customer-Provided Keys (SSE-C). The campaign did not rely on any vulnerability within AWS itself, but rather on the abuse of compromised AWS credentials with permissions to read and write S3 objects. This technique resulted in the permanent encryption of data stored in S3 buckets, for which decryption was impossible without the attacker-controlled AES-256 key, effectively locking victims out of their own data unless the ransom was paid.

The attack would begin when Codefinger gained access to valid AWS credentials — usually those leaked publicly or compromised in some other way. To perform the attack, the credentials needed to

include permissions for `s3:GetObject` and `s3:PutObject`. Using native AWS services, the attacker initiated S3 object encryption by specifying the `x-amz-server-side-encryption-customer-algorithm` header and providing a custom AES-256 encryption key. This key was never stored by AWS, and only a hash-based message authentication code (HMAC) was recorded in AWS CloudTrail logs, making forensic recovery impossible. After encryption, Codefinger configured S3 Object Lifecycle policies to mark files for deletion in seven days, creating urgency for ransom payment. The ransom note included a Bitcoin address, a unique client ID, and warnings that tampering with account settings will terminate negotiations.

Scattered Spider Ransomware Campaign Targeting Cloud Infrastructures in Financial and Insurance Sectors

On September 10, 2024, EclecticIQ published an in-depth analysis of cloud ransomware operations conducted by the financially motivated group Scattered Spider, with a focus on attacks targeting the insurance and financial sectors. The article has since been deleted from EclecticIQ's blog, but remnants of the report can be seen on *infostealers.com*. The group uses phishing and social engineering techniques, including vishing and smishing, to compromise IT service desks and identity administrators, enabling access to cloud services such as Microsoft Entra ID and AWS EC2. Credential theft is facilitated through typosquatted domains mimicking single sign-on portals and by exploiting cloud token leaks from public repositories like GitHub. SIM swapping is used to intercept MFA codes, granting access to accounts otherwise secured by two-factor authentication. At the time of writing, ExlecticIQ noted that the group shifted ransomware operations to directly target cloud infrastructure-as-a-service (IaaS) environments for enhanced scalability and impact.

Scattered Spider leverages cloud-native and open-source tools, such as AzureAD PowerShell modules, ADRecon, and PingCastle, for reconnaissance, targeting credentials, network architecture, and sensitive third-party data for extortion. Data exfiltration is performed using tools like S3 Browser and extract, transform, load (ETL) platforms to offload data to attacker-controlled infrastructure.

Persistence and lateral movement within cloud environments are achieved through abuse of Cross-Tenant Synchronization in Entra ID and federated identity providers. The group creates malicious federated domains and forges SAML tokens to maintain access after account remediation. Remote access is maintained using tools like AnyDesk, Ngrok, and Proxifier, which facilitate SSH tunneling and reverse proxy creation.

To evade detection, Scattered Spider employs residential proxies, disables Microsoft Defender and Windows Firewall using open-source scripts, manipulates mail transport rules to suppress security alerts, and reboots systems into Safe Mode to disable protections. The group also creates unmanaged virtual machines within cloud environments to host tools and stage ransomware attacks.

The campaign culminates in the automated deployment of ALPHV Ransomware within cloud environments, particularly VMware ESXi and Azure. Customized scripts are used to stop security

services, execute the ransomware payload, and encrypt data, effectively disrupting operations and increasing the impact on victims.

Ransomware Exploitation of Veeam Backup & Replication in Nigerian Cyberspace

On September 13, 2024, the Nigeria Computer Emergency Response Team (ngCERT) published an advisory identifying an active exploitation campaign targeting Veeam Backup & Replication (VBR) software by ransomware groups, with a specific focus on the Phobos ransomware group. The advisory, identified as NGCERT-2024-0033, details the exploitation of a critical vulnerability, CVE-2023-27532, which has been exploited in recent attacks against cloud infrastructure in Nigeria. The advisory notes a high risk and high damage potential, with the threat affecting Microsoft Windows OS, Linux OS, VMware ESXi, and Oracle platforms.

CVE-2023-27532 impacts Veeam Backup & Replication versions twelve and below, allowing attackers to retrieve encrypted and plaintext credentials from the configuration database. By targeting exposed Veeam instances (`Veeam.Backup.Service.exe`) operating on port 9401, adversaries can issue unauthenticated requests to access sensitive data, including administrative credentials. Attackers are able to use these credentials to perform privilege escalation and arbitrary code execution, subsequently resulting in control over the backup environment, which is often a high-value target containing sensitive data.

Attackers commonly initiate the exploit by scanning for publicly exposed, unpatched Veeam instances. Upon identifying a vulnerable system, they send crafted requests to extract credentials, leading to further malicious activity such as malware deployment, ransomware attacks, data exfiltration, or manipulation. In the documented cases, the Phobos ransomware group used this method to compromise cloud infrastructure within Nigeria. The potential consequences of a successful attack include system compromise, credential theft, script injection, data breaches, reputational damage, denial-of-service conditions, and significant financial loss.

The vulnerability exploitation underscores the importance of securing backup infrastructure, especially given its critical role in recovery and operational resilience.

Credential Abuse and Account Takeover

Key Takeaways

- Credential abuse often results in account takeover in cloud environments.
- Cloud credential abuse can include username and password combinations, API keys, one-time passwords (OTP), authentication tokens, and any other medium that allows an attacker to gain access to a cloud environment.
- Threat actors demonstrated that, following initial access of a cloud environment via compromised credentials, they will almost always perform discovery actions within the cloud environment before any other action is taken.
    - Persistence and lateral movement actions are also common shortly after gaining access to a cloud environment, although less often than discovery actions.

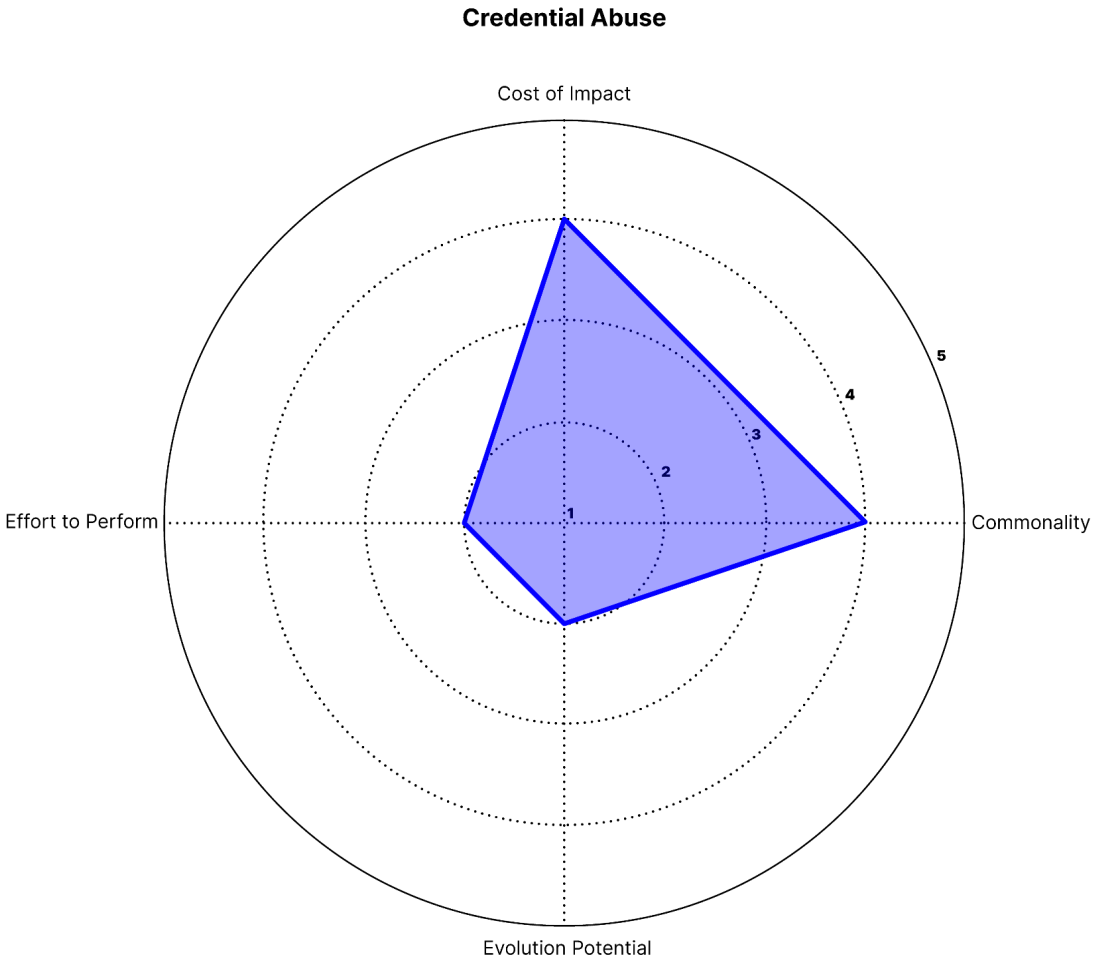

Figure 11: *Radar chart illustrating credential abuse and account takeover as a threat vector (Source: Recorded Future)*

Cost of Impact: 4 (High)

Both cloud accounts and compromised cloud credentials can be used at multiple stages during an intrusion within a cloud environment, granting threat actors increased permissions and potentially access to a broader variety of data and services within a cloud environment. This depends on the credentials and cloud account being abused, but in instances where threat actors are capable of compromising administrator-level accounts, the breadth of an attack can lead to significant data loss or potentially a total environment takeover. As such, victims may incur significant monetary, reputational, and operational losses.

Commonality: 4 (High)

Based on the evidence provided in this report, credential abuse leading to cloud account compromise or access to a cloud environment is the second most common method threat actors employ for initial access in cloud environments, behind misconfiguration exploitation.

Evolution Potential: 2 (Low)

Cloud credentials and, by extension, cloud accounts associated with them, can only be implemented in ways that are predefined by the service where they have been instantiated, limiting the attack techniques threat actors may implement when abusing compromised cloud credentials and accounts. However, threat actors have demonstrated that, by abusing these authentication materials, they are capable of gaining access to additional pieces of information and additional authentication credentials due to flaws in authorizing mechanisms.

Effort to Perform: 2 (Low)

The abuse of compromised cloud credentials is usually straightforward, since these credentials can only be used in the context of how the cloud environment or authentication mechanism expects them to be used. Additionally, due to initial access brokers (IABs), compromised cloud credentials are readily available to threat actors; however, threat actors may need to validate compromised credentials provided by IABs if this information was not previously provided.

Threat Summary

Threat actors commonly gain access to cloud environments by abusing legitimate cloud credentials. Based on Recorded Future observations, this is usually done to perform account takeover of a cloud account, which will grant the threat actor access to and permissions within the cloud environment.

Threat actors will often gain access to valid cloud account credentials via IABs and threat actors selling malware logs. Threat actors may also gain access to valid cloud credentials through a previous malicious action, such as phishing or brute force authentication attempts. Threat actors may also attempt to verify the validity of these credentials via credential spraying campaigns.

Outlook

Threat actors will almost certainly continue to identify and abuse valid cloud credentials. It is also likely that threat actors will continue to abuse valid cloud credentials to compromise cloud accounts.

Valid credential information provides threat actors with a multitude of potential attack methods that can be employed during an attack on a cloud environment — chiefly initial access. Due to the non-technical nature of credential exploitation and the added benefit that abuse of credential information allows threat actors to masquerade as legitimate entities within a cloud environment, threat actors of all kinds implement these credentials wherever possible during an attack on cloud infrastructure and will often attempt to gather more during an attack.

While threat actors will often use cloud credentials to assume the identity of an associated cloud account and gain initial access to a cloud environment, valid cloud credentials may also be used to perform additional malicious actions after gaining access to a cloud environment, especially if the attacker is capable of gaining access to a cloud admin account. The varied benefits of cloud account abuse ensure that attackers will continue to place increased importance on obtaining cloud credentials.

Mitigations and Detections

Figure 12 demonstrates a hypothetical attack chain where both abuse of a legitimate cloud account and abuse of a victim cloud account occur. Throughout this visual, Insikt Group has identified parts of the attack chain where defenders can most efficiently hunt for and mitigate behaviors associated with cloud abuse.

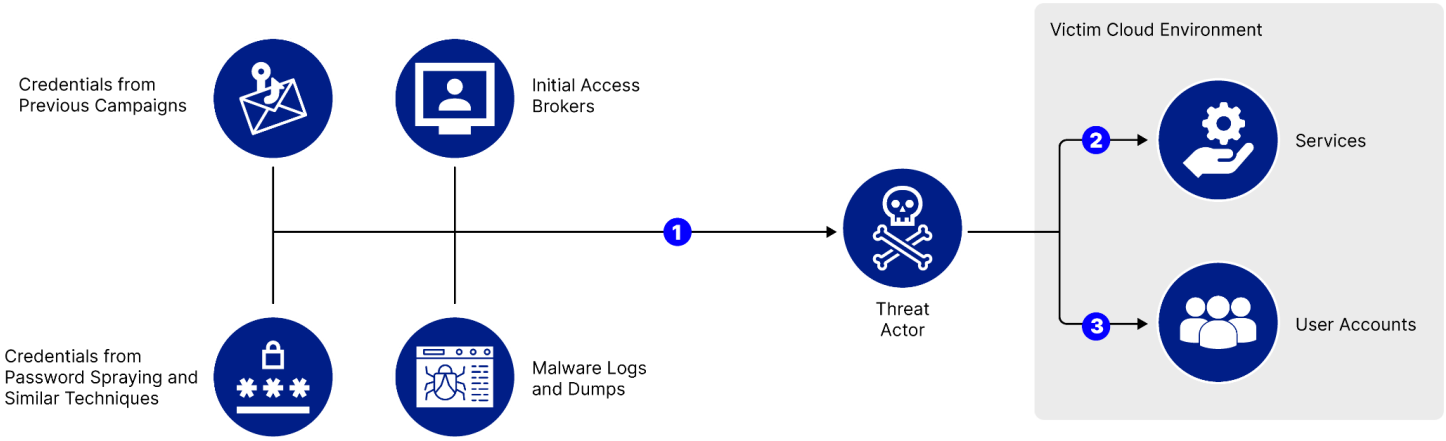

Figure 12: *Visual representation of potential credential abuse and account takeover attack vectors (Source: Recorded Future)*

① Identifying Leaked Credential Data

Threat actors use valid credential data to compromise valid cloud user accounts. To mitigate this threat, threat intelligence platforms, such as the Recorded Future® Identity Intelligence Module, can be used to monitor for exposed credential information that could grant access to a cloud environment.

② Monitoring Service Logs for Suspicious Credential Use and Mitigating the Threat of Valid Credential Abuse

In instances where legitimate credentials have been compromised, defenders can monitor service logs for the following information, which may indicate that credentials have been compromised:

- Log events where an account attempts and fails to access a service or resource due to not having permissions.
    - Identify clusters of similar activity where this behavior is demonstrated across multiple services.
- Log events where a valid account makes a singular request to multiple cloud services within a short period of time.

These behaviors are common after a threat actor has gained access to a valid account and are often associated with discovery operations followed by lateral movement within the cloud environment, privilege escalation attempts, and eventual abuse of cloud resources. Some cloud platforms provide native security services or features capable of identifying this activity by comparing expected user activity against a baseline that is passively generated by normal user activity.

To mitigate the effects of account takeover within a cloud environment, infrastructure segmentation, the implementation of user and role-based permissions (often achieved via role-based access control [RBAC] and identity and access management [IAM] policies and services) and enforcement of the least privilege principal are necessary. By effectively implementing these technologies and strategies, defenders can significantly limit the potential malicious actions threat actors can perform after gaining access to a cloud environment.

③ Monitoring User Access Logs for Suspicious Sign-Ins and Sign-In Attempts

In the event that a threat actor is attempting to use valid credentials that were previously exposed, defenders can identify malicious logins or login attempts in the following ways:

- Identify login attempts originating from a previously unknown IP address or device.
- Identify login attempts originating from a geographic location that is significantly different from where login attempts associated with the account are usually made.
- Identify the method of authentication requested using the credentials and compare this method to how the credentials are usually used.
    - For example, if cloud credentials are used for authentication in an API request when they are usually used for access to a web-based client instead, this behavior should be flagged as suspicious.

Examples in the Wild

Insikt Group curated a list of events published between H1 2024 and H1 2025 that demonstrate the threats posed by credential abuse and cloud account takeover. These events are discussed below.

Silk Typhoon Gains Access to Cloud Environments via Compromised On-Premise Technologies

On March 5, 2025, Microsoft Threat Intelligence reported that Silk Typhoon (AKA HAFNIUM), a Chinese state-sponsored espionage group, has shifted tactics to exploit IT supply chain providers, leveraging stolen API keys and credentials to pivot from on-premise environments to cloud infrastructures.

The group initially gained access by targeting unpatched vulnerabilities in IT management tools, privileged access management (PAM) platforms, and cloud data management applications. Once inside, they escalated privileges by using extracted credentials from Active Directory and key vaults, targeting Microsoft AADConnect servers to synchronize access across on-premise and cloud environments. This allowed the threat actors to manipulate service principals, create rogue OAuth applications, and exfiltrate sensitive data via Microsoft Graph (MSGraph) APIs, SharePoint, and OneDrive.

By exploiting multi-tenant applications, Silk Typhoon can move across customer environments undetected, enabling widespread data theft and persistence. Their cloud-based attack infrastructure includes covert networks of compromised Cyberoam, Zyxel, and QNAP devices, which they use to obfuscate their activities.

Mamba Two-Factor Authentication (2FA) PhaaS Observed Bypassing MFA in Phishing Campaigns

On October 7, 2024, Sekoia.io reported its discovery of Mamba 2FA, a previously unidentified phishing-as-a-service (PhaaS) platform specializing adversary-in-the-middle (AiTM) attacks. First observed in November 2023 and sold through Telegram since March 2024, Mamba 2FA enables sophisticated credential and session hijacking attacks, primarily targeting M365 users. Its phishing campaigns mimic Microsoft login pages, including customized branding for enterprise accounts, and have the ability to bypass MFA methods such as application notifications and one-time codes.

The Mamba 2FA platform generates phishing links and HTML attachments that redirect victims to phishing pages hosted on dynamically linked domains. These domains implement anti-bot detection to identify security tools or automated activity, redirecting unexpected traffic to benign pages like *https://google[.]com/404/*. The phishing page's appearance is dynamically controlled using URL-encoded parameters to emulate Microsoft services like OneDrive, SharePoint, or generic sign-in portals. User actions, including password and MFA entry, are transmitted to backend relay servers using the Socket.IO protocol over WebSockets, enabling real-time credential capture and session replay attacks.

As stated above, the platform's infrastructure consists of dynamically linked domains for user interaction and relay servers for AitM functionality. Relay servers were initially connected directly to Entra ID servers, but Mamba 2FA developers later integrated commercial proxy services from providers

like IPRoyal to obscure their origin in authentication logs. The dynamically linked domains are rotated weekly to evade detection, while relay server domains persist longer.

Valid Credentials Gathered by Infostealers Used to Compromise Customer Snowflake Instances

On June 10, 2024, Mandiant reported that threat actors compromised customer Snowflake instances in an activity cluster they initially reported as UNC5537 (arrests occurred later in 2024, identifying the individuals responsible for the attacks). Mandiant's investigation found no evidence of a direct breach of Snowflake's enterprise environment; instead, the threat actors exploited previously compromised customer credentials, many of which had been obtained from malware infections as far back as 2020. These credentials had been harvested through infostealers such as Vidar, RisePro, RedLine, Raccoon Stealer, LummaC2, and Metastealer, often from contractor-owned devices used for personal activities like gaming and downloading pirated software. By the time Mandiant publicly reported on the campaign, at least 165 organizations had been impacted, with attackers exfiltrating sensitive data from compromised accounts and attempting to sell stolen records on cybercrime forums.

Mandiant determined that UNC5537's attack methodology relied heavily on account compromise due to poor credential hygiene. Many targeted accounts lacked MFA and had not undergone credential rotation for years, allowing the stolen credentials to remain valid long after initial exfiltration. Once inside a Snowflake instance, the attackers leveraged native tools like SnowSight (web UI) and SnowSQL (CLI) to conduct reconnaissance and execute SQL queries for data exfiltration. They also employed a custom reconnaissance utility dubbed FROSTBITE to list users, roles, session details, and publicly available tools like DBeaver Ultimate to execute queries. The attackers staged stolen data using the `CREATE STAGE` command, compressed it with `COPY INTO` using `GZIP`, and finally exfiltrated it using the `GET` command.

At the time of Mandiant's reporting, UNC5537 had primarily accessed Snowflake instances via VPN services like Mullvad and private internet access (PIA) to obscure their origin. At the same time, exfiltrated data was stored on virtual private servers (VPS) from providers such as ALEXHOST SRL and cloud storage platforms like MEGA. The campaign underscored the risks posed by compromised credentials in the infostealer ecosystem and highlighted the necessity of enforcing MFA, credential rotation, and network allowlists to limit unauthorized access. Mandiant had assessed that UNC5537 was likely to continue targeting SaaS platforms using similar tactics, given the availability of extensive credential dumps in underground markets, before subsequent law enforcement actions against individuals associated with the campaign.

General Mitigations

While specific mitigation strategies have been discussed for each of the cloud threats discussed in this report, the following mitigation strategies can be implemented in any cloud environment to increase the environment's security posture:

- Robust logging services and software must be implemented, both to identify potential threats at the time they occur and to remediate the effects of an attack. Enabling logging across network activity, service activity, user activity, asset usage activity (such as virtualization metrics), and cost reporting provide insight into the various systems that are found within a cloud environment, allowing for better threat detection and hunting operations.
- Well-architected and -configured cloud environments greatly reduce the risk of initial access and follow-on malicious actions post-compromise. Implementing strict policies for access points, user account creation and permission granting, and data protection (for example, encryption at rest and in transit; access, modification, and deletion policies; backup retention configurations; and so on) allow defenders to finely forecast and mitigate the risks of cloud compromise and the cost associated with it.
- Due to the pervasive use of legitimate cloud credentials and the increased benefits they grant threat actors during an attack on a cloud environment, using threat intelligence products, such as the Recorded Future Identity Module, can assist defenders in identifying exposed cloud authentication materials prior to their abuse and rotating these materials.
- The defenders and architects of a cloud environment must work in concert to maintain a network map of cloud environments as well as an inventory of third-party or non-native cloud software that is implemented within a cloud environment. Maintaining each of these resources empowers defenders in a multitude of scenarios, such as vulnerability patching for third-party software and identification of attacker-controlled assets created or deployed in the cloud environment.
- While many organizations have opted to fully move their operations to the cloud, many continue to operate in hybrid environments, and all cloud users require a workstation or similar device to access data and systems hosted in cloud environments. As such, traditional IT security mitigations and best practices must be adhered to on such devices, with particular scrutiny applied to cloud accounts with heightened permissions, such as administrator access to a cloud environment.

The most common CSPs provide native cloud services capable of achieving the mitigation strategies discussed above, such as environment scanning for potential misconfigurations, baseline account behavior monitoring, network traffic monitoring, role-based access control (RBAC) or identity access management (IAM) suites, and data protection. However, in cloud environments where there are fewer managed cloud services and more responsibility of defense is more greatly assumed by the cloud defenders and architects, ensuring that clear policies and baselines are specified for the creation, management, and upkeep of an organization's cloud infrastructure is paramount.

Outlook

The observed continuation of cloud abuse as well as the emerging TTPs demonstrated in cloud ransomware campaigns highlights the growing sophistication of threat actors and their ability to exploit cloud-native features and misconfigurations effectively. Credential abuse, particularly involving compromised administrative credentials, will likely remain a primary attack vector, emphasizing the need for robust identity management. Endpoint misconfigurations are also anticipated to persist as significant vulnerabilities, given the operational complexity of securing expansive cloud environments. To counteract these evolving threats effectively, organizations should prioritize timely patch management, adopt comprehensive multi-factor authentication, rigorously enforce least privilege principles, and integrate threat intelligence to proactively detect and mitigate advanced adversary behaviors.

Based on the evidence discussed in this report, opportunistic threat actors appear to mainly target data hosted in cloud environments; however, more sophisticated threat actors and attacks against cloud environments appear to target not only data, but cloud services as well. The most commonly targeted of these services remains compute services, but it is likely that threat actors, especially financially motivated threat actors, will attempt to identify ways that additional, compromised cloud services can be monetized, such as in LLMjacking campaigns.

Currently, the majority of attacks carried out against cloud environments appear to be opportunistic in nature and associated with cybercriminals, with the majority of threat vectors discussed in this report leading directly to monetary gains for the perpetrator. While these attacks remain the most common, state-sponsored threat actor groups also regularly target cloud environments and often display more sophisticated attack chains. In such attacks, state-sponsored threat actors are often observed establishing long-term persistence mechanisms, demonstrating novel attack techniques, and attempting to pivot into additional cloud or on-premise environments while leaving as little trace as possible. These attacks may result in months or years of tenancy within a cloud environment and are conducted mainly to perform information-gathering operations.

While there was not a dominating or overarching vulnerability exploited, it was observed that the exploitation and operations targeting victims using cloud and hybrid environments were highly tailored to the victims' topology. Based on individual implementations, threat actors were very opportunistic in their approach to leverage disclosed vulnerabilities before victims patched their systems. Compromised credentials, whether through broker purchase, social engineering, or password spray attacks, continue to be a dominant source of weakness exploited to gain initial access, escalate privileges, and maintain a foothold within victim environments. Combining all of these tailored methods, attackers are able to successfully steal critical data and ransom cloud databases to extort victims.

Recorded Future reporting contains expressions of likelihood or probability consistent with US Intelligence Community Directive (ICD) 203: Analytic Standards (published January 2, 2015). Recorded Future reporting also uses confidence level standards employed by the US Intelligence Community to assess the quality and quantity of the source information supporting our analytic judgments.

About Insikt Group®

Recorded Future's Insikt Group, the company's threat research division, comprises analysts and security researchers with deep government, law enforcement, military, and intelligence agency experience. Their mission is to produce intelligence that reduces risk for customers, enables tangible outcomes, and prevents business disruption.

About Recorded Future®

Recorded Future is the world's largest intelligence company. The Recorded Future Intelligence Operations Platform provides the most complete coverage across adversaries, infrastructure, and targets. By combining precise, AI-driven analytics with the Intelligence Graph® populated by specialized threat data, Recorded Future enables cyber teams to see the complete picture, act with confidence, and get ahead of threats that matter before they impact your business. Headquartered in Boston with offices around the world, Recorded Future works with more than 1,900 businesses and government organizations across 80 countries.

Learn more at recordedfuture.com