Google Cloud Security

Mandiant

M-Trends

2024 Special Report

Cloud Intrusion Trends

As enterprise cloud adoption and the use of hybrid cloud/on-premises environments continues to grow, adversaries have followed similarly in their targeting. Attackers recognize the value of data stored in cloud environments, and the computing resources available that could enable future malicious operations. Mandiant continues to observe attackers of varying motivations pivot to cloud environments to target cloud-hosted data, and leverage cloud computing resources in their operations.

Targeting Identity and Access Management and Bypassing MFA Requirements

Historically, to gain initial access to cloud and hybrid environments, attackers have relied upon stolen credentials and access tokens that did not require multi-factor authentication (MFA). As security awareness and MFA adoption has increased in recent years, attackers have placed an increasing emphasis on social engineering. During targeted social engineering campaigns, attackers lure users into providing credentials and using innovative methods to circumvent MFA or exploit weaknesses in its implementation.

Mandiant has observed increased usage of adversary-in-the-Middle (AiTM) techniques to bypass MFA requirements by capturing session tokens. In an AiTM campaign, the targeted user's connection to the legitimate cloud service is proxied through an attacker controlled server. By doing so, the user's credentials and MFA method are relayed to the legitimate cloud service while the attacker can capture the access token returned to the user. There are a number of AiTM kits advertised by attackers that can help automate the construction of convincing landing pages that mimic the legitimate cloud service logon page. In the majority of business email compromise (BEC) cases Mandiant responded to in 2023, successfully targeted users had MFA configured, but it was circumvented by an AiTM phishing campaign.

Attackers are known to leverage social engineering to target users, and Mandiant continues to observe the effectiveness of these campaigns. An espionage-related investigation revealed the use of a heavily tailored spear-phishing email that impersonated an individual in the same industry as the target, and used legitimate content relevant to the user to build credibility. The email lured the targeted user to click a link and enter their credentials to access protected information. Once in possession of the username and password, the attacker triggered a MFA push notification, which the targeted user accepted.

Additionally, Mandiant observed attackers abuse the trusted role of help desk and technical support personnel. In one case, phishing messages purporting to be from technical support lured users into providing MFA approval for malicious signins to a cloud platform. In another case, the financially motivated attacker UNC3944 relied heavily on social engineering to obtain credentials of users with elevated privileges, leveraging SMS phishing and placing phone calls to the targeted organization's help desks to reset a user's password or associated MFA device.[1]

In 2023, we observed an increase in the use of targeted SIM swapping to gain access to accounts. In cases where this was effective, organizations sent time-based one-time password MFA codes via SMS messages. In other cases, organizations used SMS to verify ownership of an account before sending a password recovery link. Financially motivated attackers have been observed leveraging SIM swapping to receive both types of SMS codes to facilitate an account takeover. A financially-motivated threat cluster, UNC3786, routinely performed SIM swapping to compromise credentials and gain access to targeted organizations' Okta and/or Microsoft 365 accounts. In one UNC3786 intrusion, while performing a SIM swap, the attacker sent spam SMS messages to the targeted user's phone, likely in an attempt to distract the individual from notifications from their phone carrier related to the SIM swap. UNC3944 similarly leveraged SIM swapping to gain access to user credentials and SMS MFA codes.[2]

Mandiant continues to observe attackers perform password guessing attacks against cloud sign-in portals to identify accounts that do not have MFA configured. Often organizations will rely on users to self-enroll an MFA device. This means that if an account does not already have MFA, the first successful password authentication will immediately prompt to enroll an MFA device. We have seen espionage attackers perform these guessing attacks to find and take over dormant accounts without MFA that should have been disabled, or service accounts with no need for MFA.

Weak Credential Storage

In other cases, Mandiant observed attackers using credentials that were stored poorly to gain access to cloud environments. In one incident, default configurations on an Internet accessible server led to the discovery and compromise of clear text AWS credentials, which allowed the attacker to gain access to a target's AWS environment. In another case, an attacker leveraged account credentials believed to have been stolen during a previous incident, enabling the attacker to gain access to the targeted organization's cloud-hosted code repository, which did not require MFA. Mandiant also responded to an incident involving an attacker that accessed a targeted organization's AWS environment using a leaked AWS access key. The investigation revealed that the key likely originated from a Docker container, hosted on an EC2 instance and owned by the organization, that was exposed to the internet. Copies of the same key were also found in other public IP addressable resources not owned by the targeted organization.

Adversary Abuse of Cloud Services

After obtaining initial access to cloud environments, Mandiant observed adversaries abuse cloud native tools and services to maintain access, move laterally, and ultimately accomplish mission objectives such as stealing data. By limiting themselves to preinstalled tooling, attackers can decrease their operational profile, evade detection, and maintain presence in cloud environments for longer periods of time.

Mandiant has observed attackers using Azure Data Factory and AirByte to modify existing pipelines to steal data stored in various integrated platforms such as data warehouses, storage blobs, and SQL databases. Specifically, attackers have created pipeline jobs that export data from those data sources to an attacker-controlled SFTP server. The use of data factories provided the attackers with a stable and high-bandwidth platform to copy large volumes of data.[3]

In 2023, Mandiant observed a financially-motivated attacker, UNC3944, backdoor cloud identity providers (IDP) using techniques that were previously only observed in use by espionage groups. In multiple investigations, UNC3944 gained administrative access to Entra ID (formerly Azure AD) to configure a rogue federated identity provider, which allowed them to execute golden SAML attacks. The attacker could then authenticate to resources protected by Entra ID as any user in the organization without knowledge of their password or possession of their MFA device. In one investigation, Mandiant observed UNC3944 target an organization's Active Directory Federated Services (ADFS) server, and execute Mimikatz in an attempt to obtain the token signing certificate to conduct a golden SAML attack.

Mandiant also saw attackers target cloud compute instances to maintain stealthy persistence in target cloud environments. In multiple incidents the attackers created Azure Virtual Machines (VMs) and assigned them public IP addresses. These attacker-created VMs did not have the organization's mandated security and logging software installed on them. As such, the attackers gained unmonitored access to a trusted system inside of the organization's virtual network or virtual private cloud, which they then use to progress their intrusion.[4]

Cloud compute instances often have network connectivity to organization on-premises networks via virtual private network, and can provide an avenue for lateral movement. On multiple occasions UNC3944 moved laterally from Azure console access into an Azure-hosted VM using the Special Administration Console to connect to VMs via serial console. Attackers have employed malicious use of the Serial Console on Azure VMs to install third-party remote management software, which provided them with persistent access to the VM. This method of attack is unique in that it avoided many of the traditional detection methods employed within Azure, and provided the attacker with full administrative access to the VM.

Mandiant has also observed attackers target cloud infrastructure for the specific purposes of cryptomining based on their perceived significant processing power. In one incident, an attacker gained access to a targeted organization's Google Cloud Platform (GCP) project via a leaked service account key. After accessing the project, the attacker deployed more than 1,200 virtual machines with a startup script to execute a Monero cryptocurrency miner using the XMR-Stak miner.

Attackers also leveraged open-source offensive security toolsets to survey the environment in some cases. Tools such as Pacu, an open source AWS exploitation framework, and CloudFox, an open source command line tool that can enable the discovery of exploitable attack paths into cloud infrastructure, were seen in one case being leveraged to perform automated reconnaissance. ScoutSuite, a cloud security and auditing tool, was used in another case to access AWS API and Console to conduct operations, which included the deployment of a cryptocurrency dataminer in the AWS environment.

Recommendations

Mandiant continues to observe attackers targeting weakly implemented identity management practices and credential storage to obtain legitimate credentials and circumvent MFA. As attackers have developed new methods to bypass MFA, organizations need to implement changes to their authentication policies to maintain a strong security posture. Phishing-resistant MFA methods have gained widespread support by web browsers, operating systems, and cloud service providers. The two most commonly accepted phishing-resistant MFA methods are certificate-based authentication (CBA) and FIDO2 security keys. Organizations can limit a attacker's ability to circumvent MFA protections by phasing out legacy MFA methods such as SMS, phone calls, and TOTP codes in favor of these newer methods.

In CBA, users are provisioned a certificate identifying them and their device as well as a corresponding private key. The private key is used to prove the user's identity to the cloud service provider. The user is never prompted for the private key, and in fact they do not know it. Almost all end-user systems in use today contain a Trusted Platform Module (TPM), which is a separate chip that stores and manages cryptographic keys securely, including the private key used in CBA. The key material never leaves the hardware boundary of the TPM, so it cannot be stolen by malware on the device. When CBA is used, the connection is negotiated using mutual TLS authentication, which resists most phishing methods because the secret key is not known to the user, and they are never prompted for it. Techniques that use AiTM are also thwarted because when the attacker proxies a targeted user's login session they must terminate the TLS connection, which will break CBA.

FIDO2 security keys rely on a similar security model as CBA to be phishing resistant, with the major distinction being the portability of the keys. With a FIDO2 security key, often a USB device, the key material never leaves the hardware boundary of the physical key. Users do not know and are not prompted for the key's value. Each website configured for FIDO2 has a unique private key that is tied to a particular application. If a user visits a phishing website, the browser will refuse to prompt the user for their security key because the phishing domain does not match any configured keys.

Cloud service providers (CSP) provide many tools to organizations to help detect and prevent common cryptominer schemes. First, all CSPs support authentication methods that provide additional security features beyond access keys and API secrets. Organizations should audit their cloud accounts and establish a program to remove any access keys, especially "root" or "superuser" access keys, and move towards modern role-based programmatic access secrets. Additionally, budgeting alerts and limits are a great way to monitor cloud accounts for abnormal spending. This is often a high-fidelity signal of a cryptocurrency scheme that has hijacked a cloud account. Finally, consider using a "secure by default" design for any new cloud environments. Secure by default bakes in vendor best practices to reduce the attack surface.

Organizations should also consider implementing additional controls to restrict access to cloud resources to only trusted devices. Each CSP supports this in some form although the exact language may differ. This can be done by using mobile device management (MDM) technologies to maintain device state, and allow authentication only from enrolled devices. If organizations require access to resources using untrusted devices (for example, a hotel computer) they should put in place policies that restrict what can be accessed, and how. For example, users should not be able to access the cloud administration console from an untrusted device. Similarly, downloading documents on untrusted computers should be limited or restricted.

[1] https://cloud.google.com/blog/topics/threat-intelligence/unc3944-sms-phishing-sim-swapping-ransomware
[2] https://cloud.google.com/blog/topics/threat-intelligence/sim-swapping-abuse-azure-serial
[3] https://cloud.google.com/blog/topics/threat-intelligence/unc3944-sms-phishing-sim-swapping-ransomware
[4] https://cloud.google.com/blog/topics/threat-intelligence/unc3944-sms-phishing-sim-swapping-ransomware

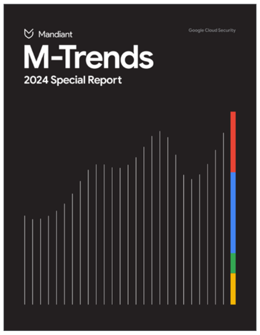

Read the full report: M-Trends 2024 Special Report

Google Cloud

For more information, visit cloud.google.com/security.