

EMBARGOED UNTIL
9AM EST JANUARY 22, 2025

Google Cloud
Security

Threat Horizons

H1 2025 Threat Horizons Report

Office of the CISO

Table of Contents

Mission Statement	03
Executive Summary	04
New Data Points to Threat Actors Increasingly Exploiting Overprivileged Service Accounts to Move Laterally	06
The Boundary of Identity	11
Database Security: Critical Cloud Protection	14
Threat Actor Spotlight: UNC2165 Ransomware and Data Theft Extortion	16
Disrupting Financially Motivated Threat Actors Conducting Cloud Hijacking Campaigns	21
Growing Threat from Data Leak Sites Enabling Extortion in the Cloud	24
Contributors	27



Mission Statement

The Google Cloud Threat Horizons Report provides decision-makers with strategic intelligence on threats to not just Google Cloud, but all providers. The report focuses on recommendations for mitigating risks and improving cloud security for cloud security leaders and practitioners. The report is informed by Google Threat Intelligence Group (GTIG), Mandiant, Google Cloud's Office of the CISO, Product Security Engineering, and various Google Cloud intelligence, security, and product teams.

Executive Summary

Evolving Ransomware and Data Theft Risks in the Cloud

Cloud environments are facing an evolving threat from threat actors prioritizing data exfiltration, exploiting identity as the new perimeter, and adapting tactics to evade detection and attribution. This iteration of the Google Cloud Threat Horizons Report provides cloud security professionals with a deeper understanding of the threat with intelligence and actionable risk mitigations from Google's security experts.

Ransomware and data threats in the cloud are not new. In Feb. 2024, Google Cloud security and intelligence experts exposed trends in the [Threat Horizons Report](#), including threat actors prioritizing data exfiltration over encryption and exploiting server-side vulnerabilities. Further, our experts cited ransomware and data theft incidents or associated risks in cloud environments in our ten [previous](#) Threat Horizons Reports.

Despite the ongoing presence of ransomware and data theft risks, the trends we observed in the last half of 2024 reveal a concerning shift. Threat actors are not only refining their tactics, techniques, and procedures (TTPs) within cloud environments, but they are also becoming more adept at obscuring their identities. This evolution makes it harder for defenders to counter their attacks and increases the likelihood of ransom payments.

Recognizing our shared fate in defending against evolving cloud threats, this Google Cloud Threat Horizons Report delivers timely analysis and actionable mitigations for the recent ransomware and data theft trends that our security and threat intelligence experts have identified and are disrupting in the current threat landscape:

- **Risks to service accounts:** Google Cloud research shows that over-privileged service accounts and lateral movement tactics are increasingly significant threats, even though credential and misconfiguration issues remain common for initial access.
- **Identity exploitation:** Compromised user identities in hybrid environments can lead to persistent access and lateral movement between on-premises and cloud environments, subsequently resulting in multifaceted extortion.
- **Cloud databases are under attack:** Threat actors are actively exploiting vulnerabilities and weak credentials to access sensitive information.
- **Increased adaptability:** Threat actors are leveraging Ransomware-as-a-Service (RaaS) offerings and adjusting tactics to evade detection and attribution.

- **Diversified attack methods:** A threat actor group we track as TRIPLESTRENGTH uses privilege escalation, including charging against victim billing accounts to maximize profits from compromised accounts.
- **Threat actors are using increasingly sophisticated tactics to steal data and extort organizations in the cloud:** Threat actors are using multifactor authentication (MFA) bypass in cloud-based services to compromise accounts and aggressive communication strategies with victims to maximize their profits.

To stay ahead of the curve in 2025, a robust cloud security strategy must prioritize data exfiltration and identity protection. The following content provides cloud security decision-makers with the latest intelligence on threat actor tactics and actionable mitigations to better inform cloud data security strategies.


New Data Points to Threat Actors Increasingly Exploiting Overprivileged Service Accounts to Move Laterally

New Google Cloud research highlights that threat actors are shifting focus. Instead of solely focusing on stealing user login information and exploiting misconfigurations to gain initial entry, they are now targeting overprivileged service accounts, or accounts that have more privileges than necessary. By exploiting these accounts, actors can more easily move laterally within an organization’s systems, potentially causing more damage from their intrusions. This research shares key internal cloud risk factors that make organizations more susceptible to these attacks.

Looking At Organizational Cloud Risks and Alerts

In H2 2024, Google Cloud extended our research focus to include instances where organizations inadvertently neglect critical risks and thus provide threat actors new avenues for exploitation. Our data on flagged instances of insecurity in organizations (Fig. 1) and subsequent analysis indicates that service account keys need further attention.

Figure 1: Cloud Risk Alerts Detected H2 2024



Nearly half (46.4%) of the observed security alerts were due to overprivileged service accounts. We also saw a notably high percentage of alerts (10.3%) from service account keys being leveraged in an unexpected location, which suggests a possible high false-positive activity, including events such as traveling employees, shifts in computing resources, or other mundane causes.

The takeaway of organizational cloud risk alert data points to the need for organizations to investigate and protect service accounts so they can prevent exploitation of overprivileged accounts and reduce detection noise from false positives.

How Threat Actors Get In: Credentials and Misconfiguration

Looking over a two year period, data for threat actor initial access trends remained largely consistent, mirroring observations from previous periods, as illustrated by the lower three lines in Fig. 2.

During H2 2024, credential-related vulnerabilities like weak or no passwords continued to be the most common entry point for attackers as shown in Fig. 3, though the frequency decreased slightly through 2024 (Fig. 2). Misconfiguration of cloud environments (services or software) remained a significant security gap.

Figure 2: Initial Access Vectors of Concern 2022-2024





Figure 3: Initial Access Vectors of Concern H2 2024



We also observed a new trend in the second half of 2024: a sharp rise in compromised APIs and UIs due to threat actor targeting. These attacks accounted for 17.1% of observed incidents, a substantial increase from the approximately 13% observed in the first half of 2024. This upward trend in API/UI exploits reflects a persistent and longstanding security concern cited in

[previous](#) Google Threat Horizons Reports, highlighting the need for organizations to bolster security measures specifically designed to protect against unintentionally exposed APIs and UIs. Google Cloud offers [Advanced API Security](#) to address this risk.

Post-Initial Access Efforts: Lateral Movement Dominates


New data also reveals valuable insights into the actions threat actors take once they've gained access to an organization's systems, as illustrated in Fig. 4.

More than half (62.2%) of threat actor movements once they gained access involved attempting lateral movement within an environment and downloading tools designed for this purpose. To help address the need for effective detective controls capable of identifying and remediating lateral movement on cloud

assets, Google Cloud Security Command Center offers customers [Event Threat Detection](#) in the Premium tier.

We also saw a significant trend in threat actors searching for insecure private keys (13.7%), reinforcing the need for organizations to prioritize the security and proper management of private keys. Access token manipulation also appeared often enough to be notable (11.3%), highlighting the ongoing importance of Identity and Access Management (IAM) as a critical security focus area.

Figure 4: Observed Impact of Intrusions H2 2024



Mitigations

We recommend the following risk mitigations to enhance your Google Cloud security posture to help protect against threats to service accounts:

- **Reduce service account key risk:** Consider [alternative solutions](#) to using service account keys to [reduce](#) this attack surface. When they cannot be removed, review [best practices](#) for managing service account keys.
- **Restrict service account key creation:** Use organization policies to restrict service account key creation and limit the roles assigned to service accounts. Google Cloud [restricts](#) creation by default for new customer organizations created after May 2024.
- **Optimize identity and access management (IAM) policies:** Ensure only necessary services have access to critical assets, and regularly review IAM policies to apply the principle of least privilege. Consider using [IAM Recommender](#) to help navigate proper permissioning for roles.
- **Enhance internal threat monitoring:** Reinvigorate lateral movement detection technologies and policies for internal-facing sensors.

The Boundary of Identity

As organizations expand the cyber boundary to cover a hybrid plane of on-premises, multi-cloud, and multi-Software as a Service-based applications, the common “boundary” has shifted from the network perimeter to the identity plane. With the expansion to cloud, the scope of what represents an “[identity](#)” has also expanded, including managed identities (typically associated with human interaction to resources) and identities associated with workloads (including programmatic/automated interaction). Without proper controls and processes, a single compromised identity could cause a disproportionately impactful cyber event, including data theft and/or ransomware deployment, causing significant damage to organizations.

Identity Threats

Identity compromise is no longer limited to password theft based upon [misconfigurations or weak passwords](#). Threat actors are now gaining access by intercepting or stealing post-authenticated tokens or cookies, effectively bypassing traditional authentication criteria. The most common methods of identity compromise include brute-forcing using common/guessable passwords, replaying stolen credentials from a previous breach, credential stuffing, phishing, and social engineering.

Organizations have responded to these growing threats by enhancing authentication requirements like enforcing multifactor authentication (MFA), but threat actors continue to adapt their techniques by invoking SIM swapping, MFA fatigue (push/text-based notifications), Adversary in the Middle (AitM) attacks, and targeted social engineering—masquerading as a trusted resource to convince someone to provide MFA codes or accept an MFA validation prompt.

A single stolen credential can initiate a chain reaction, granting attackers access to applications and data, both on-premises and in the cloud. This access can be further exploited to compromise infrastructure through remote access services, manipulate MFA, and establish a trusted presence for subsequent social engineering attacks. Stolen credentials can also be used to register malicious applications for persistent access to communication platforms, or to obtain long-lived credentials like access keys and certificates, further solidifying a foothold. Ultimately, an initial credential compromise can enable attackers to pivot across on-premises or cloud infrastructure, escalate privileges, and establish persistence, resulting in stolen data, extortion, and destructive activities.

Mitigations

To protect against the impact of attacks on the identity plane, organizations must level up their authentication processes and playbooks for responding to threats in the following ways:


Combine strong authentication with attribute-based validation: The authentication process should not be based on a single identity attribute like a password. Rather, organizations should shift towards the concept of a **positive identity transaction**, which requires **strong authentication** (e.g., phishing

resistant MFA methods, passwordless) combined with attribute-based validation (Fig. 5), which may include:

- Geo-verification for where the authentication request was initiated
- Identity risk reviews and verification (suspicious logins, leaked credentials, atypical travel, recent changes to identity attributes following a large scope of access attempts)
- Time-based access enforcement (Just-in-Time) or predefined session durations based upon sensitivity of what is being accessed
- Device state review and verification (pre-defined attributes, trusted health status)

Figure 5: Positive Identity Transaction

Passphrase (Password) + MFA Method (Device) + Contextual Information =
Positive Identity Transaction



Comprehensive identity incident response:

In addition to fortifying authentication processes, organizations must also modernize playbooks and processes for proper containment and remediation for identities, which may include:

- Enforcing [mandatory MFA](#) for an account if not already configured
- Disabling and rotating credentials for an account
- Revoking access tokens within the identity provider (IdP)/cloud platform(s)/accessible application(s)
- Revoking cookies for authenticated identities within applications
- Reviewing, revoking, and regenerating programmatic/long-lived identities (access keys/certificates)
- Reviewing registered devices associated with compromised identities, and revoking any devices which are unauthorized/recently added
- Reviewing enforced MFA methods associated with compromised identities, and removing any methods that are weak or subject to MFA bypass techniques
- Reviewing and revoking (credentials/access) for any newly registered applications associated with a compromised identity

Database Security: Critical Cloud Protection

Threat actors are increasingly targeting identities and databases, exploiting misconfigurations and vulnerabilities to gain access to sensitive information and resources. Insecure databases containing critical business data and personally identifiable information (PII) are particularly attractive targets. Once inside, attackers can leverage compromised credentials to move laterally and access potentially even more valuable data, which may lead to additional attack paths with accesses that may contain even more valuable information.

Our Google Cloud experts observed this pattern, wherein the same group repeatedly targets environments using a combination of authentication attacks and exploitation of software vulnerabilities. The objective remains consistent: gaining unauthorized access to the underlying cloud infrastructure to conduct follow-on malicious activities. Threat actors often automate the process of scanning and identifying vulnerable databases. They scan for open ports, weak passwords, and other misconfigurations. Once they compile a list of potential targets, they sort and prioritize a target list based on perceived ease of exploitation and potential return on investment (ROI). These threat actors operate with exceptional speed, exploiting vulnerabilities within their 'n-day' window and capitalizing on older, unpatched vulnerabilities before remediation can occur.

This article examines a large-scale campaign by threat actors that used a layered initial access method and exploited password spraying, misconfigurations, and brute-forcing. Our analysis of these tactics, techniques, and procedures (TTPs) reveals how prevalent the risk is across all cloud providers.

The Convergence of Threats

During the second quarter of 2024, we observed a widespread campaign where a group of threat actors used Kinsing malware, also known as H2Miner, to target publicly exposed PostgreSQL databases. Notably, these threat actors typically deploy the Kinsing malware on compromised cloud resources. The group demonstrated increasing sophistication in their TTPs, specifically around maximizing illicit profits by targeting Monero—a decentralized privacy-focused cryptocurrency that uses cryptographic techniques to mask user identities and transaction details—rather than other cryptocurrencies. The group's use of a cronjob with a backup command and control (C2) server highlights their focus on ensuring persistent access and evading detection. They then targeted insecure PostgreSQL databases, achieving initial access through brute-forcing weak credentials.

Targeting Monero is indicative of the group's diverse capabilities, expanding its toolset to conduct more complex operations. The choice to target a privacy-centric cryptocurrency may suggest a deliberate

intent to obfuscate financial transactions and further evade detection. This intentional targeting of privacy-enhanced technologies, combined with other TTPs, signals a higher level of operational security and technical expertise on the part of the threat actor. This pattern of persistence mirrors tactics observed in previous attacks on Apache Solr instances, specifically the creation of a new cronjob linked to a secondary C2 server. The tactic is suggestive of the operator's familiarity with techniques used to target this specific software or a broader campaign impacting multiple systems. By tracking this campaign, we discovered additional compromised servers beyond the initially compromised databases.

Mitigations

Google Cloud offers a variety of services and products to help ensure the security and integrity of managed databases. Recommendations include:

- **Secure private connections:** If you must use a [private IP](#) with Cloud SQL, configure authorized networks on the SQL instance to restrict access.
- **Enable logging & monitoring:** Detecting brute-force attacks requires monitoring failed login attempts, which [Cloud Monitoring](#) can help with by focusing on log-based metrics and alerts. Enable Database Audit Logs within [Cloud SQL and MySQL Audit Logs](#) for your Cloud SQL instance to track administrative access and system changes to complement database logs which may not capture individual failed login attempts. [Cloud Logging](#) can also help by collecting and analyzing logs from Cloud SQL for PostgreSQL instances on database activity, user access, and more. Additionally, the 'pgAudit' extension could be configured to log failed PostgreSQL attempts.
- **Use robust Identity and Access Management (IAM):** [Google Cloud IAM](#) provides comprehensive IAM controls, including [PostgreSQL auth proxy](#) for secure authentication and authorization. This feature centralizes resource control and minimizes unauthorized data access by authenticating users through IAM, reducing password leakage risks.
- **Proactively approach vulnerability management:** [Google Cloud SQL security recommenders](#) automatically monitor your databases for vulnerabilities like public IP access, missing patches, and unencrypted connections. These timely alerts and actionable remediation steps can help prevent incidents by ensuring your databases are always protected.
- **Enhance data protection with Virtual Private Cloud (VPC) service controls:** Google Cloud [VPC Service Controls](#) enable data exfiltration protection and allow list management which controls access to your Google Cloud resources within a perimeter. [Think Outside the Perimeter: Bug Hunting in Google Cloud's VPC Service Controls](#) offers insights on the Google Cloud Vulnerability Reward Program.

Threat Actor Spotlight: UNC2165

Ransomware and Data Theft Extortion

UNC2165 is a set of financially motivated threat actor activity dating to at least 2019 that abuses cloud services to host data exfiltrated from victim environments. The threat actors behind this activity have shifted to using new ransomware families over time, likely in response to [sanctions](#) and their desire to hinder attribution efforts by security defenders. UNC2165 has notable [similarities](#) to operations publicly attributed to Evil Corp, including a heavy reliance on FAKEUPDATES infections to obtain initial access to victims and overlaps in their infrastructure, and use of particular ransomware families. UNC2165 has impacted nearly every industry, including healthcare, retail, construction, engineering, legal and professional services, with victims located in North America, Europe, Asia Pacific, and the Middle East, according to Mandiant.

Threat Actors Abusing Cloud Storage Services to Host Data Exfiltrated from Victim Environments

Beginning in December 2023, UNC2165 resumed their intrusion operations after a period of dormancy dating back to April 2023. UNC2165 has almost exclusively obtained initial access to victims' networks from UNC1543, financially motivated threat actors that have distributed FAKEUPDATES since at least April

2018. Consistent with previous campaigns, UNC2165 leveraged UNC1543 distribution channels, involving search engine optimization (SEO) poisoning and FAKEUPDATES, to deliver the COLORFAKE.V2 in-memory dropper and MYTHIC payloads. Historically, UNC2165 operations heavily relied on BEACON for lateral movement and to maintain access to the victim environment. However, as of late 2023, UNC2165 has used the MYTHIC post-exploitation framework in intrusions.

In one incident, a user who visited a legitimate website was redirected to a malicious website that led to a drive-by download of an UNC1543 FAKEUPDATES payload. After obtaining access to a victim network from the FAKEUPDATES infection, UNC2165 deployed MYTHIC to establish a foothold. UNC2165 performed reconnaissance using built-in Microsoft Windows utilities and leveraged RDP for lateral movement to connect to an on-premises file server. UNC2165 then used Rclone to exfiltrate data to an attacker-owned Azure blob storage before deploying RANSOMHUB ransomware. Threat actors commonly use the Rclone utility to synchronize files with cloud storage providers. Mandiant also observed UNC2165 access the victim's cloud-based data backup instance where they deleted backup routines, deleted existing data and backups, and modified user permissions to hinder response and recovery efforts.

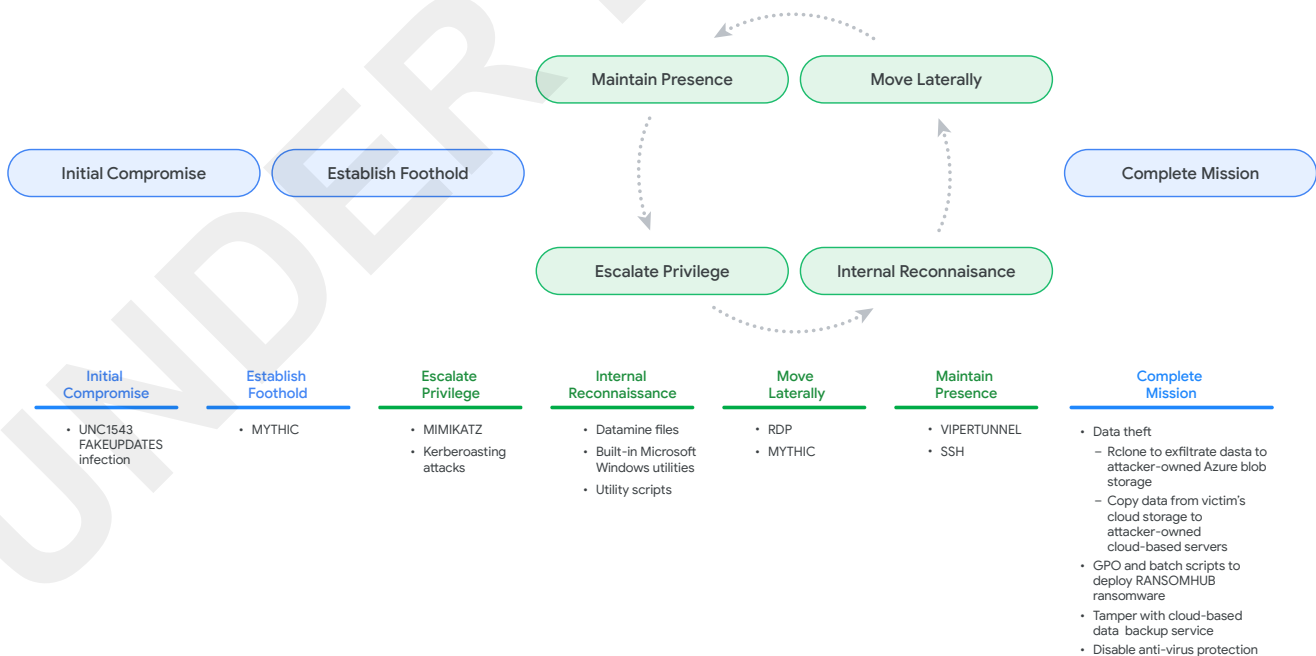
In a separate incident in 2024, UNC2165 gained access to a victim environment via an UNC1543 FAKEUPDATES infection, then deployed their Python tunneler, VIPERTUNNEL to maintain presence in the environment, and executed a series of utility scripts to perform reconnaissance and disable anti-virus protection (Fig. 6). UNC2165 then accessed the victim’s Azure blob storage and directly copied data to multiple attacker-owned cloud-based servers. Having successfully exfiltrated sensitive data, UNC2165 moved on to the disruptive phase of their operation leveraging Group Policy Objects (GPOs) to deploy malicious scheduled tasks which subsequently executed RANSOMHUB on Windows systems in the

victim’s on-premise environment. UNC2165 also deployed an Azure run command to execute a bash script on Linux systems hosted in Azure to download and execute a sample of the Linux variant of RANSOMHUB ransomware.

Threat Actors Demonstrate Adaptability with New Ransomware

Mandiant has observed UNC2165 evolve their use and adoption of new ransomware families in their operations, including HADES, LOCKBIT, CONTI, and RANSOMHUB (Fig. 7). Based on the overlaps between UNC2165 and Evil Corp, Mandiant assesses with

Figure 6: UNC2165 Attack Lifecycle in Observed Intrusions Leading to RANSOMHUB Deployment

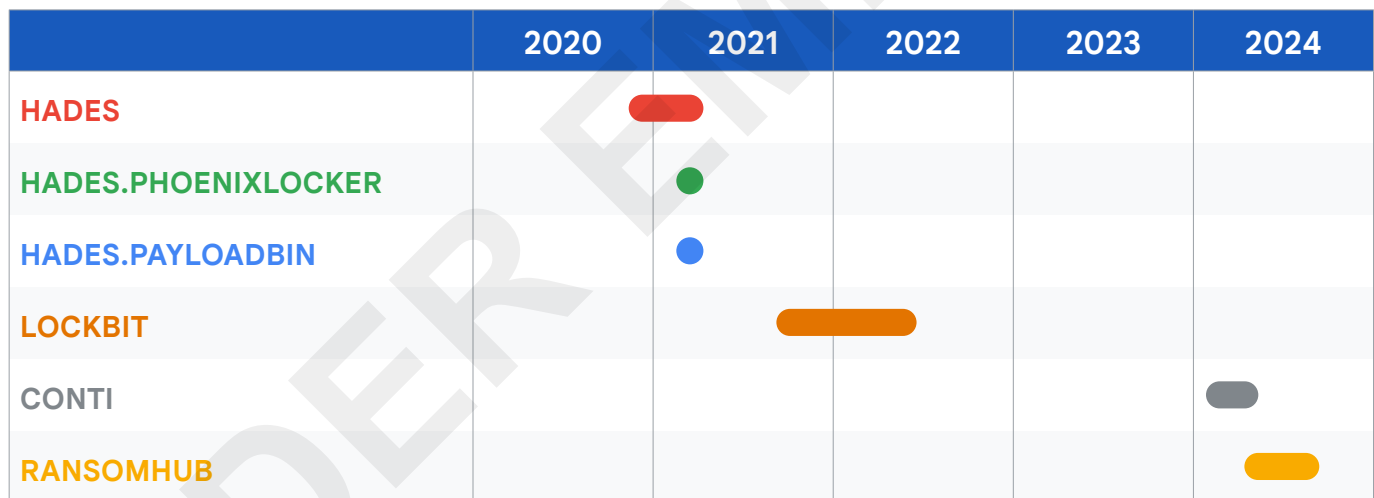


high confidence that actors affiliated with UNC2165 have shifted away from using exclusive ransomware variants and instead are employing well-known ransomware-as-a-service (RaaS) offerings in their operations, likely for the purposes of hindering attribution efforts in order to evade sanctions.

UNC2165 shifted from the use of HADES, a private ransomware family, to LOCKBIT in October 2021. Following the LOCKBIT disruption in February 2024, UNC2165 briefly leveraged CONTI—almost certainly

from the leaked CONTI source code—before shifting to deploying RANSOMHUB by April 2024. Frequent code updates and rebranding of HADES required development resources. It is plausible that UNC2165 saw the use of LOCKBIT as a more cost-effective choice as it would eliminate the ransomware development time and effort allowing resources to be used elsewhere, such as broadening ransomware deployment operations.

Figure 7: UNC2165 Ransomware Observation Timeline




By using common RaaS offerings, UNC2165 effectively obscures their identity and increases the likelihood of receiving ransom payments. A victim would need visibility into earlier stages of the attack lifecycle to properly attribute the activity, compared to prior UNC2165 operations that could be attributable based on the use of an exclusive ransomware family. Given the sanctions against Evil Corp, victims might be hesitant to pay a ransom if they suspected actors associated with Evil Corp were behind the attack, as doing so could have legal repercussions.

Mandiant has observed recent changes in the ransomware landscape, pointing to the likelihood that at least some actors who previously leveraged LOCKBIT, including UNC2165, have shifted to other

ransomware offerings, such as RANSOMHUB. Notably, in Q3 2024, RANSOMHUB displaced LOCKBIT as the most impactful RaaS brand accounting for approximately 15% of victim organizations listed on data leak sites (Fig. 8). While LOCKBIT was consistently the highest volume RaaS offering through the first half of 2024, Mandiant observed a decrease in the proportion of victims appearing on its data leak site in the months following its disruption. Based on victims posted to prominent data leak sites, LOCKBIT accounted for approximately 21% of victims in Q1 2024, 17% of victims in Q2 2024, 7% of victims in Q3 2024, and less than 1% of victims in Q4 2024. In Q4 2024, RANSOMHUB maintained its place as the most impactful RaaS brand, seemingly replacing LOCKBIT as the most prominent brand in the RaaS market.¹

Figure 8: Distribution of LOCKBIT and RANSOMHUB Victims in 2024



¹ Mandiant tracks multiple data leak sites (DLS) dedicated to releasing victim data following ransomware and/or data theft extortion incidents in which victims refuse to pay a ransom demand. These websites are intended to pressure victims to pay the ransom demand or give threat actors additional leverage during ransom negotiations. While the data is skewed to victims who refused to pay the attackers' ransom demand, it does provide visibility into the most active operations during this time frame. These sites also provide insight into broad victimology trends, although not all the data leak sites or victims appearing on them are associated with operations in which ransomware was deployed.

Mitigations

Google Cloud offers a number of capabilities to help customers protect against threat actors similar to UNC2165 conducting ransomware and data theft extortion operations.

- **Regularly review user permissions:** [Google IAM](#) allows you to review and remove unnecessary permissions allowing users or roles to create, modify, or execute serverless resources. [IAM recommender](#) can help identify and remove excess permissions from principals in Google Cloud.
- **Enhance cloud security posture with Google SecOps:** Leverage automated alerts from [curated detection](#) rules, such as the “Fake Updates-related Wscript Execution and Network Connection” rule, to identify and mitigate threats from FAKEUPDATES targeting Windows systems.
- **Use Virtual Private Cloud (VPC) Service Controls:** Enhance your defense-in-depth strategy by using [VPC Service Controls](#) to restrict access based on IP address, identity, and trusted client devices; and by leveraging [Cloud Logging](#) to log access denials for review.
- **Define access policies:** Google Cloud [Access Context Manager](#) allows an administrator to create detailed access control rules based on user attributes like device, location, and identity.
- **Strengthen ransomware defense with protection and containment strategies:** [Practical Guidance for Hardening and Protecting Infrastructure, Identities and Endpoints](#) offers detailed steps to identify vulnerabilities and harden your environment, helping you proactively prevent ransomware attacks.

Disrupting Financially Motivated Threat Actors Conducting Cloud Hijacking Campaigns

Since 2023, teams across Google Cloud have worked to disrupt a financially motivated actor that the Google Threat Intelligence Group tracks as “TRIPLESTRENGTH”. This actor engaged in a variety of threat activity, including cryptocurrency mining operations on hijacked cloud resources and ransomware activity. Additional opportunistic threat activity includes:

- Account hijacking focused on cloud service accounts to mine cryptocurrency
- Ransomware and extortion operations, including activity that overlaps with multiple ransomware deployments and efforts to recruit partners in blackmail operations
- Advertising access to servers, including those from cloud platforms such as Google Cloud, Amazon Web Services, Microsoft Azure, Linode, OVHCloud, and Digital Ocean

Cloud Account Hijacking Activity

To take over cloud service accounts, TRIPLESTRENGTH leverages stolen credentials and cookies, at least a portion of which have come from Raccoon infostealer logs, to gain access to victim cloud environments. Once authenticated, the actor has used hijacked cloud projects to mine

cryptocurrencies. TRIPLESTRENGTH has adapted their abuse of compromised accounts over time. Initially, they abused legitimate, compromised accounts to create compute resources for mining. In subsequent iterations of the campaign, the threat actor abused highly privileged accounts to invite attacker-controlled accounts as a billing contact on the victim’s cloud project, and then used enhanced billing privileges to spin up large compute resources.

In response to TRIPLESTRENGTH’s activity, we collaborated on a multi-quarter effort to disrupt the actor and enhance Google Cloud’s security features. These efforts include the phased roll out of [mandatory multifactor authentication \(MFA\)](#) to reduce the risk of account takeover. Other improvements include [improved logging for sensitive billing actions](#), detection pipelines to identify and mitigate new compromise attempts, and threat actor tracking to identify and disrupt their account infrastructure.

Connections to Ransomware Operations

TRIPLESTRENGTH is also involved in ransomware deployment operations. However, they appear to keep their ransomware activity separate from their cryptomining efforts. Notably, we have not

seen their ransomware operations impact cloud infrastructure, rather these operations have appeared to target on-premises resources. The Google Threat Intelligence Group identified shared email addresses, infrastructure overlaps, and underground forum posts indicating that the actor's activity overlaps with PHOBOS, RCRU64, and LOKILOCKER ransomware deployments.

In Telegram channels focused on hacking, actors linked to TRIPLESTRENGTH have posted advertisements for RCRU64 ransomware-as-a-service and also solicited partners to collaborate in ransomware and blackmail operations. In one Telegram post, for example, a persona linked to the actor indicated they were looking for partners with access to servers for "blackmail," suggesting they were likely seeking to conduct ransomware and/or extortion operations

Mandiant has also linked TRIPLESTRENGTH to operators in a RCRU64 ransomware incident in May 2024. In that incident, the actors gained initial access via remote desktop protocol, which they accessed via brute-force password guessing. After gaining access, the threat actor moved laterally, disabled the anti-virus, and then executed RCRU64 ransomware on multiple hosts.

Cryptomining Operations

TRIPLESTRENGTH leverages the unMiner application alongside the unMineable mining pool to conduct the mining portion of their operations. The unMiner application packages several popular cryptocurrency

miners into a single application, while the unMinable pool allows payouts in over 90 different cryptocurrencies, regardless of which miner is used.

They have used a variety of both CPU- and GPU-optimized mining algorithms. However, all observed funds were paid out in TRX, the native token for the TRON blockchain. The use of both CPU- and GPU-optimized algorithms indicates the threat actor likely chooses the most optimized algorithm for the target system, with CPU-optimized mining being consistent with mining operations abusing cloud and/or corporate infrastructure.

Targeting Cloud Service Providers Beyond Google Cloud

Google Cloud's security and threat intelligence teams have focused on tracking TRIPLESTRENGTH's activity on Google Cloud. However, it's worth noting that the threat actor targets other cloud providers as well, and their operations may impact organizations across a wide range of industry sectors and regions. Based on analysis of attacker-owned infrastructure, the Google Threat Intelligence Group determined that the actor has relied on Raccoon infostealer logs as the source of at least a portion of the stolen credentials and cookies used in cloud hijacking activities, and that the actor had access to credentials for Google Cloud, Amazon Web Services, and Linode. Additionally, in monitoring Telegram channels, Mandiant has observed personas connected to the group routinely advertise access to servers, including those from prominent hosting providers and cloud platforms such as Google Cloud, Amazon Web Services, Microsoft Azure, Linode, OVHCloud, and Digital Ocean.

Mitigations

We recommend the following risk mitigations to enhance your Google Cloud security posture to protect against threats like account takeover, which could lead to threat actor ransomware or data extortion operations.

Help prevent cloud account takeover:

- **Enroll in multifactor authentication (MFA):** Google Cloud's [phased approach to mandatory MFA](#) makes it harder for attackers to compromise accounts even if they have stolen credentials and/or authentication cookies.
- **Use automated sensitive monitoring and alerting:** The [Sensitive Actions Service](#) within Google Security Command Center (SCC) automatically detects and alerts for potentially damaging actions.
- **Implement robust Identity and Access Management (IAM) policies:** Use IAM policies to grant users only the necessary permissions for their jobs. Google Cloud offers a range of tools, including [Policy Analyzer](#), to help implement and manage IAM policies.

Help mitigate ransomware and extortion risks:

- **Establish a [cloud-specific backup strategy](#):** Disaster recovery testing should include configurations, templates, and full infrastructure redeployment and backups should be immutable for maximum protection.
- **Enable proactive virtual machine scanning:** [Virtual Machine Threat Detection \(VMTD\)](#) is a built-in service in Google SCC that scans virtual machines for malicious applications to detect threats like ransomware.
- **Monitor and control unexpected costs:** With Google Cloud, you can [identify and manage unusual spending patterns](#) across all projects linked to a billing account.

Growing Threat from Data Leak Sites Enabling Extortion in the Cloud

Mandiant has observed threat actors increasingly extorting victim organizations by exposing their stolen data on Data Leak Sites (DLS). This threat actor tactic is alarming because in the last year we have seen this activity impact victim organizations who rely on cloud technologies across multiple cloud service providers, not just those with on-premises systems. The expanded use of these extortion tactics combined with the prevalence of DLS poses a growing threat for all organizations, regardless of where their data is stored.

In Q3 2024, Mandiant observed 1,242 victims posted on DLS with , almost reaching the peak observed in Q3 2023 with 1,329 victims. Although the numbers are slightly lower in Q3 2024, the overall posts to DLS from Jan. through Sept. are higher in 2024 with 3,546 victims posted, than in 2023 with 3,385 during the same time period. We expect the number of posts to DLS in 2024 to surpass the number of posts to DLS in 2023 (4,521 victims) by a few hundred posts.

As more extorted victims and information are posted to DLS, this potentially increases the risk exposure to organizations because of the sensitive information released. Due to the nature of the attack and the use of the information from the DLS, this article provides a

closer look at a threat actor tactics, which were used to compromise cloud environments.

For example, Storm-0501, a threat actor group that is publicly tracked and has possible links to the EMBARGO Ransomware-as-a-Service (RaaS) group, gained initial access to a cloud-based identity provider, deployed ransomware, and amplified their accomplishments on their DLS. Notably, Storm-0501 actively develops and customizes their tools for each victim.

Since April 2024, Mandiant has observed 11 different postings to the EMBARGO DLS, which includes the following activity:

- **Data Exfiltration:** An internal database of a mortgage lender was breached in a major ransomware attack that leaked the personal data of multiple customers on the dark web.
- **Extortion Attempt:** Alleged victims appear on the EMBARGO DLS. Affiliates of EMBARGO can create a victim blog post with the company's name, their logo, a description of the company, a description of the incident (e.g., what and how much was stolen), any screenshots, and a possible link to the data.

Similarities in Threat Actor Tactics for Data Extortion

Two threat groups, UNC3786 and UNC5791, have used similar tactics to compromise organizations that resulted in exposure of their sensitive data on a DLS.

In January 2024, Mandiant observed an attack by UNC3786, a threat group that has been active since at least early 2022. The group primarily gains initial access with compromised credentials and uses MFA bypass techniques (e.g., including SIM swapping) to conduct data extortion.

In September 2024, UNC5791 leveraged a similar tactic to the UNC3786 January 2024 attack to obtain initial access to Microsoft 365 accounts. Mandiant observed the threat actors exfiltrating sensitive information from business software and cloud services organizations, and in some cases, leveraging the ALPHV DLS, which contains data on 53 victim organizations spanning multiple verticals.

In both cases, UNC3786 and UNC5791 employed three notable tactics:

- **MFA Bypass in Cloud-based Services:** Both UNC3786 and UNC5791 gained initial access to Microsoft 365 accounts by SIM swapping targets and exploiting Microsoft's self-service password reset (SSPR) to enroll their own devices for MFA. They then covered their tracks by deleting emails, including password reset notifications.
- **Data Exfiltration in the Cloud:** Both UNC3786 and UNC5791 exploited compromised accounts to access and steal data from victim organizations'

SharePoint servers. UNC3786 downloaded approximately 25GB of data and modified group permissions where the compromised account likely held elevated privileges. Evidence suggests they may have used Veeam Backup for Microsoft 365 for exfiltration, as deleted emails contained Veeam evaluation key and website registration messages.

- **Data Extortion Attempts with Aggressive Tactics:** Both UNC3786 and UNC5791 employed aggressive tactics to attempt to extort victims, including contacting employees and their families. Notably, UNC5791 used aggressive communication tactics to increase pressure on their data theft extortion victims, including calling and texting victims' personal devices.

Mitigations

Organizations can leverage multiple Google Cloud products to enhance protection against ransomware and data theft extortion:

- **Leverage Google Security Command Center (SCC):** [SCC](#) provides a multi-cloud security solution that can help detect data exfiltration events through [Event Threat Detection](#) and misconfigurations through [Security Health Analytics](#). SCC also leverages threat intelligence from Google.
- **Prevent data exfiltration:** [Sensitive Data Protection](#) (SDP) protects potential data exfiltration by identifying and classifying sensitive data within your Google Cloud environment, enabling you to monitor for unauthorized access or movement of data.

- **Incorporate automation and awareness strategies:** Implementing strong password policies, enforcing multifactor authentication (MFA), regularly reviewing user access and monitoring the dark web for leaked credentials associated with your organization's users, implementing account lockout mechanisms, and educating users about security best practices to prevent credential compromise.
- **Enhance security with government insights:** Follow guidance from the U.S. Dept. of Homeland Security, Cybersecurity and Infrastructure Security Agency's Ransomware Vulnerability Warning Pilot ([RVWP](#)), which proactively identifies and warns about vulnerabilities that could be exploited by ransomware actors.

Contributors

Jason Bisson

Charles DeBeck

Elliot Eaton

Nigel Gardner

Cris Brafman Kittner

Dima Lenz

Crystal Lister

Noah McDonald

Matthew McWhirt

Luca Nagy

Zachary Riddle

Alyse Rothman

Josh Stern

UNDER EMBARGO

UNDER EMBARGO

Google Cloud