



Cybersecurity
Action Team

**THIS REPORT IS UNDER EMBARGO
UNTIL THURSDAY, NOVEMBER 2 AT 9AM**

Threat Horizons

Q3 2023 Threat Horizons Report

Table of contents

Mission statement	03
Good Cloud Hygiene is Not a One Time Event	04
Credentials continue to factor into over half of incidents in Q2 2023	05
Threat Actors Adapting Tactics to Target Data on Cloud-hosted Software-as-a-Service (SaaS)	08
Threat Actors potentially abusing Google Calendar to host C2 infrastructure	12
Cybersquatting Abuse Across WHOIS and Cloud Storage Platforms	15
Cloud Industry Review: Healthcare	19



Mission statement

The Google Cloud Threat Horizons Report provides decision-makers with strategic intelligence about threats to cloud enterprise users, along with cloud-specific research. Most importantly, the report delivers recommendations from Google's intelligence and security teams.

Letter from the Editor

Good Cloud Hygiene is Not a One Time Event

As defenders, the most interesting attacks are the advanced ones that make headlines across the industry. For example, in 2022, [Mandiant wrote about a sophisticated attack campaign](#) that leveraged two zero-day vulnerabilities, a novel hypervisor malware, and a new technique for running malicious software on virtual machines.

However, based on the latest data captured in this and previous Threat Horizons Reports, the majority of victims in the cloud are not compromised by these types of advanced attacks. Rather, cloud intrusions are resulting from common and well-known threat actor attack techniques, such as obtaining and using stolen credentials, and from security weaknesses, such as misconfigurations. It may not be as exciting, but by focusing on simple cloud security hygiene, defenders have an opportunity to dramatically reduce the risk of a cloud compromise.

Practicing good cloud hygiene is not a one time event. As your cloud environment matures, it is common for security to drift away from its baseline. Build guardrails into your environment to ensure cloud hygiene is monitored and enforced.

The Threat Horizons Report will continue to highlight advanced threats to the cloud, sophisticated attack campaigns, and novel techniques used to target victims in the cloud. By focusing on good cloud hygiene, defenders will raise the bar necessary for attackers to be successful while reducing the risk of becoming a victim to a common attack. Now let's get into it.


Credentials continue to factor into over half of incidents in Q2 2023

Cloud Compromises Q2 2023

The cloud compromise methods of initial access observed in Q2 2023 were largely similar to previous quarters and consistent with the last 12 months of reporting. This quarter, weak credentials continue to represent the largest compromises where many observed instances were a result of attackers brute forcing default accounts, Secure Shell (SSH), and the Remote Desktop Protocol (RDP). The default Google Cloud Virtual Private Cloud (VPC) network (an auto mode VPC network) has [pre-populated rules](#) that can be useful for early exploration, however custom mode VPC networks are better suited for most production environments. They allow incoming connections from other instances on the same VPC network, SSH, RDP, and also the Internet Control Message Protocol (ICMP). Google Cloud's [OS Login](#) feature helps customers simplify SSH access to instances by using Identity and Access Management (IAM) without having to create and manage SSH keys to help harden remote access. Organizations can also reduce the risk of these services being exposed publicly by disabling the creation of this default network by creating [an organization policy](#) with the `compute.skipDefaultNetworkCreation` constraint and following [best practices for VPC design](#).

Cloud Compromises: Initial Access

- Weak or no password
- Sensitive UI or API exposed
- Vulnerable Software
- Misconfiguration
- Other




In the Q2 2022 Threat Horizons Report, we highlight that a disproportionate percentage of attackers opportunistically use coin mining across Cloud products and alter their tactics to evade discovery. This is consistent with this quarter's findings, as this is the most observed impact from compromises.

(Credentials, cont'd.)

This underscores the importance of enabling [security-related logs](#) along with routine monitoring and alerting. In addition to Google Cloud's cloud-native security suites like Chronicle and Security Command Center, customers could leverage [Community Security Analytics](#) with pre-built queries to detect anomalous behavior and reduce the time to detection, and routinely audit [organization contacts](#) to assure the right stakeholders will receive important notifications.

This complements the observations from our incident response teams and suggests that organizations operating in a multi-cloud environment can expect to see similar threats in their AWS and Azure environments. Organizations can mitigate the risks of databases exposed to the public internet by following security [best practices guides](#) and also considering managed database services that handle the security and scaling for customers.

Cloud Compromises: Impact



This quarter our teams observed an 8.5% increase in vulnerable software compromises led primarily by PostgreSQL being the most exploited. An internet based search filtered on Google, Microsoft, or Amazon shows thousands of cases where SSH and RDP are exposed to the internet along with popular database services such as MySQL and PostgreSQL.

Compromise Initial Access	Previous 12 months average
Weak or no password	48.5%
Sensitive UI or API Exposed	16.5%
Misconfiguration	14.9%
Vulnerable software	11.3%
Leaked Credentials	5.2%
Other	3.1%
Remote Code Execution	0.5%
Trojan	0.0%

Compromise Impact	Previous 12 months average
Coin Mining	70.1%
Intrusion Attempt	24.8%
Other	2.9%
Account Leaked Credentials	2.2%
DOS	0.0%


Statistics are based on observations by our Google Cloud incident response teams, which will be skewed to the platforms and sample size and might not be representative of all customer environments and verticals on Google Cloud, but should be representative of general trends.

(Credentials, cont'd.)

Chronicle Security Operations alerting trends

In Q2, alerts from Chronicle Security Operations – Google Cloud’s modern, cloud-native SecOps platform – complimented the data above. Nearly 65% of alerts across organizations were related to risky use of service accounts. These accounts have associated permissions where if compromised, could lead to attackers gaining persistence and subsequently using this access for privilege escalation in cloud environments. In the [Q1 2023 Threat Horizons Report](#) we highlighted in more detail how attackers could abuse service account keys and included several mitigations techniques such as [evaluating alternative authentication methods](#).

Chronicle Security Operations alert trends



Threat Actors Adapting Tactics to Target Data on Cloud-hosted Software-as-a-Service (SaaS)

Our research and analysis team has observed persistent threat actor activity targeting data stored on cloud-hosted Software-as-a-Service (SaaS) systems across multiple industries, at times with new attack methods. This article explores the implications of emerging attack methods targeting cloud-hosted SaaS systems and provides risk mitigations to help prevent attacks.

Increasing Use of Cloud-hosted SaaS Expands Attack Surface

As organizations increasingly adopt cloud-hosted Software-as-a-Service (SaaS), they have additional attack surface to manage and secure because sensitive data is distributed to more locations to conduct critical business functions. [The 2023 Thales Cloud Security Study](#) indicates a 41% increase in the mean number of SaaS applications used by survey respondents from 2021 to 2023. Additionally, cybersecurity industry reporting indicates that on [average](#), in a 10,000 SaaS-user organization, there are approximately over 4,000 applications connected to both Microsoft 365 and Google Workspace.

Looking at cloud-hosted SaaS environment security incidents in the last two years, over 55% of security executives have experienced cases involving data breaches and leaks, malicious applications,

ransomware, corporate espionage, or insider attacks, according to the [Cloud Security Alliance](#). Some [industry reporting](#) indicates that large enterprises are not fully prepared to protect their SaaS data from ransomware attacks, even though SaaS data lost in successful ransomware attacks was the least likely type of data to be fully recovered.

Threat Actors Exploiting More than One SaaS System, Leveraging Diverse Methods to Conduct Intrusions

Over the last nine months, we observed threat actors leveraging diverse tactics to gain access to, and exfiltrate data from, cloud-hosted SaaS systems.

Multi-SaaS Cloud Exploitation: In addition to conducting intrusions by exploiting one cloud-hosted SaaS system, some threat actors are conducting intrusions by exploiting more than one system at a time. During Q1 2023, Mandiant Intelligence observed several incidents in which the majority of a cloud-based intrusion occurred within one or more SaaS system (See Fig. 1).

(Threat Actors Adapting Tactics to Target Data on Cloud-hosted Software-as-a-Service (SaaS), cont'd.)




Figure 1: Model of a Multi-Software-as-a-Service Cloud Exploitation Attack

Unauthorized Access: In Q1 2023, Mandiant responded to an incident in which a threat actor accessed a victim's SaaS system used for customer engagement. The attacker then generated reports containing customer data that were automatically uploaded to the victim's Amazon Web Services (AWS) environment. The threat actor, whose motivations are unknown, subsequently downloaded some of the reports via links sent to the victim's compromised email account.

In a separate incident in May 2023, security [researchers](#) observed the Omega ransomware group conduct an attack using an administrator account instead of a compromised endpoint to infiltrate an unnamed company's environment, elevate permissions, and exfiltrate sensitive data from SharePoint libraries.

Supply Chain Compromise: Mandiant Consulting responded to a compromise in July 2023 attributed to Democratic People's Republic of Korea (DPRK) actors that leveraged a SaaS provider, Jumpcloud, to conduct a [targeted supply chain attack](#). SaaS providers were also targeted earlier in the year by suspected financially-motivated DPRK actors in order to gain access to [downstream victims](#).

Malicious OAuth Application: Security researchers discovered a vulnerability within Microsoft's OAuth application registration in February 2023 that would allow a new [attack](#) using Exchange's legacy API to create hidden forwarding rules in Microsoft 365 mailboxes (See Fig. 2). [The Q2 2023 Google Threat Horizons Report](#) provides more information on OAuth application risks and mitigations.

(Threat Actors Adapting Tactics to Target Data on Cloud-hosted Software-as-a-Service (SaaS), cont'd.)




Figure 2: Model of an Attack on Email Data Stored on Cloud-hosted Software-as-a-Service

Vulnerability Exploitation: Mandiant observed financially-motivated threat group FIN11 conduct two notable campaigns [exploiting](#) zero-day vulnerabilities in managed file transfer software from [Fortra](#) and [MOVEit](#) to conduct data theft extortion operations affecting hundreds of organizations in the first half of 2023.

(Threat Actors Adapting Tactics to Target Data on Cloud-hosted Software-as-a-Service (SaaS), cont'd.)

Mitigations

Opportunities and resources to help defenders strengthen security for cloud-hosted SaaS systems include:

Use SaaS Security Configuration Guides: Investigate how to leverage native security controls embedded into each SaaS application and configure them according to industry best practices and standards, such as the principles of least privilege and multifactor authentication (MFA). Some resources include the Cloud Security Alliance's [SaaS Governance Best Practices](#) for Cloud Customers, the National Institute of Standards and Technology's [Guide to a Secure Enterprise Network Landscape](#), and the UK National Cyber Security Centre's [Cloud Security Guidance](#).

Use a SaaS Security Posture Management (SSPM)

Tool: A SSPM will automate the protection of SaaS applications and can help identify misconfigurations, unused user accounts, unnecessary user permissions and other cloud security risks.

Incorporate Robust Identity and Access

Management (IAM): For identities and permissions, closely manage accounts with high privilege and administrator access and apply [least privilege](#) principles to ensure each user has the minimum required permissions. Ensure SaaS account management in identity lifecycle processes, including deactivating unused user accounts and revoking data shares after a certain period of time. If granting a SaaS application access to resources in your Google Cloud environment, tools such as Google Cloud's [Policy Analyzer](#) can help confirm the minimum IAM role necessary.

Limit User Access to SaaS Systems: Implement a Cloud Access Security Broker (CASB) to broker the connection between users and SaaS. This provides the ability to monitor SaaS usage, control or terminate SaaS access appropriately and instantly, and apply appropriate governance policies around SaaS.

Threat Actors potentially abusing Google Calendar to host C2 infrastructure

Rather than rely on infrastructure they operate themselves or lease from criminal services, many threat actors increasingly favor legitimate cloud services to host their infrastructure. Public cloud services provide cheap, reliable infrastructure trusted by enterprises and consumers, allowing threat actor activity to evade detection by blending into high volumes of legitimate traffic.

All cloud vendors and their products are affected by this type of abuse. Google's Threat Analysis Group (TAG) tracks and regularly disrupts serious cyber threat actors and malware abusing legitimate cloud services. These services range from cloud-based storage and compute services to workplace productivity services like email and calendar tools.

Threat actors have abused cloud-based storage to host campaign infrastructure, to deliver malware, to act as malware command and control (C2), and to upload exfiltrated data. This trend goes back several years - in 2021 Cisco Talos [reported](#) on threat actors using Microsoft Azure and Amazon Web Service to deploy and deliver variants of malware with information stealing capabilities. In April 2023, TAG detected a People's Republic of China (PRC)-backed actor using malicious PowerShell scripts that communicated with Dropbox to retrieve commands

and exfiltrate data. TAG has also observed threat actors pivot between providers in response to policy changes and disruption efforts.


In June 2023, an independent developer [published proof of concept code](#) to Github for "Google Calendar RAT (GCR)." At this time, TAG has not observed the use of GCR in the wild.

The red teaming tool uses Google Calendar events for C2. The tool enables an attacker to place commands in the event description field of Google Calendar events.

While we have not seen the use of GCR in the wild to date, Mandiant has noted multiple actors sharing the public proof of concept on underground forums, illustrating the ongoing interest in abusing cloud services.

GCR, running on a compromised machine, periodically polls the Calendar event description for new commands, executes those commands on the target device, and then updates the event description with command output. According to the developer, GCR communicates exclusively via legitimate infrastructure operated by Google, making it difficult for defenders to detect suspicious activity.

(Threat Actors potentially abusing Google Calendar to host C2 infrastructure, cont'd.)



Google Calendar RAT attack flow diagram, published by the developer on Github

A demo video posted to Github shows the developer using two specific accounts to operate GCR. The first is a GCP service account with a corresponding API key for beaconing, and the second is a Gmail “attacker account” used to manually interact with Calendar.

TAG has previously observed threat actors abusing Google products in their campaigns. In March 2023, TAG observed an Iranian government backed actor use macro docs to infect users with a small .NET backdoor, BANANAMAIL, for Windows that uses email for C2. The backdoor uses IMAP to connect to an attacker-controlled webmail account where it parses emails for commands, executes them, and sends back an email containing the results. TAG identified and disabled attacker-controlled Gmail accounts that the malware was using as a C2 mechanism.

(Threat Actors potentially abusing Google Calendar to host C2 infrastructure, cont'd.)

Mitigations

- **Architect systems with a defense-in-depth approach** to reduce risk if threat actors bypass controls by evading detection such as when using valid cloud services as noted above.
- **Use an Intrusion Detection System (IDS) and network monitoring tools** to detect application level or network level C2 traffic or even exfiltration with tools such as [Cloud IDS](#) or open source alternatives [Suricata](#) in conjunction with [Zeek](#).
 - » **Segment networks** to reduce the impact of adversaries gaining access to additional resources in your environment. Consider Google Cloud's [best practices and reference architectures for VPC design](#).
 - » **Develop baselines for network traffic** and monitor for connections to user facing cloud services to aid defenders in identifying low prevalence and/or anomalous behavior.
- **Implement robust centralized logging** and regularly monitor your environment for anomalous behavior. The Q3 2023 compromise metrics section outlines security-related logs organizations should consider enabling in their Cloud environment along with a link to Community Security Analytics with example queries and YARA rules organizations could use for detections.

Cybersquatting Abuse Across WHOIS and Cloud Storage Platforms

Cybersquatting – the practice of registering domain names in violation of trademark rights – has surged significantly over the past ten years. This is because domains are inexpensive and give threat actors a large return on investment.

Recently, threat actors have evolved their tactics to include typosquatting – a form of cybersquatting – attacks on cloud storage platforms such as Google Cloud Storage, Amazon S3, and Azure Blob.

Typosquatting is the most common form of cybersquatting. It relies on mistakes (such as typos) made by internet users when entering a website address into a browser. Should a user accidentally enter an incorrect address, they would be directed to an alternative website owned by a cybersquatter. The following are examples of typosquatted domains:

Omission	Insertion	Substitution	Transposition	Hyphenation
"oogle.com"	"googgle.com"	"goog1e.com"	"googel.com"	"g-oogle.com"
The first "g" is omitted	An extra "g" is inserted	The letter "l" is substituted with the number "1"	The letters "e" and "l" are transposed	An arbitrary hyphen is inserted

Threat actors use typosquatted domains to accomplish a range of objectives. Typosquatting can be leveraged in phishing attacks, and when users visit the site, they may encounter phishing pages impersonating legitimate services, prompting them to disclose sensitive information such as login credentials, credit card details, or personal data. Typosquatting can also be a means to distribute malware. By luring unsuspecting users to a typosquatted domain, threat actors can infect their systems with malicious software such as ransomware, spyware, or Trojans. Overall, typosquatting is a subtle yet powerful technique used by threat actors to exploit human fallibility and carry out a wide range of cybercrimes.

Cybersquatting can also be used to launch identity theft attacks. For example, threat actors could register a company's literal name and attempt to operate as the organization itself. In the United States, personal names can be trademarked, and domains representing individuals that have marketplace significance can lead to reputational damage.

(Cybersquatting Abuse Across WHOIS and Cloud Storage Platforms, cont'd.)

Typosquatting Across Cloud Platforms

Threat actors have been observed using typosquatting techniques to abuse cloud storage names across Amazon S3 Buckets, Google Cloud Storage, and Azure Blob. Although cloud storage names are globally unique across cloud providers, they are not bound specifically to an organization. As a result, if a bucket name is available, threat actors could forge a company's name. And if the precise spelling of a company's name is not available, the threat actor could resort to a typosquatting attack similar the following:

Amazon S3 Bucket	Google Cloud Storage	Azure Blob
https://gogle.com.s3.amazonnews.com	https://g00gle.storage.googleapis.com	https://googel.blob.core.windows.net
Omission	Substitution	Transportation

In this example, a typosquatting attack on "google" was performed within the cloud storage URL. These URLs may be used as staging areas to harbor malware, phishing attacks, and other nefarious activity. If a typosquatted cloud storage name closely resembles a company's name, there is a chance malicious activity could go unnoticed by the Security Operations Center (SOC) team. Below is a comparison of legitimate and typosquatted log telemetry:


<— Outbound connection 1 —> <https://google.storage.googleapis.com>

<— Outbound connection 2 —> <https://google.storage.googleapis.com>


Would you catch it? Outbound connection 1 is typosquatted where a capital "i" is used in place of a lowercase "l" in "google". Outbound connection 1 could easily be classified as benign by the SOC – due to it being virtually indistinguishable from outbound connection 2.

(Cybersquatting Abuse Across WHOIS and Cloud Storage Platforms, cont'd.)

Cybersquatting continues to plague organizations. A random sample of ten Fortune 100 companies found that 100% of the organizations had one or more typosquatted domains. Additionally, 60% of the sampled organizations had one or more typosquatted cloud storage URLs:



Research by [Orca Security](#) (2023) found that malware is approximately three times more prevalent in cloud storage buckets than Virtual Machines. Recently, a threat actor identified a once active, yet abandoned, cloud storage bucket and seized it to launch malicious payloads. [Keeper Security](#) (2023) discovered an NPM package named “bignum” contained a component used for downloading binary files hosted on an Amazon S3 bucket. Users who downloaded the package also downloaded malicious binaries, which were used to steal user IDs, passwords, and perform data exfiltration. Ultimately, the threat actor discovered the bucket name was still in use within the package and cybersquatted an Amazon S3 Bucket name to successfully launch the attack:



(Cybersquatting Abuse Across WHOIS and Cloud Storage Platforms, cont'd.)

Mitigations

There are a range of options to tackle cybersquatting attacks:

- It is important to understand the degree to which a domain has been typosquatted as well as monitor for newly registered domain permutations. This can be accomplished using a third-party domain monitoring service.
- Develop a playbook to address typosquatted domains as well as cloud storage cybersquatting abuse.
- Directly report cloud storage cybersquatting to each associated cloud provider ([Google](#), [Amazon](#), and [Microsoft](#)).
- Proactively registering domain permutations can significantly reduce the risk of typosquatting attacks.
- Attack Surface Management efforts should include domain and cloud storage typosquatting checks.
- Internationally, victims can file their complaints under the [Uniform Domain-Name Dispute-Resolution Policy \(UDRP\)](#) with the World Intellectual Property Organization (WIPO). The United States has adopted the [U.S. Anticybersquatting Consumer Protection Act \(ACPA\)](#) to provide protection against cybersquatting for individuals as well as corporate owners of distinctive trademarked names.

Cloud Industry Review: Healthcare

This is the second article in a new series from our threat analysis team exploring the implications of cloud services adoption and security concerns across various industries. Healthcare organizations should utilize these data-driven insights and actionable cloud security risk management recommendations to enhance their defensive posture against threats.

Introduction

The healthcare industry is a well known target for cybersecurity attacks. Hackers take advantage of the industry's possession of significant amounts of personally identifiable information (PII) and protected health information (PHI), and healthcare entities' various legal and ethical obligations to protect that data.

Reviewing cybersecurity incident summaries gathered from Mandiant and Google and research from various public sources, for the period 2021-2023, we see that the cloud is both an attack target, and is increasingly used by threat actors as a platform to stage attacks on victims. Examining the data more deeply, we found that our collected set of incidents had operational impacts on healthcare entities, including hospitals, national health systems, and similar institutions.

The attacks themselves weren't new. Ransomware was often used, victim extortion continued, and credential abuse was a common initial infection vector. But what was especially concerning from the data we gathered, is that such attacks were beginning to negatively affect patient safety and life-saving medical care – thereby giving leverage to attackers.

Healthcare organizations adopt cloud computing to improve business efficiencies, automate clinical processes, and enhance patient outcomes. As per a 500-person 2022-2023 survey by DuploCloud, seven out of 10 healthcare IT professionals indicate that their organization already uses the cloud; and an additional 20% of survey respondents hope to use the cloud within the next two years.¹ Healthcare organizations use cloud technologies to rapidly analyze large volumes of health records, extend care delivery into people's homes via cloud-powered telehealth, prototype new healthcare products and solutions, and handle other use cases.^{2,3}

Yet the cloud, like other IT environments, also presents various security risks to healthcare organizations. First, cloud-hosted healthcare organizations are targeted by attackers. According to a 2022 study, 63% of healthcare organizations have suffered at least one account compromise in their cloud platforms previously.⁴ Non-cloud environments have had analogous security and network visibility concerns. Ninety percent of participants in a 2023 Gigamon survey experienced a data breach in their on-premise and cloud hybrid infrastructure within the past 18 months.⁵ And only roughly 33% of the survey participants had visibility from the network to the application level within these hybrid infrastructures.

More disturbingly, cybersecurity attacks on cloud-hosted healthcare entities can impact patient care (as in other IT environments). A 2022 analysis by Netwrix found that 61% of healthcare organizations experienced an attack on their cloud infrastructure in

(Cloud Industry Review: Healthcare, cont'd.)

the past 12 months, often via ransomware.⁶ And the Cybersecurity & Infrastructure Security Agency (CISA), analyzing its own 2020–2021 healthcare industry data, determined that such successful ransomware attacks on hospitals in particular degraded their operational capacity. The attacks reduced the number of beds staff could attend to, potentially diverted critical care patients to hospitals farther away, and had a number of other harmful health effects.⁷

At the same time, cloud services, besides being a target, are also leveraged by malicious actors to attack healthcare organizations. A Netskope March 2023 study showed that the healthcare vertical, when compared to other industries, was receiving an increasing share of malware delivered from popular cloud applications like OneDrive and AWS S3.⁸ Organizations and users appeared to place additional “trust” in such known applications and data repositories.

¹DuploCloud. “70% of Healthcare Businesses Have Adopted Cloud Computing: DuploCloud Report”, GlobeNewswire, 23 February 2023, <https://www.globenewswire.com/news-release/2023/02/22/2613339/0/en/70-of-Healthcare-Businesses-Have-Adopted-Cloud-Computing-DuploCloud-Report.html>. Accessed 3 August 2023.

²Schnittfink, Theo. “How Technology Puts The ‘Care’ In Healthcare: The Role Of The Cloud During The Pandemic”, Forbes, 10 May 2022, <https://www.forbes.com/sites/forbestechcouncil/2022/05/10/how-technology-puts-the-care-in-healthcare-the-role-of-the-cloud-during-the-pandemic/?sh=c8b4d7011137>. Accessed 8 August 2023.

³Security Boulevard. “Cloud Computing The Prescription for Modern Healthcare Challenges”, Security Boulevard, 22 May 2023, <https://securityboulevard.com/2023/05/cloud-computing-the-prescription-for-modern-healthcare-challenges/>. Accessed 8 August 2023.

⁴Ponemon Institute and Proofpoint. “Cyber Insecurity in Healthcare: The Cost and Impact on Patient Safety and Care”, Proofpoint, March 2022, <https://www.proofpoint.com/sites/default/files/threat-reports/pfpt-us-tr-cyber-insecurity-healthcare-ponemon-report.pdf>. Accessed 5 August 2023.

⁵Gigamon. “2023 Hybrid Cloud Security Survey”, Gigamon, 2023, <https://www.gigamon.com/content/dam/gated/wp-gigamon-survey-hybrid-cloud-security-2023.pdf>. Accessed 13 September 2023.

⁶Netwrix. “2022 Cloud Data Security Report”, Netwrix, March 2022, https://www.netwrix.com/download/collaterals/Netwrix_Cloud_Data_Security_Report_2022.pdf. Accessed 8 August 2023.

⁷Cybersecurity & Infrastructure Security Agency. “Provide Medical Care is in Critical Condition: Analysis and Stakeholder Decision Support to Minimize Further Harm”, Cybersecurity & Infrastructure Security Agency, September 2021, https://www.cisa.gov/sites/default/files/publications/CISA_Insight_Provide_Medical_Care_Sep2021.pdf. Accessed 7 August 2023.

⁸Netskope. “Netskope Threat Labs Report Healthcare”, Netskope, March 2023, <https://www.netskope.com/wp-content/uploads/2023/03/threat-labs-report-healthcare-march-2023.pdf>. Accessed 7 August 2023.

(Cloud Industry Review: Healthcare, cont'd.)

Analyzing Google and Mandiant Cloud Security Incidents

To provide guidance to healthcare organizations on securing their cloud environments, we looked closer at 2021-2023 private incident response data for various healthcare incidents, as well as overall attack campaigns handled by Mandiant and Google. This included Mandiant's additional analyses of various public healthcare incidents and campaigns. From this corpus, we selected a data set, so as to better examine and recommend customer cloud controls. Our data set included incidents or campaigns where at least one healthcare cloud-hosted system or storage repository was attacked; or at least one but often many incidents (as in a campaign) leveraged some part of cloud infrastructure to attack healthcare institutions.

First, we can summarize our data set—to observe overall trends. The table below aggregates overall attacker strategies, from our data, showing threat actors' motivations and methods for attacking healthcare institutions.

Attacker motivations for compromising healthcare organizations	Geographies of targeted organizations	Initial attack vectors used in the compromises	Follow on compromises (post initial entry)
Most common: - Financial gain Infrequent: - Espionage - Security researchers uncovering vulnerabilities (e.g. non-harmful "attacker" motivations) - Hacktivist (political)	Most common: - North America - Asia - Western Europe Less common: - South America - Eastern Europe - Australia - Africa - Middle East	Most popular: - Stolen credentials Less popular: - Phishing - Third-party vulnerabilities - Denial of Service attacks - Web exploits - Misconfigurations	Frequently used: - Ransomware - Data extortion Infrequent: - Selling captured data - Destroying data and infrastructure - Sharing initial infection vectors with others (i.e. becoming an Initial Access Broker (IAB—which specializes in providing initial infection vectors to purchasing stakeholders), to other groups)

We can also summarize attackers' interactions with the cloud, as per our data, to better understand the role that cloud plays in our collection of attacks. Attackers frequently targeted cloud-hosted healthcare organizations by trying to capture their sensitive business data and PHI. When the cloud facilitated attacks on healthcare institutions, it was used to host attacker infrastructure, and to hide and ascribe trust to attacker assets. The attackers' aggregated interactions with the cloud, per our data, are shown below.

(Cloud Industry Review: Healthcare, cont'd.)

Attacker goals when compromising cloud-hosted healthcare organizations or processes	Attacker aims when cloud was facilitator of attacks
Extracting credentials and data from the cloud—often by targeting: - Outlook Web Access application - AWS resources like S3	- Cloud used to host malicious files, and is a source of malicious downloads - Cloud stores data stolen by attacker - Cloud used to generate Denial of Service traffic against targets - Cloud platforms provide “trusted brands” for attacker domains to impersonate - Cloud used as trusted ‘front’ to mask attacker IP addresses

Threat Actor and Attack Details

The tables above summarize the breadth of attacker motivations and tactics. Below we describe unique attacker motivations and/or the uncommon TTPs they use to add considerable detail to such summaries. The diversity helps to better assess and address risks in the cloud.

MDS. Threat actor MDS has been the biggest seller of stolen data on the Chinese dark web marketplace DeepWebChinese since May 2021. A number of the data sets are healthcare-specific, and they often contain demographic information, phone numbers, names, and associated data. Once purchased, data sets are downloaded from a cloud service. On MDS’ web store, certain data sets are labeled “real time”, suggesting the actor maintains live data access—and can support real-time requests.

EXOTIC LILY. EXOTIC LILY is an Initial Access Broker (IAB) that sells initial infection vectors to different purchasers.⁹ Establishing an initial foothold within organizations, EXOTIC LILY sells the access to other attackers. Through November 2021, Google’s Threat Analysis Group observed EXOTIC LILY targeting certain industries, including healthcare; in 2022, the actor further expanded its attacks to other industries. Much of the actor’s work involves social engineering to build trust with victim organizations. To create required “legitimate”-looking identities for its compromises, EXOTIC LILY would create false online profiles, sometimes including synthetic profile photos created using an AI service.

Utilizing spoofed identities, EXOTIC LILY would email victim users. Building rapport, the actor would upload a malicious payload to a well-known file sharing or cloud service, such as TransferNow or OneDrive. At an appropriate time, the actor would send the payload to victim users—for detonation later. The actor would send the users a download notification message from the file sharing or cloud service, to entice the download via a “trusted” source.

⁹Stolyarov, Vlad, and Benoit Sevens. “Exposing initial access broker with ties to Conti”, Google Threat Analysis Group, 17 March 2022. <https://blog.google/threat-analysis-group/exposing-initial-access-broker-ties-conti/>. Accessed 8 August 2023.

(Cloud Industry Review: Healthcare, cont'd.)

Ransomware Operations and UNC2190. In our data set, ransomware affiliate programs—wherein the ransomware program operator provides ransomware software, and recruits customers (e.g. affiliates) to deploy it—were frequently responsible for attacks on healthcare institutions. There was usually an agreement between operator and affiliate on how to divide the required ransomware attack workflows and tasks—including the apportionment of payment for any paid ransoms.

In October 2021, Mandiant observed financially motivated ransomware program operator UNC2190 initiate new ransomware operations. UNC2190 had been observed targeting various critical infrastructure institutions, including healthcare entities, since June 2021. Through its operations, UNC2190 stole large volumes of victim data and deleted many data backups. To protect its IP addresses, UNC2190 also masked them using a cloud service.

FIN12 and UNC2727. FIN12 was a financially motivated threat group that was active in 2021 and several years prior, targeting healthcare entities in about 20% of its attacks. FIN12 frequently obtained initial access via UNC2053, a threat group that often hosted their malicious payloads on cloud-based platforms. By 2021, many ransomware actors had shifted to multifaceted extortion, incorporating additional extortion methods—such as stealing data and threatening to publicly release it—into their operations. Yet FIN12 continued to rely solely on ransomware deployment in most intrusions. While multifaceted extortion increases the pressure on the victim, it can also be time consuming to identify and steal relevant data. By focusing on ransomware deployment without data theft, FIN12 prioritized speed. Ransomware attack efficiency

is often measured as ‘time-to-ransom’ (TTR)—the time between an attacker’s initial compromise and ransomware deployment. The average TTR across all 2021 Mandiant-handled ransomware incidents was 7 days, while FIN12’s 2021 TTR was less than 2 days.

Mandiant assesses with high confidence that threat actor UNC2727 includes one or more members formerly related to FIN12. UNC2727 conducted extortion operations from at least early 2021 to late 2022, with suspected activity also in 2023. Analogous to FIN12, UNC2727 has disproportionately impacted healthcare organizations. In its attacks, UNC2727 has also abused cloud assets, including uploading stolen data to a cloud-based file-sharing service. Unlike FIN12, UNC2727 performed data theft prior to encrypting victims’ systems. While the use of data theft increased UNC2727’s TTR compared to FIN12, it has likely allowed them to apply additional pressure to their victims, thus making ransom payments more likely. Ransomware operators — as other attackers — choose their attack tactics, mindful of the risks and benefits of the impact they wish to have.

Anonymous Sudan. The hacktivist group Anonymous Sudan appeared to begin operations in January 2023, and has been active at least through August 2023. It launches Denial of Service (DOS) attacks against entities providing important infrastructure and services within countries, including healthcare organizations. Anonymous Sudan’s attacks are periodically motivated by political, Muslim-associated events occurring in a country. While the actor has successfully attacked systems with robust security, such as OneDrive and Outlook, it prefers to attack smaller organizations which have less capability to defend their operations. To perform DOS attacks,

(Cloud Industry Review: Healthcare, cont'd.)

the actor uses a set of coordinated cloud servers to generate the substantial traffic against its victims, and open proxy infrastructure to conceal attack sources.

Below are shorter summaries from our data set of other cloud-associated healthcare attackers—that have some unique TTPs.

UNC3810. Espionage-motivated threat actor UNC3810 has conducted campaigns by actively trying to harvest credentials from various organizations, including pharmaceutical companies. The actor had added credential theft code to victim organizations' cloud-based Outlook Web Access login pages, sending the captured credentials to a C2 server.

UNC3774 and UNC4017. Financially motivated threat actors UNC3774 and UNC4017 have opportunistically targeted various industries, and have attacked healthcare organizations. They distributed information stealers which looked like legitimate applications or pirated software, that were popular with users. The information stealers were hosted in cloud platforms like Discord, Telegram, and Github. The stealers would target data from multiple web browsers, cryptocurrency wallets, chat programs, and other applications.

APT22. Espionage-motivated threat actor APT22 has conducted attacks against various industries, including healthcare. It has used malicious domains impersonating cloud services from Amazon Web Services, Alibaba Cloud, and other CSPs in its attacks.

UNC961. UNC961 is a financially-motivated actor that has targeted various industries, including healthcare. It has exfiltrated data to the cloud in at least one prior healthcare compromise. UNC961 attacks using

a variety of custom malware, although it also uses publicly-available malware. UNC961 has also been an IAB for at least two other groups and their follow-on compromises.

Mitigations

Reviewing attacker operations above—healthcare (and other) customers can mitigate such risks, by implementing the recommendations below.

- 1. Mitigate against cloud service credentials and session abuse, to protect against the TTPs used by MDS, UNC3810, and analogous actors.** Follow robust Google Cloud IAM [authentication](#) practices, including using [Policy Intelligence tools](#)—to ensure least privilege credentials are defined, set up, and monitored for the cloud services utilized. MFA should also be set up for key resources, to ensure that stolen credentials, by themselves, don't lead to proper, sufficient authentication. This includes setting up MFA for [important Google Cloud-hosted web applications](#); and setting up [2-Step Verification](#) for the Workspace administrator account and for key Workspace users. Workspace administrators should also implement [appropriate session expiration](#) for key Google Cloud services—to help mitigate threats like “real time” data access, as utilized by actor MDS or other actors.
- 2. Reduce data exfiltration and extortion risk, to protect against the TTPs used by UNC2727 and other actors.** Follow a variety of [prescribed Google Cloud strategies](#) to reduce data theft from cloud services. These include using Digital Rights Management on sensitive cloud files to prevent remote data use via file-embedded controls, and monitoring and alerting security administrators when the data read rate from storage devices

(Cloud Industry Review: Healthcare, cont'd.)

exceeds a particular threshold. Strategies also include making permissions to critical data temporary and subject to frequent reviews—such as by using [App Engine](#) quotas to enable different permission “limitations” and use thresholds.

3. Mitigate effects of ransomware and data destruction, to protect against the TTPs of UNC2190, FIN12, UNC2727, and similar actors.

Maintain appropriate disaster recovery capabilities, to successfully recover from ransomware, data deletion, and related attacks. Keep backup data [isolated](#), such as in Google Cloud zones away from the production zone. For critical data, maintain offline backups. Periodically test system resiliency by performing “whiteboard” or live business continuity tests to ensure infrastructure destruction and similar attacks, don’t affect production work.

4. Mitigate web exploits and third party vulnerabilities as used by a variety of threat actors. Run regular vulnerability scans against cloud instances and perform penetration testing against key cloud-hosted web applications. Patch any identified vulnerabilities in native services, third party software, and web apps in a timely fashion.

5. Mitigate cloud instance misconfigurations as exploited by a variety of threat actors. Ensure that the [Security Health Analytics \(SHA\) scanner](#) within the Security Command Center is turned on. The scanner probes a variety of configuration parameters in Google Cloud instances at a pre-set, documented cadence. The scanner identifies misconfigurations in containers, IAM settings, and other artifacts.

6. Protect against Denial of Service attacks to mitigate against Anonymous Sudan and analogous threat actors. Use [Cloud Armor](#)

to protect network load balancers, protocol forwarding, or VMs with public IP addresses in a cloud instance against network or protocol-based volumetric attacks. Cloud Armor allows only well-formed requests to flow through an instance’s load balancers—so that internal VPC resources are not affected by the (substantial) inbound traffic.

7. Mitigate against malicious files entering the environment, including TTPs from EXOTIC LILY, UNC961, UNC3774, UNC4017, and many other threat actors. Enable endpoint protection (e.g., anti-malware and endpoint detection and response (EDR) tools) on an instance’s Internet-facing hosts, to prevent malicious files downloading into the instance.

8. Mitigate social engineering attacks, to protect against TTPs from EXOTIC LILY and various similar threat actors. Workspace administrators can turn on [Advanced Phishing and Malware Protection capabilities](#) to protect against suspicious email attachments, links, and images; as well as the spoofing of employee names, emails pretending to come from one’s organization, and similar threats. For Gmail, Google Account, and all other email systems, [train employees to be vigilant against phishing attacks](#). Create incident response processes to manage credential capture or abuse incidents on key production systems.

(Cloud Industry Review: Healthcare, cont'd.)

9. Restrict interactions with malicious sites, including those hidden by IP masking or domain impersonation—and mitigate threats from dark web data sales—which protects against the TTPs from APT22, UNC2190, MDS, and similar threat actors. Leverage Chrome with its [Safe Browsing](#) protections to guard against malicious sites. Use threat intelligence (TI) platforms to help identify malicious sites behind masked IPs and impersonated domains—feeding them into Security Web Gateways to considerably reduce interaction with them. TI services can also help to identify stolen assets being mentioned or re-sold in dark web forums. If found, the risk to the organization can be assessed—and mitigations undertaken, as required.

Google Cloud