



Cybersecurity
Action Team



Threat Horizons

September 2022 Threat Horizons Report

Google Cloud

September 2022
For more information, visit gcat.google.com

Table of contents

Mission statement	03
Letter from the editor	
Strategic perspective: The impact of intelligence sharing on adversary operational planning	04
Summary	
View from the frontlines: Threat actors use weak or compromised creds to target SSH, WordPress in Q2	06
Threat trends	
Social engineering campaign disrupted and moved off of Google	08
Malicious files and URLs slipping by IT governance controls	09
Cloud and SaaS-enabled environments increasingly vulnerable to the connectedness of SSO	14
Cloud application security reviews and data discovery	17
Hijack of cloud infrastructure service accounts set to rise	22
Responding to the next SolarWinds: Logging tools	26



Mission statement

The Google Cloud Threat Horizons Report brings decision-makers strategic intelligence on current and likely future threats to cloud enterprise users and the best original cloud-relevant research and security recommendations from across Google's intelligence and security teams.

Letter from the editor

Strategic perspective: The impact of intelligence sharing on adversary operational planning

When you face a compromise from an Advanced Persistent Threat (APT) group, it can feel overwhelming facing off alone against hackers working for a major military or intelligence agency or organized crime group. Moreover, those threat groups are often multi-disciplinary, with the resources to recruit insiders, conduct sophisticated open source research, and execute social engineering activities, using teams of developers creating or adapting their own exploits and tool sets, and ready to exploit any gains with a robust propaganda apparatus. They bring all the scale of a nation-state to bear on a single, outmatched victim.

But that very sense of scale also gives network defenders certain advantages, especially in cloud deployments.

Government-backed hackers often attempt to collect intelligence or prepare for attacks against a wide variety of networks to meet a single requirement common to all their victims. For example, multiple government ministries in several different countries might be hacked, with the common underlying goal of obtaining intelligence on foreign diplomatic efforts or military activities in the region. Private companies of various sizes and locations might all be compromised because of their importance to an entire sector to enhance the economic competitiveness of the country sponsoring the operation. Personally identifying information (PII)

from hotels, healthcare, and financial institutions could all be stolen in order to support counterintelligence investigations or the targeting of dissidents.

Each victim organization feels specifically targeted, but from the government sponsor's point of view, these are all just single points of information gathering in a much larger collection plan that probably involves other tools as well. To meet their government mission, state-sponsored malicious cyber actors often need to be ready to collect important information on short notice, and they want to be pre-positioned on as many victims' networks as possible to ensure the loss of access to any one network doesn't prevent them from accomplishing that mission.

Personally identifying information (PII) from hotels, healthcare, and financial institutions could all be stolen in order to support counterintelligence investigations or the targeting of dissidents.

(Letter from the editor, cont'd.)

In many cases, these networks will be targeted with a common tool set, with operations launched from some overlapping infrastructure. They might even have developed their own 0-day exploits in an attempt to gain access to as many critical targets as possible before victims patch their systems. As we have seen in recent years with log4j and Microsoft Exchange vulnerabilities, state-sponsored actors are also skilled at rapidly adapting new vulnerabilities discovered by others for their own purposes.

What this means for network defenders is simple: Because these adversaries are attempting to operate at scale and as part of a larger intelligence apparatus, threat-intelligence sharing and patching impairs the scale at which even the most sophisticated threat groups can gain access with their latest tools. Attackers might feel confident that they can gain access to any particular victim with a new exploit, but the very power and usefulness of that exploit means they also face high opportunity costs: if their first targets rapidly warn other would-be targets, who then inoculate themselves, those future operations won't be as successful.


With enough persistence, these APT groups can eventually gain access to some of even the best-defended networks. Turning an operation from one that affects 1,000 organizations to one that only impacts 10, while still very serious, is an achievable and real improvement for defenders as a community. Rapid sharing and analysis of intelligence on new operations and exploits through ISACs, community-driven efforts, with commercial security providers, with built-in security tools like [Chronicle's integration with BigQuery](#), and with public-sector partners turns the assumptions of the benefits of scale on their head. Together, we can make the bad guys make tough choices.

Moving infrastructure to the cloud can improve resilience as enterprises gain visibility into their own systems and defense strategies can be deployed more reliably than on-premises techniques. The defensive advantages of scale in the cloud also apply to stopping less sophisticated but more widespread threats from cybercriminals, skilled individuals, and less-spectacular nation-state efforts. [Distributed denial-of-service attacks](#), especially if targeted against time-critical services or on key dates, can cause significant financial damage. To paraphrase National Cyber Director Chris Inglis, we want to create a situation where threat actors have "to beat all of us to beat any of us."

Christopher Porter is the Head of Threat Intelligence for Google Cloud.

Summary

Cloud Compromise Factors (Q2 2022)



View from the frontlines: Threat actors use weak or compromised creds to target SSH, WordPress in Q2

Despite relative incidents being down due to customer controls overall, in Q2 threat actors frequently targeted weak and default-password issues for initial compromise, factoring in over half of identified incidents. A wide array of software suites were targeted for exploitation, with SSH, WordPress, and RDP software all frequently targeted. Once inside, threat actors frequently engaged in cryptomining, accounting for nearly two-thirds of incidents (65%).

Targeting password issues is a low bar for entry for threat actors, making it an often-successful technique for malicious actors that vary greatly in terms of their technical sophistication. The use of cryptominers indicates that the actors behind this type of targeting frequently expect to be promptly found and blocked

and they rely on volume of compromises rather than the potential longevity of any specific compromise. In addition, cryptominer attacks are often partially or fully automated, dramatically reducing their time to exploit an available vulnerability. As a result, while prevention is always paramount, clients should also be focused on minimizing their response and mitigation time metrics, as the longer a cryptominer infection persists the greater the potential cost. The high level of SSH activity suggests that organizations are using either no credentials or default credentials when spinning up cloud instances. Adjusting this approach via policy shifts can substantially improve an organization's cyber risk exposure while using cloud services.

Threat trends

Social engineering campaign disrupted and moved off of Google

In June 2022, Google's Threat Analysis Group (TAG), in coordination with the Trust and Safety team, mitigated a social engineering campaign hosted on Google Cloud Storage. The group had been tracked by TAG since mid 2021 as COLEUS and is known externally as the "[Stolen Images Evidence](#)" campaign.

The campaign involved actors sending social engineering emails with legal threats for copyright infringement through the contact forms on various company websites. The malicious emails alleged they contained proof of stolen content and directed users to a file hosted on Google Cloud Storage or Firebase Storage which masqueraded as a Google Drive page. The file contained a malware payload and it is suspected that this initial access broker operates a distribution-as-a-service model catering to both APT and cybercrime groups.

Customers can report spam, malware, or phishing by reporting [suspected abuse on Google Cloud](#), phishing websites to [Safe Browsing](#), and requesting content be [investigated and removed from Google](#).



File 'document' is ready for download

Your download should begin automatically.
Didn't work? Try downloading again.

Download my file

Tip: Click 'Keep' or 'Save' first, and then click the file name to open it

Screenshot of the file download hosted on Firebase Storage disguised as a Google Drive page.

Malicious files and URLs slipping by IT governance controls

Issue description


Cloud and on-premises users can be harmed by malicious files and URLs attempting to hide from organizations' IT governance protocols. A [recent report](#) from VirusTotal, analyzing 2021 through the middle of 2022, demonstrates how malware and phishing websites bypass corporate IT security controls. The controls fail to identify the malware assets' nefarious nature as they check the assets' context and external characteristics, instead of exploring their content in more depth. The report doesn't describe trends for each evasion technique, but two described trends are becoming more popular: how malware incorporates icons visually similar to legitimate Windows software, and how malware embeds and subsequently executes, from its packaging, installation programs for legitimate software. We further found, via separate raw VT data searches, evidence of these suspect techniques involving Google Cloud-related and more general Google software.

(Malicious files and URLs, cont'd.)

The VT report discusses four methods of how malicious files and URLs avoid detection.

- A. First, 10% of well-known, popular websites are seen to be distributing malware. Malware “legitimacy” is inherited from credible hosts.
- B. VT also found over 1 million signed malware samples, 87% of which were signed via legitimate certificates (the rest used revoked certificates, and so on). Apparently, attackers often fraudulently accessed signing workflows or signing authorities to sign their code – increasing the likelihood of its downstream acceptance.
- C. Attackers also created malicious applications incorporating icons visually similar to popular Windows applications, and designed malicious websites using URL favicons visually similar to well-known websites – for example, to build trust via “similarity.” For example, below is a graphic from the report on the growing number of submitted malicious samples with icons similar to popular Windows software.
- D. Finally, attackers also incorporated malware into legitimate software installers, or included legitimate software installers in their own malware “packaging” – again, to encourage installation via “good” association.

Such techniques can impact Cloud clients. We reviewed various Google-related VT data in the context of the techniques described above. We found seven malware packages that included the Cloud-executable `GoogleCloudSDKInstaller.exe`, which installs the Google Cloud CLI capability on Windows machines.



(Malicious files and URLs, cont'd.)

Dropped Files ⓘ				
	Scanned	Detections	File type	Name
✓	2022-08-17	2 / 69	Win32 DLL	C:\Users\user\AppData\Local\Temp\insjF10.tmp\insResize.dll
✓	2020-06-30	0 / 73	Win32 EXE	D:\Desktop\Desktopest\2.0 - AI Dungeon 2 Unleashed\GoogleCloudSDKInstaller.exe
✓	2022-08-16	1 / 66	Win32 DLL	C:\Users\user\AppData\Local\Temp\insuC102.tmp\UAC.dll
✓	2022-08-17	1 / 69	Win32 DLL	C:\Windows\Temp\inswA1A2.tmp\System.dll
✓	2022-03-28	62 / 70	Win32 EXE	C:\Users\<USER>\Downloads\l_cache_Synaptics.exe
✓	2022-08-17	0 / 69	Win32 DLL	C:\Users\user\AppData\Local\Temp\insd210F.tmp\insDialogs.dll

Here is VT data from one such package, Synaptics.exe, and the legitimate GoogleCloudSDKInstaller.exe as one of its dropped files.

We found several Google non-Cloud examples of such “file inclusion” in our VT data searches as well. Other legitimate Cloud files could also be packaged with other malicious installers.

Other Cloud-related attacks are also possible. For example, it is possible for websites to contain credible, almost-indistinguishable logs compared to Google Workspace logs – and therefore become phishing attack-vectors against Cloud clients. (Our preliminary VT data analysis did not find any Cloud URL favicons imitated by URLs submitted to VT. Of course, as VT doesn’t contain all URLs, this may happen with other websites.) Also, if an organization has weak VM/ container governance controls (for instance, not using antivirus software on VMs ingesting internet-originating files), or user (or administrator) access has been compromised, the malicious software described in the VT report could be downloaded into Google Cloud VMs – potentially leading to container abuse. Finally, if any of such malware infects on-premises environments, it may also compromise Cloud services via commands “issued

by” authenticated users. Using logged-in employees as, effectively, “conduits,” malware could mount Cross-Site Request Forgery-type attacks via HTTP, CLI, or similar channels, undermining Cloud resources. Use vigilance regarding the techniques discussed in the VT report.

(Malicious files and URLs, cont'd.)

Suggested mitigations for Google Cloud clients

Given the variety of potential attack vectors arising from the malware assets, as discussed earlier, the controls recommended below provide for a multilayer defense-in-depth strategy. The detective and preventive controls should help monitor and block the malware assets.

They look more deeply into assets' operations so as to better comprehend behavior. And while all the controls in the following list are suggested, the top three are particularly recommended.

1. Only install software from trusted hosting locations and manufacturers. Use [VirusTotal](#), check the software's hash code, and security-test files as further validations when installing software.
2. Consider turning on the [Advanced Protection Program](#) (APP) to protect cloud users from different online attacks. The APP protects against phishing attacks by requiring users to log in to Google Accounts using their security keys, flags or blocks users from downloading potentially malicious files, and offers other safeguards. The APP might have an impact on user experience – such as via additional interactions with security keys – but should still be seriously considered, at least for staff with significant sensitive data access.
3. Use [Chrome Enterprise](#) (CE) to secure users' interactions with general Cloud services. CE components include the Chrome browser, with its many security features against malicious URLs (for example, for this post, arising from a "similar favicon" attack) – such as site isolation, which prevents malicious sites from stealing data from the other websites a user is working with. CE also includes ChromeOS, which, among its security features, is a read-only OS, thereby substantially preventing local malware installation.
4. Use [Chronicle](#) for threat analysis. Chronicle merges the significant security and network telemetry generated within a customer's Google Cloud instance with Google and non-Google threat intelligence feeds to show individualized threats within each customer's environment. Use [Simplify](#), Google's recently acquired SOAR platform, to automate responses for specific threats (for example, shutting down a rogue VM instance if certain conditions were met).
5. The [Event Threat Detection](#) (ETD) capability within the Security Command Center (SCC) can be turned on to quickly detect Google Cloud threats based on logged cloud events. ETD monitors the Cloud Logging stream as well as Google Workspace logs, analyzing the creation, modification, and so on, of Google Cloud instance resources, Google Workspace domain user sign-ins, and many other activities. ETD deduces threats from logged event patterns.
6. [Control which third-party and internal apps can access Google Workspace data](#). Using Google Admin console settings, customers can permit, restrict, or block connectivity to third-party and internally made cloud apps by, respectively, relaxing or restricting the OAuth 2.0 scopes if/as used by such apps. Cloud data can be protected should such apps become compromised.
7. Use [Container Analysis](#) to perform vulnerability scans on container images in Google Cloud's Container Registry, the newer Artifact Registry,

(Malicious files and URLs, cont'd.)

or on local images customers build to place into these two repositories. Scans produce severity levels, including CVSS scores, permitting prioritized remediation.

8. Ensure that the [Security Health Analytics scanner](#) within the SCC is turned on. The Security Health Analytics scanner probes a variety of parameters in the Google Cloud environment at an asset-dependent cadence, and when the configuration for certain assets change, identifying configuration deficiencies. The scanner looks for misconfigurations in containers, IAM settings, and other assets.

Cloud and SaaS-enabled environments increasingly vulnerable to the connectedness of SSO

As cloud hosting becomes more prevalent, and the volume of cloud-optimized cross-application capabilities increases, it is our assessment that the volume of cloud-targeted attacks leveraging corporate single-sign-on (SSO) logins will rise. This is true of SSO for enterprise in many forms, including those with Software/Security-as-a-Service (SaaS), but especially impactful to cloud due to the ability to access and provision cloud resources.

Kimsuky, a nation-state threat actor, has been observed by researchers at Volexity accessing user Gmail account data through a hidden Chrome browser extension known as SHARPEXT.¹ The group, which reportedly works toward Democratic People’s Republic of Korea (DPRK) interests – according to members of the Korea Internet & Security Agency² – was able to install a malicious browser extension via phishing, leveraging pre-authenticated browser activity to read and exfiltrate data from other services such as Gmail content.³ The convergence of technologies, enabled by SSO, should be a point of concern for those transitioning to a cloud-centric environment, as the increased productivity provided by seamless SSO also provides broader access for attackers to otherwise confidential data.

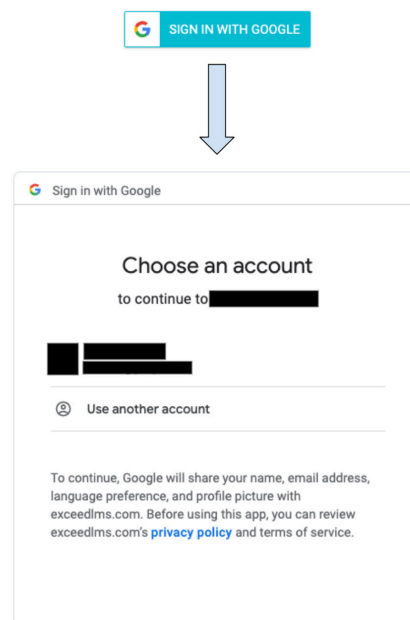


Figure 1. Legitimate Google-account SSO user flows can be triggered by an application (such as a browser extension) or accessing a website.

¹ Rascagneres, Paul, and Thomas Lancaster. "SharpTongue Deploys Clever Mail-Stealing Browser Extension "SHARPEXT."" Volexity, Jul 28, 2022, <https://www.volexity.com/blog/2022/07/28/sharptongue-deploys-clever-mail-stealing-browser-extension-sharpext/>. Accessed Aug 9, 2022.

² Lee, Taewoo, et al. "Kimsuky, STOLEN PENCIL, Thallium, Black Banshee, Velvet Chollima, Group G0094 | MITRE ATT&CK®." MITRE ATT&CK®, Aug 26, 2019, <https://attack.mitre.org/groups/G0094/>. Accessed Aug 9, 2022.

³ Goodin, Dan. "North Korea-backed hackers have a clever way to read your Gmail." Ars Technica, Aug 3, 2022, <https://arstechnica.com/information-technology/2022/08/north-korea-backed-hackers-have-a-clever-way-to-read-your-gmail/>. Accessed Aug 9, 2022.

(Cloud and SaaS-enabled environments, cont'd.)

Observations

In this case, the initial infection vector is via phishing and results in the installation of a developer-mode browser extension which, through a DevTools workaround, has its security warnings suppressed and targets a user's cloud-accessed data (such as online email applications).

While this malware specifically targets certain resources (Gmail and AOL mail), this is a design feature to reduce the likelihood of detection and need not be so targeted. The value of this method is to embed with the browser, which maintains the SSO and allows permission to the apps to leverage that access. This could theoretically be achieved with any application and it is assessed that future Trojans may use similar techniques to evade detection and inherit the valuable SSO accesses.

Key risks

Attackers gaining legitimate user account verification can cause serious problems:

1. **Legitimate access credentials.** The initial compromise leverages existing credentials, which makes it difficult to distinguish from legitimate admin access.
2. **Multi-factor authentication (MFA) is ineffective.** Since the user has already logged in and is simply allowing access permissions, traditional MFA methods do not protect the user.
3. **Ease of user permission flow.** A user login (inputting a password) may have already been conducted, so providing permission to the malicious application is single-click.
4. **Expected and trusted application.** Access requests are only submitted as part of an access to a web app, browser plugin, or similar that the user intends to access. Hence, the access request is expected and has high trust with the user.
5. **Third-party data leakage.** Apps and browser extensions that leverage these permissions may also be retaining user account data to third-party hosting, resulting in unintended data leaks.
6. **Single point of access and failure.** The potential breadth of an SSO compromise is vast and, without swift remediation, can result in an unknown level of compromise as the attacker can access a wide range of company infrastructure and data.
7. **Lack of user awareness.** Depending on the initial infection vector, the user may be unaware of the application's activity on their device, as warnings may be suppressed.

```

52 packetProc = function(request, body){
53     var url = request.request.url;
54     var mimeType = request.response.content.mimeType;
55     if( ( url.indexOf("https://mail.google.com/mail") != -1 ||
56         url.indexOf("https://mail.google.com/sync/u") != -1) &&
57         (mimeType == "application/json" || mimeType == "text/html") ) ||
58         url.indexOf("https://mail-attachment.googleusercontent.com/attachment/u") != -1 ||
59         url.indexOf("https://mail.aol.com/") != -1)
60     {
61         chrome.runtime.sendMessage({
62             "action": "packet",
63             "request": request,
64             "body": body
65         });
66     }
67 }

```

Figure 2. Javascript module of Kimsuky's "SHARPEXT" Chrome extension malware, which specifically searches for SSO-authenticated email account access. (Image: Volexity⁴)

⁴Rascagneres, Paul, and Thomas Lancaster. "SharpTongue Deploys Clever Mail-Stealing Browser Extension "SHARPEXT." Volexity, 28 July 2022, <https://www.volexity.com/blog/2022/07/28/sharptongue-deploys-clever-mail-stealing-browser-extension-sharpext/>. Accessed 9 August 2022.

(Cloud and SaaS-enabled environments, cont'd.)

Mitigations

1. **Raise user awareness.** We are yet to normalize user awareness around SSO access permissions to the same level as link verification in phishing emails, and it is an important addition to any SSO-enabled environments' user security training.
2. **[SSO account and access monitoring](#).** It is vital to implement effective monitoring of SSO account activity and automated remediation methods as a standard in order to detect and remediate SSO account compromise.
3. **Managing SSO-enabled software installation.** Avoid allowing installation of unvetted software on corporate devices that could erroneously request user access permissions. Manage such software through centralized software provisioning processes that allow for the verification of legitimate software, and regularly review and verify installed software. Zero-trust environments such as [BeyondCorp Enterprise](#) implement these sorts of policies at a fundamental level.
4. **Enterprise browser management.** This attack, as with many that target the browser, succeeds due to stealthy manipulation of the users' settings and the ability to run scripts based on browser activity. Correctly configured, implementing an enterprise-wide browser management solution (such as that implemented in [Chrome Enterprise](#)) can help prevent the successful running of malicious extensions.

Cloud application security reviews and data discovery

Issue description

A mid-2021 International Data Corporation (IDC) survey indicated that almost three-fifths of organizations felt that inadequate visibility and identity management were major cloud infrastructure concerns,⁵ as attackers typically exploit excessive application access rights, misconfigured data-sharing capabilities, and similar attack vectors to abuse cloud assets. The IDC survey highlighted a key impact from such vulnerabilities: 63% of the organizations actually had sensitive data exposed.

Cloud application security reviews and a data discovery process facilitate secure cloud workloads and user interactions and effective data use.

And this grew to 85% for firms spending at least \$50M annually on cloud infrastructure. Cloud application

security reviews and a data discovery process facilitate secure cloud workloads and user interactions and effective data use. Understanding application security threats – and having more visibility into cloud data locations and purpose – reduces vulnerabilities and permits effective data analysis.

Cloud threats

Organizations continue to upload diverse data, such as business documents, PII, and healthcare data – as well as onboard more applications – to the cloud. The goal is to use the cloud’s scalability, analytics, and other features. But security issues can arise without appropriate IT governance:

- In Q1 2022, according to the FBI, Russian-state-supported attackers gained access to a non-governmental organization (NGO) cloud instance as they brute-forced the instance’s password and enrolled their own device in the NGO’s multi-factor authentication (MFA) process.⁶ Capturing the password and using their “own” MFA device, the attackers ultimately exfiltrated the NGO’s data.

⁵ Ermetic, “IDC Survey Report: State of Cloud Security 2021,” Ermetic [web survey access], 2021, <https://l.ermetic.com/wp-idc-survey-results-2021/>, (accessed May 12, 2022).

⁶ Gatlan, Sergiu, “FBI warns of MFA flaw used by state hackers for lateral movement,” BleepingComputer, Mar 15, 2022, <https://www.bleepingcomputer.com/news/security/fbi-warns-of-mfa-flaw-used-by-state-hackers-for-lateral-movement/>, (accessed Jul 29, 2022).

(Cloud application, cont'd.)

- According to a security publication, and security vendor TrendMicro, in 2021, cybercriminal group TeamTNT launched a campaign against misconfigured Docker REST APIs exposed in many public cloud instances.⁷ Exploiting weak API passwords, promiscuous instance firewall rules, and other flaws, the threat actor pulled and installed malicious container images from the Docker Hub central repository into these cloud instances. The malicious container code then tried stealing AWS login credentials, installed cryptomining software, and tried other means of VM/container abuse. By late 2021, 150K of TeamTNT's malicious images were downloaded via Docker Hub, suggesting relatively widespread compromise.
 - In a Q4 2021 security-industry survey sponsored by vendor CyberArk, security leaders indicated that over 80% of an organization's internal applications – including Google Cloud and other cloud-based platforms such as AWS, Salesforce, and ServiceNow – experienced access abuse from internal staff over the past 12 months.⁸
- examined the financial implications stemming from certain cloud architecture choices. For instance:
- The NY Department of Financial Services (NYDFS) Cybersecurity Regulation, 23 NYCRR 500, requires banks, insurance companies, and other financial institutions operating in NY State to maintain an information security program. Providing sufficient funding to implement modern data security standards, maintaining an incident response process, and other tenets are the requirements. If there is a breach due to improper implementation of such requirements, financial penalties can rise to \$1,000 per user per violation – which could be significant if many users or records are compromised. In June 2022, for example, the NYDFS fined several Carnival Cruise Line companies (CCLC) a total of \$5M. During years 2019–2021, CCLC was impacted by ransomware, phishing, and other data breaches and a NYDFS examination indicated that CCLC lacked MFA processes, procedures for monitoring unauthorized network traffic, and other key security controls.⁹
 - Without a data oversight process, a firm may also overspend on data storage. If data-access frequency is not examined, companies may store data in a CSP's more expensive "standard" storage tier. However, storing infrequently used data in the "cold" or "archive" storage tiers would cost less. It costs more to retrieve data from such tiers, but storing infrequently used data in such tiers leads to considerably higher savings overall.

Regulatory and business risk

Improper cloud oversight leads to regulatory and business risks, too. Some organizations have not devoted proportional security resources to understand corresponding vulnerabilities. For example, they may forget that operations by low-skill cybercriminal groups can lead to much larger fines on the organization than potentially anticipated. Or organizations might not have

⁷ Gatlan, Sergiu, "Cryptojacking worm steals AWS credentials from Docker systems," BleepingComputer, Aug 18, 2020, <https://www.bleepingcomputer.com/news/security/cryptojacking-worm-steals-aws-credentials-from-docker-systems/>, accessed Jul 27, 2022; and Trend Micro Research, "Compromised Docker Hub Accounts Abused for Cryptomining Linked to TeamTNT," TrendMicro, Nov 9, 2021, https://www.trendmicro.com/en_us/research/21/k/compromised-docker-hub-accounts-abused-for-cryptomining-linked-t.html, accessed Jul 26, 2022.

⁸ CyberArk, "CyberArk Research: Lack of Security Controls and Visibility Into User Activity Continue to Put Organizations at Risk," CyberArk, Nov 2, 2021, <https://www.cyberark.com/press/cyberark-research-lack-of-security-controls-and-visibility-into-user-activity-continue-to-put-organizations-at-risk/>. Accessed Jul 12, 2022.

⁹ Peter Baldwin, Bob Mancuso, and Jane Blaney, "New York Department of Financial Services Announces \$5 Million Penalty in Most Recent Cybersecurity Enforcement Action," Faegre Drinker Biddle & Reath LLP, Jul 11, 2022, <https://www.discerningdata.com/2022/new-york-department-of-financial-services-announces-5-million-penalty-in-most-recent-cybersecurity-enforcement-action/>. Accessed Sep 2, 2022.

(Cloud application, cont'd.)


Providing governance

Appropriate cloud governance can address the risks discussed earlier. As data volumes and application capabilities grow, cloud administrators are interested in understanding:

- The cloud data's location, including the projects it is being shared with
- If the data contains any sensitive information, and how data contents are changing over time
- If applications and workloads are implementing appropriate security controls, given data classifications, including providing support for regulations – like data residency – as required

Application security reviews, which examine application architecture, data security controls, likely threat scenarios, and similar aspects can answer such queries. However, given the considerable manual effort involved, such assessments are performed less frequently and on a smaller number of key applications. When there are many applications to examine, or for real-time monitoring, a more automated security process should be created.

One “cost-effective” cloud governance process to assess different kinds of application workloads based on known a priori risks is shown below.



This process matches the effort of application security reviews to the risk that different workloads present. It also permits data to be annotated with its “characteristics” (like backup locations, and so on), and redacted – at a regular cadence. This kind of an approach (with suggested Google Cloud tools mentioned on the next page), lets IT owners:

- Keep enterprise asset and application inventories (for example: CMDB) up to date
- Keep data classifications current, including as sensitive data is added/removed from repositories
- Align application security controls – including data redaction – to better-understood information security risks, including adhering to regulations such as data-breach notification laws when in-scope data has been identified

LLP, Jul 11, 2022, <https://www.discerningdata.com/2022/new-york-department-of-financial-services-announces-5-million-penalty-in-most-recent-cybersecurity-enforcement-action/>, (accessed Sep 2, 2022).

(Cloud application, cont'd.)

Google Cloud mitigations

The following are suggestions for Google Cloud customers:

1. Perform cloud application security reviews based on the risk of onboarded applications, workloads, or data feeds. For high-sensitivity and key moderate-sensitivity assets, in-depth security reviews should be carried out annually. Less sensitive or less critical assets can be assessed less frequently – such as every two, three, or more years. Such reviews can reduce risk. In 2018, a financial firm correlated BitSight's Security Ratings, which track organizations' publicly facing security controls (like patched web servers and so on) with security breaches affecting over 3,000+ firms.¹⁰ Compared to firms with higher Security Ratings, those with low Security Ratings were four times as likely to be breached. Application security reviews would improve an organization's security controls, thereby increasing its Security Ratings. This in turn should reduce the firm's susceptibility to breaches, as per the study.
2. Use a DLP solution that supports automated scanning and classifying of sensitive data, and real-time data redaction and de-identification using tokenization or other techniques. Tokenization converts a sensitive data element, like a Social Security Number, into an obfuscated string. The process is consistent – for example, the same data element is converted into identical tokens and basic database operations, like JOINS or search, continue to work despite the "anonymized" tokens. For Google Cloud customers, Google has introduced [Cloud Data Loss Prevention](#) to address such needs. Cloud DLP can automatically classify and de-identify data within Cloud Storage, BigQuery, and Datastore. It can also support additional data sources. Cloud DLP's ability to tokenize and do other data redactions will reduce the potential for data abuse, as inappropriate data access will yield, effectively, relatively valueless information.
3. Use [Dataplex](#) to search for and tag Google Cloud data with technical and business metadata, such as its lineage, freshness, any transformation process modifying it, and so on. Such visibility gives users data "context" to answer various governance questions. For example, if a data set's confidentiality is compromised, knowing its lineage would identify related data assets that also need to be inspected/secured as they're at risk now, too. Further, cataloging data also reduces breach potential – as knowing the data's sensitivity and usage allows for its selected security controls to match its likely abuse scenarios.
4. Use [Data Access audit logs](#) and [Admin Activity audit logs](#) within the Google [Cloud Logging](#) tool set to determine which users or accounts are viewing or writing data, or changing data permissions. Monitoring what data access events are occurring this way also reduces breach potential, as organizations can detect misuse by monitoring for unexpected logged event patterns. Such audit logs may also assist with forensic analysis, and could prevent a repeat of prior data security incidents by understanding why they were successful originally.

¹⁰ Research Signals, "Cybersecurity factors powered by BitSight," IHS Markit, Mar 31, 2022, <https://cdn.ihs.com/www/blog/Cybersecurity-factors-powered-by-BitSight.pdf>, (Accessed Jul 14, 2022).

(Cloud application, cont'd.)

5. Follow [GCAT's policy intelligence tool recommendations](#) to optimize different Google Cloud user identity and data management controls, including helping automate user access re-certifications, placing data into cold storage if it hasn't been accessed in 30 days, and improving analogous capabilities. Google Cloud permits expressing various security policies as code so administrators can monitor, and action, policy drifts in a more automated manner. For example, the [role recommendations](#) tool uses machine learning to recommend user-permission modifications by reviewing a user's prior 90-day resource access patterns. If permissions haven't been used, the tool may recommend their removal. An administrator can "apply" such recommendations [via Google Console or API calls](#), to therefore enforce least-privilege in a more streamlined manner. Other policy tools that streamline Google Cloud end-to-end user access and data management can be found [here](#).
6. The [GCAT Risk Governance of Digital Transformation in the Cloud guide](#) suggests how cloud security control effectiveness can be more readily measured. Data-driven analysis can occur because the controls cover the entire cloud environment, due to the creation of standardized, widely deployable workloads with relevant integrated controls – the embedding of which can also be measured – and the integration of security into software development, ensuring that appropriate controls are integrated into code (which can also be measured before code release). Using such "metrics," Google Cloud customers can manage control effectiveness (for example, after application security and data-discovery assessments have completed). They can check if controls operate properly after the need for them is established.

Hijack of cloud infrastructure service accounts set to rise

Cyberthreat groups can reduce the risk their attack will be detected by leveraging poorly defended and trusted servers, domains, or applications. The use of compromised legitimate infrastructure allows actors to disguise their operations as legitimate traffic¹¹ and has been observed as an effective technique against even the most experienced of security professionals.¹² Leveraging this attack vector can be much simpler for the attacker if they can manipulate the infrastructure to better fit the disguised approach. As organizations increasingly transition toward cloud hosted, this presents a new attack surface that can offer an even greater effect for an even lower exposure: compromising the service-account credentials for the underlying cloud project that hosts an organization's enterprise network itself, allowing architectural control over trusted infrastructure in a target's supply chain, and the potential for cloud admin access to go undetected for long periods of time. Threat groups have been observed leveraging compromised service-account credentials to run expensive cryptomining workloads in customer environments, but greater concern would arise should they choose to keep these actions covert and leverage the access for other nefarious activities. During May–July inclusive,

Google Cloud has detected projects affected and their corresponding exposed service-account credentials. These compromises often get detected and remediated before they provision high volumes of resource, but even the few that avoid detection can cause significant impact. The disproportionate return on investment for such attacks is assessed to increase their prevalence in the coming years and has the potential to impact across all platforms where credential exposure is a risk.




Figure 1. Accessing the legitimate infrastructure's underlying cloud hosting environment provides greater opportunities for attack tailoring.

¹¹ Galloway, Jeremy, and Mitre. "Compromise Infrastructure, Technique T1584 - Enterprise | MITRE ATT&CK®." MITRE ATT&CK®, Apr 20, 2022, <https://attack.mitre.org/techniques/T1584/>. Accessed Aug 8, 2022.

¹² Kopeytsev, Vyacheslav, and Seongsu Park. "Lazarus targets defense industry with ThreatNeedle." Securelist, Feb 25, 2021, <https://securelist.com/lazarus-threatneedle/100803/>. Accessed Aug 8, 2022

(Hijack of cloud infrastructure, cont'd.)

Key risks

The combination of techniques provides several advantages to the attacker:

- 1. Legitimate access credentials.** The initial compromise leverages existing credentials, which makes it difficult to differentiate from legitimate admin access.
- 2. Cloud-access keys in public code repositories.** While capabilities to protect access credentials exist for cloud (such as Google Cloud [Secret Manager](#)), they are not always implemented in infrastructure-as-code deployments and credentials may be exposed through public code repositories.
- 3. Customization of onward attack.** With cloud admin access, the fake web server (or equivalent infrastructure) can be customized to further obfuscate the malicious activity or improve the technical implementation of the attack, increasing the chances of successful compromise.
- 4. Stealthy pre-positioning.** Once initial cloud admin access is confirmed – provided that the admin credentials are not changed – there is very little detectable “presence” of the attacker in the environment, which could allow them to remain hidden for long periods.
- 5. Maskable persistence.** Cloud infrastructure provides an additional layer of persistence options for an attacker as they can create cloud access accounts, account reset options, and – as a fallback – traditional persistence on new network infrastructure (such as the new fake web server). The increased scope of access and low detectable presence increases the complexity of persistence detection.

(Hijack of cloud infrastructure, cont'd.)

Observed events

A recent incident we tracked observed threat actors leveraging compromised service-account credentials to abuse targets' infrastructure by provisioning cloud resources. These resources can be used for any purpose, including but not limited to cryptomining, hosting malicious content, or pre-positioning for an onward attack.

In this instance, the use of service-account access was limited only to provisioning vast amounts of resources. It is very feasible, however, that more sophisticated actors could leverage this access for less detectable and longer-term goals, such as establishing "jump boxes" in victim environments to avoid detection or the creation of additional service accounts as a form of persistence.




Figure 2. CPU resource usage over time showing a threat actor provisioning cloud resources in a customer environment using stolen service-account credentials.

(Hijack of cloud infrastructure, cont'd.)

Mitigations

Taking the time to implement good security configuration in cloud environments is vital:

- Set up [private connectivity](#) for your service accounts so that, even if compromised, your service-account credentials can only be accessed through approved connections and connectivity is restricted from IP ranges not in pre-approved ranges
- Maintain a policy of [least privileged access](#) for all of your accounts, including service accounts, to reduce the impact of a successful compromise
- Monitor for [leaked service-account credentials](#) in public code repositories and similar data stores where cloud API development may occur
- Ensure that service-account credentials are treated with appropriate levels of protections: Encourage best-practice [Secret Manager](#) usage for development work and ensure that developer workflows avoid leaking those credentials to public repositories
- As a further indicator of anomalous activity, ensure appropriate [budget alerts](#) to signal where your organization may be billed for resources that it did not use

Responding to the next SolarWinds: Logging tools

One of the key lessons learned from the SolarWinds incident was the importance of logging. In response, the US government later published a [memorandum](#) of logging requirements with a three-tier maturity model. The SolarWinds incident went undetected for months and impacted over 18,000 of its customers.

- [DNS logs](#), when available, proved valuable in identifying instances of SolarWinds Orion talking to command-and-control servers. Organizations can review their DNS logs for connections to known IP and domain Indicators of Compromise (IOCs).
- [Authentication and authorization logs](#) uncovered the attacker successfully authenticating through the identity provider and later using this access for privilege escalation.
- Email logs captured changes to email-forwarding rules. Robust logging, monitoring, and alerting can help organizations [prevent and mitigate email exfiltration](#).

During the SolarWinds incident, attacks enumerated a list of Active Directory (AD) accounts during

reconnaissance to identify privileged accounts and move laterally. Organizations can improve detection of anomalous behaviors by reviewing and enabling [security-related logs](#) and using [pre-built queries for BigQuery and YARA rules for Chronicle](#) to regularly monitor cloud security threats such as enumeration of cloud resources and IAM permissions, or attempts to export service-account keys or edit the metadata on compute instances. Google Cloud also provides [lateral movement insights](#) to help identify roles that give a service account from one project the ability to impersonate service accounts on another project.

The SolarWinds incident demonstrated attackers' tactics searching for and compromising privileged accounts to move laterally. When working with Google Cloud Identity and Access Management (IAM), organizations should be aware that broad permissions of basic roles such as Project Viewer, Project Editor, and Project Owner include thousands

(Responding to the next SolarWinds, cont'd.)

#	Cloud Security Threat	Log Source	Audit	Detect	ATT&CK® Techniques
1	Login & Access Patterns				
1.01	Login from a highly-privileged account	Google Workspace Login Audit (Cloud Identity Logs)		✓	T1078.004
1.02	Suspicious login attempt flagged by Google Workspace	Workspace Login Audit (Cloud Identity Logs)		✓	T1078.004
2	IAM, Keys & Secrets Changes				
2.21	Permissions granted to impersonate Service Account	Audit Logs - Admin Activity	✓	✓	T1484.002
2.22	Permissions granted to create or manage Service Account Keys	Audit Logs - Admin Activity	✓	✓	T1484.002
3	Cloud Provisioning Activity				
3.01	Changes made to logging settings	Audit Logs - Admin Activity	✓	✓	T1562.008
3.11	Unusual number of firewall rules modified in the last 7 days	Audit Logs - Admin Activity		✓	T1562.007

Sample use cases of the security analytics

of permissions and can give attackers the ability to delete log sinks, filter logs, delete the log bucket, or exfiltrate log data. Following the principle of least privilege, utilize either [pre-defined roles and permissions](#) for Cloud Logging or create custom roles with fine-grained permissions – rather than using the available basic roles. When using a group for IAM rules, organizations should implement a directory-syncing mechanism with their source of truth to remove access

and permissions for accounts that are suspended or deleted. Orphaned accounts with IAM roles and permissions create more opportunities for attackers to achieve lateral movement within an organization.

(Responding to the next SolarWinds, cont'd.)

Google Cloud Mitigations

- Configure [Google Workspace logs](#) to be sent to Cloud Logging to centralize your visibility. Review [Google Workspace Login Audit logs](#) for events such as email forwarding outside your domain
- [Centralize logs and at the organizational level](#), which not only protects logs from being tampered with by attackers but also gives defenders visibility across the entire organization, making it easier to identify threats and anomalous events
- Utilize infrastructure as code to reduce configuration errors and [automate deployments of logging agents](#) for VMs, apply IAM roles and permissions, and even validate their scripts
- Review [Cloud Audit Log best practices](#), which include key actions to minimize your organization's risk

Google Cloud