

Threat Horizons

Cloud Threat Intelligence
July 2022. Issue 3



Office of the CISO

Providing Strategic Threat Intelligence to Those in the Cloud

Google's Office of the CISO team is pleased to publish another issue of our *Threat Horizons Report*. The report is based upon strategic risks, threat intelligence, and various security observations from Google's Threat Analysis Group (TAG), Trust & Safety, Cloud Threat Intelligence teams, and various other internal teams. The goal of this report is to provide security leaders and other executives strategic, actionable information that enables organizations to better understand emerging risks and further secure their cloud environments against ever changing threats.

Summary of Observations

In this issue, we provide an update on the evolution of cloud misconfiguration, cryptomining, and phishing. While these have been covered in previous reports, we will share our latest observations on how attack methods have evolved over time.

It is often not enough to just understand how attackers are changing; it is just as important to understand what the most effective defense and mitigation strategies are. As such we will also share a Cloud Security Checklist to help organizations establish a baseline for their Cloud defenses. This checklist will be an ongoing series to provide operational guidance to better secure your Cloud infrastructure. We will also go in-depth on a few actionable programs (such as Web Risk) organizations can implement to mitigate their risks.

Threat Trends			
01	02	03	04
Beyond the 1st Click: Establishing Trust in Phishing Attempts	An Update on Cyber Activities Related to the War in Ukraine	Brute Force and Vulnerable Software Attacks Drive Continued Diligence in Cloud	Trending Abuse Tactics And How To Defend

Defensive Strategies				
05	06	07	08	09
Open Source Supply Chain Security	Cloud Security Checklist: When was your last backup?	Zero Trust: Lessons and Misconceptions	Going Deeper: Logging	Reminders and Recommendations

01 Beyond the 1st Click: Establishing Trust in Phishing Attempts

In the industry, phishing continues to be an area where bad actors adapt their techniques. While phishing is not a cloud-only threat, it can introduce vulnerabilities to cloud-hosted applications. For example, domains similar to a company's cloud-hosted domains can be purchased by a threat actor to launch different phishing attacks. In previous Threat Horizons reports, we covered specific phishing campaigns that had been identified by Google's Threat Analysis Group. In this report, the latest trends show phishing emails are being designed to look more benign and that bad actors are establishing trust by increasing seemingly legitimate interactions with victims before compromise.

Similar Domains

A resurgent trend we're observing is the purchasing of similar-sounding or similar-spelled domains by bad actors. Purchasing domains similar to a company's domain (such as using a different Top Level Domain suffix, i.e., ".us", ".co", or ".biz" instead of ".com") continues to be a popular tactic along with "typo-squatting" where slightly misspelled domains are used. Many organizations do not own domains that are similar sounding or similarly spelled to their organization, nor do they keep track of similar domain names that have been purchased by bad actors. Attackers send phishing emails from the "same" domain, and host phishing web pages that look identical to a company's website with the explicit goal of stealing sensitive data.

An example of this was observed in March 2022, where Google's Threat Analysis Group (TAG) observed ancillary domains of various companies being purchased by threat actor Exotic Lily¹. A dozen of these domains were identified to be very similar to current Workspace customers. Exotic Lily used these ancillary domains to send phishing emails to other companies. It also sent malicious payloads to its victims using legitimate, well-known file hosting services to further limit suspicion by the victims.

Relationship Building

Another common tactic that continues to be observed is when bad actors actively impersonate legitimate sounding organizations (especially in journalism or education) with the objective of interacting with the target in a trusted manner before launching an attack.

TAG analyzed such trust-building activities by a threat actor called Kimsuky. Kimsuky has been targeting North Korean experts with fake news interview requests. Their phishing campaigns typically start as a non-malicious email impersonating a journalist, asking if the target would like to comment on an article or participate in an interview. If the target agrees, they are

emailed a link directing them to a site containing additional interview questions which also tries capturing their credentials.

Recommended Mitigations

Customers operating in Google Cloud or Workspace can continue to protect themselves against phishing threats through adopting a multi-layer approach:

- Companies should consider using a brand protection service, or purchasing domains ancillary to their domain(s) themselves to prevent domain spoofing. Additionally, active monitoring should be set up for domains that sound similar, are spelled similarly, and ccTLDs of company names for bad actors.
- Companies may also consider using Google Cloud's [Web Risk](#) service. Companies can check any impersonating sites against the Web Risk Evaluate API, which returns confidence scores based on the maliciousness of webpages. In addition, urls can be sent to Google through the Submission API, which, after a Google review, will be added into the Safe Browsing service if they fall under [Safe Browsing policies](#). The offending sites will be blocked or trigger warnings on major Google products, including Google Search, Android apps, Chrome, Gmail, and several other major browsers.
- Companies should [configure SPF, DKIM, and DMARC capabilities](#) for their email to prevent email spoofing.
- Turn on free [Advanced phishing and malware protections](#) within Workspace. These protections scan attachments and links in email bodies, protect against spoofing of employee names, quarantine emails for further analysis, and offer other features.
- Customers should consider enrolling in Google's free [Advanced Protection Program](#) (APP). APP uses an even more rigorous version of Safe Browsing in Chrome to protect against malicious file downloads; permits only Google apps and verified third-party apps to access a user's Google Account data, and only with their permission; and offers other protections.
- Individuals should always verify an outreach from a previously unknown contact through another trusted method (i.e. contacting the company directly to verify).
- Customers may also use [Google's Work Safer](#) (WS) program, which provides firms with access to leading security for emails, meetings, files, and other assets. The security provided includes zero-trust access for corporate resources; endpoint protection for BYOD or company-owned devices; and physical security keys to prevent account takeover.

02 An Update on Cyber Activities Related to the War in Ukraine

Google continues to partner closely with government and industry partners in response to the Russian invasion of Ukraine. During the course of the conflict, we have observed threat actors targeting Ukrainian critical infrastructure entities including oil and gas, telecommunications and manufacturing.


Key attack trends include DDoS, data destruction, and credential phishing. For DDoS, there has been a large uptick in DDoS as a service offerings which has enabled less sophisticated actors with sufficient funds access to capabilities to disable large enterprise services. The observed services utilize attack methods targeting layer 4 and layer 7. This situation is being closely monitored in partnership with [Google Cloud Armor](#) and [Google Safe Browsing](#) teams. Additionally, Google takes enforcement action against GCP resources that are not operating within Google Cloud's [Acceptable Use Policy](#).

In addition, please refer to [this recent blog post](#) by Google's Threat Analysis Group for a deep dive on individual threat actors and IoCs related to the war in Ukraine.

Journalists, media organizations, human rights entities, elections-related entities, and government organizations impacted by the war in Ukraine should consider applying for Google's free DDoS protection service - [Project Shield](#). Project Shield has helped defend over 180 Ukrainian websites that provide critical information like current news and evacuation resources from DDOS attacks.

03 Brute Force and Vulnerable Software Attacks Drive Continued Diligence in Cloud

As covered in previous Threat Horizons reports, Google continues to observe in the cloud industry overall that attackers are automatically scanning for and compromising misconfigured Cloud resources. The primary goal for attackers in compromising Cloud resources has not changed; cryptomining and ransomware continue to be the most common purposes of compromise. In [Threat Horizons Issue #1](#), we list out the most common actions after compromise. While our focus in the previous issue was **what** the attacker does with compromised instances, we are following up with an assessment of **how** attackers are getting in to install cryptominers, and our recommendations for mitigation.



Common Attack Vectors

In recent months (Q1 2022), one of **the most common attack vectors used across cloud providers was brute force of cloud services that are exposed to the internet and have a weak or default password**. Due to the large number of public breaches in the past few years and the continued use of weak passwords in many organizations, bad actors are aggregating login information and launching brute-force logins against cloud admin accounts and privileged users. Advances in automation and computing speed in recent years have led bad actors to be able to launch these brute-force attacks intelligently and at a significant scale.

Close behind brute force attacks was the **exploitation of vulnerable software**. Beyond attackers exploiting vulnerable software, many organizations continue to struggle to patch and update their systems with the latest versions to protect themselves from compromise. With the continued increase in industry-wide vulnerabilities such as Log4J coinciding with an increasingly complex technology supply chain, known vulnerabilities which are not addressed in time will continue to be a common vector for compromise.

Zero-days continue to proliferate as new research emerges and technologies reveal latent risks. Google's Project Zero has recently published a [report](#) indicating the highest ever number of zero-day vulnerabilities reported in 2021. Despite the large increase in zero-days reported, Google has not observed significant wide-scale exploitation prior to patches being released.

Finally, the third most common attack vector we've seen in cryptomining is the **leveraging of stolen credentials** which may have been inadvertently leaked to public code repositories. This is most often due to human error, where keys may be hardcoded into a code commit, or are accidentally leaked from a test environment. As more engineering teams work together via

publicly-hosted code repositories, these human errors will continue to lead to additional compromises.

Weak firewall rules played a small role in cryptomining compromises, however it did act as a gateway to other common attack vectors such as brute force.

Shifts in Targets

There has been a distinct shift in ransomware targets post-initial compromise. Previously, the primary focus was specific files or file systems in general. **Attackers are now moving on from file systems to databases where there may be a higher chance of hitting critical corporate data.** The most common technique observed was where attackers were seen brute forcing SQL databases, cloning a database table into a new table, encrypting the data, and proceeding to drop the original table. Attackers have been observed leaving instructions in the new table that instruct the victim to transfer funds to a specified crypto wallet to recover their data. Similar tactics have been seen around cloud project takeover, with threats to delete data & resources. These attacks were most commonly observed in developer and proof of concept (POC) instances. **In many instances, these were targeted due to fewer security controls being placed in non-production environments due to their perceived lower risk.**

Recommended Mitigations

We see Google Cloud customers substantially mitigating these and other risks by adopting these practices:

- Implement detective and preventive controls to ensure Cloud resources are not being inadvertently exposed to the internet. GCP's Security Command Center provides a consolidated view into overall security health/risk of your assets, as well as what is currently publicly accessible. [Org policies](#) can also be used to provide preventative controls.
- To help detect and mitigate crypto mining, Google has added [Virtual Machine Threat Detection \(VMTD\)](#) to Security Command Center Premium. VMTD is a detection capability that **provides agentless memory scanning to help detect threats like crypto mining malware inside your virtual machines running in Google Cloud.** To get started with VMTD, open the Settings page in Security Command Center Premium. Click on "MANAGE SETTINGS" under Virtual Machine Threat Detection. You can then select a scope for VMTD. To confirm that VMTD is working for your environment, you can download and execute a test binary that simulates cryptomining activity.
- [Enable the detection of security keys](#) to prevent improper storage of secrets like tokens and private keys. Consider using [Secret manager](#) to securely store secrets.
- Set up [container](#) and [web security](#) scanning to help detect insecure configurations and vulnerabilities to reduce the exploitable attack surface.


04 Trending Abuse Tactics And How To Defend

There continues to be an increasing trend of cryptomining abuse across cloud providers and on-prem environments, including a shift from typical Proof of Work (PoW) to other new forms of mining such as Proof of Capacity (PoC), Proof of Space and Time (PoST) and bandwidth based miners. These trends directly mirror shifts in the crypto market as it continues to evolve.

“Proof of” Overview

The most common type of cryptomining is based on a concept called “**Proof of Work**”. Proof of Work is where, in order to “mine” or derive the underlying value that a cryptocurrency is based on (bitcoin, for example), a mathematical puzzle must be solved successfully in order to be added to the blockchain. Solving this mathematical problem requires significant computational power, and is the basis for cryptocurrency mining, both legitimate and illegitimate. Most of the well known cryptocurrencies are based on Proof of Work.

The difference, at a high level, between Proof of Work, Proof of Capacity and Proof of Space and Time is how the “solve” for the problem is less compute-intensive and more based on various other factors, such as available hard drive capacity (**Proof of Capacity**), the unchanged capacity used to solve it over a period of time (**Proof of Space and Time**), or available internet bandwidth. Many of these were designed with the legitimate use of reducing the power consumption and environmental impact of traditional Proof of Work mining, but are used by bad actors during compromises to mine cryptocurrency from victims.



Observations

Proof of Work (PoW)

Over 85% of all cryptocurrency mining cases in the cloud are based on the Proof of Work (PoW) consensus mechanism resulting in high resource utilization often seen as a spike in CPU usage.

Proof of Capacity (PoC) / Proof of Space and Time (PoST)

PoC requires storage to be allocated to the mining network, which can impact overall availability of cloud storage capacity.

The emergence of Proof of Space and Time (PoST) first started as early as June 2021. In both instances of storage based mining, Google was able to rapidly deploy detections & mitigations.

Bandwidth based mining

In 2022, actors have attempted to commit cloud network resources to mine cryptocurrencies that require network bandwidth.

Cryptocurrency mining impacts customer resource availability, spans from compute to storage, and now bandwidth. This can consequently result in reduced performance or even downtime if resources are exhausted as a result of a compromised instance used for mining.

Recommended Mitigations

We see Google Cloud customers substantially mitigating these and other risks by adopting the following practices:

- Become familiar with how to [handle and invalidate compromised credentials](#) which is often the precursor to cryptocurrency mining.
- Google Cloud's [Security Command Center](#) (SCC) Premium includes [Virtual Machine Threat Detection](#), which provides agentless runtime threat detection for cryptomining, as well as [Event Threat Detection](#) (ETD) which alerts when there are network detections indicative of mining activity.
- Use the Cloud Monitoring functionality within Google Cloud's Operations suite to [collect metrics](#) and [alert](#) based on resource consumption such as CPU, network, and storage.
- [Set up budget and budget alerts](#) within Billing to detect spikes in costs that are often an indicator of compromise.
- Continue to stay updated on the latest crypto mining trends, as the crypto industry continues to shift away from computationally intensive resources to more readily available resources which will lead to shifts in incidents and abuse.

Defensive Strategies


As we evolve the Threat Horizons Report, we have received feedback that our Google Cloud Customers want not only information on the latest threats, but recommendations and more information on how better to defend their environments. As such, we will be adding this “Defensive Strategies” section to further inform and remind customers and security professionals alike of best practices and recommendations.

05 Open Source & Supply Chain Security

A little over a year ago we published [Know, Prevent, Fix](#), which laid out a framework for how the software industry could address vulnerabilities in open source software and improve software supply-chain security.

The landscape has changed greatly since then:

- Prominent attacks and vulnerabilities in critical open source libraries such as Log4j, Codecov, and a [650% year-over-year increase](#) in cyberattacks aimed at open source suppliers made headline news, **bringing a new level of awareness to the issue** and unifying the industry to address the problem.
- The **US government formalized the push for higher security standards** in the May 2021 [Executive Order on Cybersecurity](#). The release of the [Secure Software Development Framework](#), a set of guidelines for national security standards on software development, sparked an industry-wide discussion about how to implement them.
- Last August, technology leaders including Google, Apple, IBM, Microsoft, and Amazon **invested in improving cybersecurity** — and Google alone pledged [\\$10 billion over the next five years](#) to strengthen cybersecurity, including \$100 million to support third-party foundations, like OpenSSF, that manage open source security priorities and help fix vulnerabilities.



The figure above details the many stages of the software supply chain for an open source dependency. Organizations may approach this many different ways - some may build packages from source themselves, while others pull packages from repos that they trust.

A few organizations including Google centralize control and actively secure each step of the end-to-end process. In our case, we start by maintaining separate secured copies of the source code for our dependencies and perform our own vulnerability scanning. We continuously fuzz [550 of the most commonly-used open source projects](#), and as of January 2022 have found [more than 36,000 vulnerabilities](#). This makes us one of the largest contributors to the OSV.

Google manages an end-to-end build, deploy, and distribution process that includes integrated integrity, provenance, and security checks. Based on our internal security practices, we have created the [SLSA framework](#) to enable organizations to assess the maturity of their software supply chain security and understand key steps to progress to the next level.

We recognize that most organizations do not have the resources or experience to construct and operate such a comprehensive program. Instead, their development teams might individually decide where they get third-party source code and packages, how they are built, and how to redistribute them within their own organizations according to their goals, threat and risk model, and resources. However, the lack of an end-to-end process creates risk exposure [each step of the way](#).

This month Google joined the Open Source Security Foundation (OpenSSF), Linux Foundation and industry leaders for a meeting to continue progressing the open source software security initiatives discussed during [January's White House Summit on Open Source Security](#). During this meeting, Google announced the creation of its new ["Open Source Maintenance Crew"](#) — a dedicated staff of Google engineers who will work closely with upstream maintainers on improving the security of critical open source projects.

In addition, to help organizations Google will be introducing Google Cloud's [Assured Open Source Software service](#). Assured OSS allows enterprise and public sector users of open source software to directly benefit from the in-depth, end-to-end security capabilities and practices we apply to our own OSS portfolio by providing access to the same OSS packages that Google depends on. Users will also be able to submit packages from their own OSS portfolio to be secured and managed through the Google Cloud managed service.

Packages curated by the Assured OSS service:

- are regularly scanned, analyzed, and fuzz-tested for vulnerabilities
- have corresponding enriched metadata incorporating [Container/Artifact Analysis](#) data
- are built with [Cloud Build](#) including evidence of verifiable SLSA-compliance
- are verifiably signed by Google
- are distributed from an [Artifact Registry](#) secured and protected by Google

As a result, Assured OSS lets organizations benefit from Google's extensive security experience and can reduce their need to develop, maintain, and operate complex processes to secure their open source dependencies. Assured OSS is expected to enter Preview in Q3 2022.

06 Cloud Security Checklist: When was your last checkup?

One of the most common questions Cloud customers ask us is “What should I be doing operationally, day to day, to address security risk?”

While there are general strategies and recommendations that can address large enterprise risks, it is often the smaller operational checks that can have the largest impact on the overall cloud risk of an organization.

Below we provide a short checklist of some operational Cloud defensive controls, strategies and considerations that will help security leaders and cloud teams more effectively address the risks and attack trends. While not comprehensive, this list should provide ways of taking immediate action to increase security posture.

While we provide Google Cloud specific mitigations at the end of this section, these questions for CISOs and security teams to consider are universal and are not specific to Google Cloud. The Cloud Security Checklist will be an ongoing series in the Threat Horizons report, focusing on often overlooked controls & methods of reducing risk.

☐	Does your organization mandate the use of an enterprise-wide browser, and enforces a version standard?
Why do you ask?	Many organizations do not enforce the use of one common enterprise-wide browser leading to insecure and risky browsing sessions by users There has been a record number of browser-based vulnerabilities found in the past two years; outdated browsers pose a significant risk to organizations. Google’s Work Safer Program enables multi-layered defenses, including Chrome Enterprise which offers ○ Continuous and automatic updates ○ Centralized policy management to reduce risk across endpoints ○ Increased end-user warnings and safe browsing capabilities ○ Mandatory patching & updates ○ Protect your data with site isolation

☐	Does your organization perform periodic recertification of critical admin/legal/support/security contacts, and work with vendors & CSPs to ensure they're accurately reflected in their systems?
Why do you ask?	A common trend that has been observed is a lack of long term maintenance around critical account contacts - ensuring that leavers are promptly replaced so that security alerts can be reviewed in a timely manner.

☐	Does your organization enforce strong access control policies?
Why do you ask?	Many organizations do not consistently enforce strong access policies across their devices and environments. This can lead to increased impact during a compromise. Two-factor authentication and solutions such as Context Aware Access can significantly reduce the blast radius during an attack.

☐	Do your security operations teams test triggering & receipt of alerts? Work with your vendors and CSP account teams if needed to run synthetic tests to ensure critical notifications/alerts reach their intended audience.
Why do you ask?	During an incident it is critical to have internal playbooks to define key decision makers and points of contact. Conduct tabletop drills (bi-annual to start) to ensure teams know who to reach out to and what the process for communication/escalation is. Is there a number to call upon discovery of an incident? A console to submit a ticket to? An email DL that should be cc'd? Assure key stakeholders receive notifications in order to ensure important information from Google Cloud reaches the right people. Google provides monitoring & alerting functionality, such as the detection of leaked credentials. Timely receipt and action of such alerts are crucial to minimizing potential impact.

☐	Have you recently reviewed your application/system/service logging to ensure it is consistently and sufficiently enabled?
Why do you ask?	When a security incident arises, the first place your Incident Response team will want to look is at relevant log data. Ensuring both enablement and familiarity with log sources where critical services & data sit can dramatically reduce overall time to respond. Having security champions embedded in your SWE teams can help reinforce consistent practices in both logging critical security events, and assuring that privacy of your data is protected by not over-logging sensitive fields (consider a DLP product to detect any sensitive data leaking into logs).

07 Zero Trust: Foundational Lessons and Misconceptions

As more employees return to the office in 2022, securing a hybrid workforce that blends lessons from before and after the pandemic continues to be critical.


As the world eases back to work in the offices, organizations continue to allow employees to use their own devices (BYOD) to access corporate resources. In addition, they are faced with the challenge of allowing the extended workforce (3rd parties contractors, customer support agents, and others) to access corporate resources from non-corporate devices that are unmanaged. With the increase in complexity comes the need for access controls to be context and content aware, and for the securing of data not to be at the authentication layer but at the authorization layer without adding friction.

Zero Trust (and [BeyondCorp](#)) models were created in response to the threats and attacks Google experienced directly. As part of our [“Shared Fate” model](#), it has become critical to providing our customers and partners with the lessons learned and tools to increase their security via the Zero Trust model. Key zero trust principles of BeyondCorp include

1. No privileged networks – do not rely on the network location to be the primary factor to determine trust. This is the case for most VPN-based security models.
2. Understand identities and devices – instead of relying on privileged networks, authorization should be based on deep understanding of the user identities and the devices they are using.
3. Continuous authorization – instead of authorizing only once at login time, to the extent possible, continuously validate the context of the identities and devices for every request.

Access and data inventory: It's 10pm, do you know who's accessing your data?

One of the most common challenges organizations face (and Google faced during its transition) is the importance of access models that consider data and device inventories. In order for Zero Trust to work, there must be a continuous authorization model that can determine **who** is authorized to access **what** data using **which** devices. This requires an accurate account of all users and devices that need to access an organization's resources and data, and also having a proper accounting of the location and sensitivities of your data and resources. Therefore creating an accurate inventory of data, users and devices is a critical component of creating a Zero Trust Architecture.



As previously mentioned, having a deep understanding of the devices that are accessing corporate resources is one of the key principles of zero trust. A common pitfall is that many assume that simply having an inventory of devices in a CMDB is sufficient to enable zero trust. Creating an accurate and authoritative device inventory is an iterative process. Organizations, especially in this hybrid work environment, must also consider

1. Unmanaged devices that are used by employees (BYOD) or the extended workforce must also be accounted for. Traditional CMDBs typically only have company-owned assets.
2. Device security postures must be collected continuously so access policies can be constructed based on the contextual knowledge. Traditional access control policies are mostly network-based and do not consider device context.
3. Device and user risks must be analyzed continuously so that access authorization can be allowed or denied based on the latest information. Traditional solutions typically do not take risk analysis into consideration for access control.

08 Going Deeper: Log Generation, Collection, and Managing Costs

Logging is an integral part of an organization's security posture, particularly in regards to detection and incident response. Logging failures were identified in the [OWASP Top 10 for 2021](#) among the most critical security risks. There are recent additional legal and regulatory requirements such as [Executive Order 14028](#) - sec 8 that compel government agencies to implement standardized logging, and provide a roadmap for the private sector to follow.



Many organizations understand the core ideas behind logging and enable logging for their critical systems. However with a deeper look most would likely agree there are opportunities to improve the effective analysis and actionable nature of the logs they collect.

To address this gap and help your organization think of its logging strategy, this series, released over the next several issues, will cover the lifecycle of logging, monitoring, and alerting on GCP. Google Cloud's Operation Suite integrates logging and monitoring for services on GCP and beyond with the logging API. GCP's Security Command Center provides a centralized vulnerability and threat reporting service, which paired with logging will help strengthen security. In this first installment, we'll focus on log enablement highlighting **log generation, collection, and how to manage costs**.

The National Institute of Standards and Technology ([NIST publication 800-92](#) sec 2.4) outlines key practices organizations can follow to navigate challenges in log management:

- Prioritize log management appropriately throughout the organization.
- Establish policies and procedures for log management.
- Create and maintain a secure log management infrastructure.
- Provide adequate support for all staff with log management responsibilities.

In this issue, we focus on the following [basic concepts](#): log entries, logs, and log types. Log entries record events and often include a timestamp, the monitored resource, the message, and the name of the log. Logs are a collection of log entries. Logging can be found in all layers of an organization's hierarchy such as at the resource, project, folder, and organization level. [Centralizing these logs](#) across your organization is not only critical to investigations, but will also make it easier for you to monitor your Cloud environment.



Some security logs are enabled by default on GCP and cannot be modified or disabled, while other logs require an organization to enable, configure, and fine tune. There are various types of logs on GCP and can be categorized as:

Google Cloud platform logs	User-written logs
VPC flow , Firewall rules , Cloud NAT gateways , and Load balancer	Written to Cloud Logging using the logging agent , the Cloud Logging API , or the Cloud .

More platform logs [found in this table](#).

Security logs	Multi-cloud and hybrid-cloud logs
Audit logs: • Admin Activity • Data Access • System Event • Policy Denied • Access Transparency logs	Supports: - Multi-cloud, ingesting logs from other cloud service providers - Hybrid clouds, integrating your on-premise infrastructure and apps.

GCP's Cloud Logging ingests audit logs and platform logs by default. [Audit logs](#) help answer the question of "Who did what, where, and when?" Among the audit logs: Admin Activity, System Event, and Policy Denied audit logs are enabled by default and cannot be configured, or disabled. Data access logs are disabled by default due to high volumes and personally identifiable information potentially being present in the logs, and require configuration. It is recommended to explicitly enable VPC flow logs platform logs and to take the [needed steps to enable Data Access Audit logs](#) to aid your organization with forensics and real-time security analysis.

You can collect telemetry from your Google Compute Engine instances by installing the [Ops Agent](#). When operating a multi-cloud model, you can use the Cloud Logging agent to collect telemetry from your Amazon Elastic Compute instances. These agents can be deployed at scale with automation tools such as Terraform, Ansible, Chef, Puppet, and Agent Policy using the gcloud CLI.

When creating custom application logs, use [structured logging](#) to significantly simplify searching and querying for logs as this nested structure integrates with logging tools. Additionally if you have embedded security champions within your organization, discuss opportunities with them to enrich your custom application logs to best serve your needs for monitoring and response.

It is important to ensure that critical & actionable logs are prioritized to help ensure a balanced signal-to-noise ratio and minimize costs around ingestion and storage. An example from Google's SRE book [Building Secure and Reliable Systems](#) chapter 15: Investigating Systems notes that firewalls routinely block many packets, many of which are harmless and may not be worth paying attention to. Over-enablement of such logs can not only unnecessarily increase cost but also make investigations more challenging by introducing noise along with alert fatigue for analysts.

Cloud Billing Reports allow you to filter on a project and the Cloud Logging service to see how much logging charges have been incurred and the metrics explorer in Cloud Monitoring will allow you to see the volume of logs ingested by your project. One method of reducing logging costs includes using filters such as log [exclusion filters](#) or [sampling flow logs](#)

In summary, we discussed various logs that are available on the platform, strategies for collecting relevant logs, and suggestions on how to manage your logging costs. In future reports we'll cover other aspects of the logging, monitoring, and alerting lifecycle and how it can be used to secure your organization.

09 Reminders and Recommendations

Google Cloud continues to operate with a “shared fate” model that exemplifies a true partnership with its customers. We will continue to provide our customers and the general public with reminders of some of the best practices they can implement to secure their environments. Many of these recommendations are straightforward and have been covered in other whitepapers and industry publications in greater detail, but are good reminders nevertheless for everyone.

Threat	Recommended Countermeasures
Spear-phishing	- Engage in email best practices. - Employ 2-Step Verification. - Enroll in the Advanced Protection Program. - Use Google's Work Safer and BeyondCorp Enterprise. - Deploy Context-Aware Access.
Cloud instance vulnerabilities	- Follow password best practices and best practices for configuring Cloud environments. - Update third-party software prior to a Cloud instance being exposed to the web. - Avoid publishing credentials in GitHub projects. - Use Container Analysis to perform vulnerability scanning and metadata storage. - Leverage Web Security Scanner in the Security Command Center to identify security vulnerabilities in App Engine, Google Kubernetes Engine, and Compute Engine. - Use service accounts with Compute Engine to authenticate apps instead of using user credentials. - Implement Policy Intelligence tools to help understand and manage policies. - Use predefined configurations through Assured Workloads to reduce misconfigurations. - Set up conditional alerts in the Cloud Console to send alerts upon high resource consumption. Enforce and monitor password requirements for users through the Google Admin console.
Downloading software updates	- Establish a strong chain of custody by hashing, verifying and security testing software downloads. - Google has also published the SLSA framework, which is complimented by the Software Bill Of Materials (SBOM)--both of which help organizations secure

	different parts of the software supply chain.
Ensure sensitive credentials are not in source code	• Audit projects published on GitHub and other sites to ensure credentials and certificates were not included. Note that with Google's current partnership with GitHub, repositories are automatically scanned for secrets, and customer are notified if found. • Scan code as part of a CI/CD pipeline, or perform code reviews to look for hardcoded keys or other credentials. • Provide more awareness to developers on proper use of such credentials during development. • If credentials are leaked, follow GCP best practices to recover.

¹ Stolyarov, Vlad, and Benoit Sevens. *Exposing initial access broker with ties to Conti*. 17 March 2022.

Google Threat Analysis Group [web blog],

<https://blog.google/threat-analysis-group/exposing-initial-access-broker-ties-conti/>