

Threat Horizons

Cloud Threat Intelligence
November 2021. Issue 1



Office of the CISO

Providing threat intelligence to those in the Cloud

Part of offering a secure cloud computing platform is providing cloud users with cybersecurity threat intelligence so they can better configure their environments and defenses in manners most specific to their needs. Office of the CISO is pleased to publish the first issue of *Threat Horizons* report. The report is based on threat intelligence observations from the Threat Analysis Group (TAG), Google Cloud Threat Intelligence for Chronicle, Trust and Safety, and other internal teams. It provides actionable intelligence that enables organizations to ensure their cloud environments are best protected against ever evolving threats. In this and future threat intelligence reports, Google will provide threat horizon scanning, trend tracking, and *Early Warning* announcements about emerging threats requiring immediate action.

Summary of Observations

While cloud customers continue to face a variety of threats across applications and infrastructure, many successful attacks are due to poor hygiene and a lack of basic control implementation. Most recently, our team has responded to cryptocurrency mining abuse, phishing campaigns, and ransomware. Given these specific observations and general threats, organizations that put emphasis on secure implementation, monitoring and ongoing assurance will be more successful in mitigating these threats or at the very least reduce their overall impact.

Spear-phishing and phishing campaigns are not new to the threat landscape; TAG observed recent attacks that targeted Gmail accounts and impersonated employment recruiters with the goal of stealing user credentials. Attackers also continue to exploit poorly configured Cloud instances with the goal of obtaining profit through cryptocurrency mining and traffic pumping. The universe of ransomware also continues to expand with the discovery of some new ransomware that appears to be offshoots of existing malware with mixed capabilities.

01	02	03	04	05	06
Compromised Google Cloud instances used for crypto-currency mining	Russia group launched Gmail phishing campaign	Fraudsters employ new TTP to abuse Cloud resources	North Korea actors impersonate employment recruiters	Black Matter ransomware rises out of DarkSide	Lessons Learned
Some poorly configured GCP instances are compromised quickly and used for cryptocurrency mining and other malicious activity.	Attackers, who typically targeted Yahoo! users, launched a campaign against Gmail accounts.	Fraudsters sought to abuse Cloud resources to generate traffic to YouTube.	Attackers impersonated employment recruiters in an attempt to steal credentials.	Black Matter ransomware found to be formidable; however, it does not exfiltrate data.	As Google Cloud partners with its customers in a shared fate security model, some valuable trends and lessons-learned emerge.

While a variety of threats exist, there are a number of ways to mitigate them as well, which includes [Container Analysis](#) to perform vulnerability scanning and metadata storage for containers and the [Web Security Scanner](#) in the [Security Command Center](#) to identify security vulnerabilities. Additionally, Google Cloud customers should [employ two-factor authentication](#), enroll in the [Advanced Protection Program](#), whenever possible, and use [Google's Work Safer](#), which provides companies with access to best-in-class security for email, meetings, messages, documents, and more. Work Safer brings together cloud-native, zero-trust solutions of Google Workspace with [BeyondCorp Enterprise](#).

Detailed Observations

Compromised GCP instances used for cryptocurrency mining

Threat Description / TTPs

Malicious actors were observed performing cryptocurrency mining within compromised Cloud instances. Of 50 recently compromised GCP instances, **86% of the compromised Google Cloud instances were used to perform cryptocurrency mining, a cloud resource-intensive for-profit activity**, which typically consumed CPU/GPU resources, or in cases of Chia mining, storage space. Additionally, 10% of compromised Cloud instances were used to conduct scans of other publicly available resources on the Internet to identify vulnerable systems, and 8% of instances were used to attack other targets. Table 1 lists the outcomes of compromised Google Cloud instances with an observation that, in some instances, multiple malicious actions were performed from within a single compromised instance. While data theft did not appear to be the objective of these compromises, it remains a risk associated with the cloud asset compromises as bad actors start performing multiple forms of abuse. The public Internet-facing Cloud instances were open to scanning and brute force attacks.

Table 1: Compromised GCP instances

Resultant actions after compromise	Percentage
Conduct cryptocurrency mining	86%
Conduct port scanning of other targets on the Internet	10%
Launch attacks against other targets on the Internet	8%
Host malware	6%
Host unauthorized content on the Internet	4%
Launch DDoS bot	2%
Send spam	2%

Note: Totals do not add up to 100% as some compromised instances were used to perform multiple malicious activities.

Malicious actors gained access to the Google Cloud instances by taking advantage of poor customer security practices or vulnerable third-party software in nearly 75% of all cases. As shown in Table 2, 48% of compromised instances were attributed to actors gaining access to the Internet-facing Cloud instance, which had either no password or a weak password for user accounts or API connections. As a result, these Google Cloud instances could be easily scanned and brute forced. 26% of compromised instances were attributed to vulnerabilities in third-party software, which was installed by the owner.

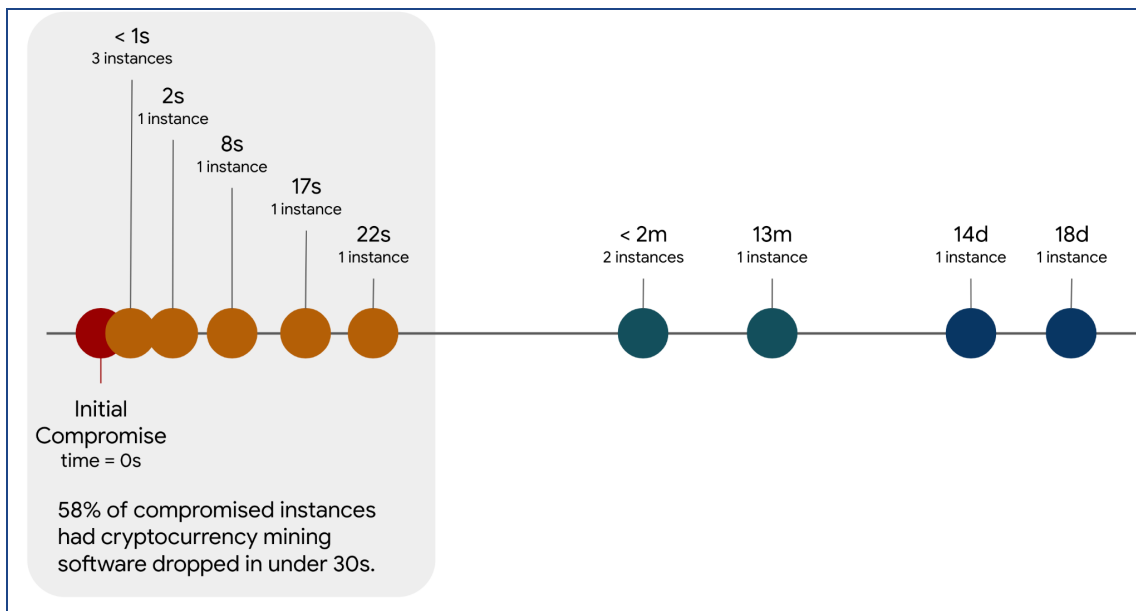
Table 2: Exploited vulnerabilities in Cloud instances

Exploited vulnerabilities	Percentage
Weak or no password for user account or no authentication for APIs	48%
Vulnerability in third-party software in the Cloud instance was exploited	26%
Other issues	12%
Misconfiguration of Cloud instance or in third-party software	12%
Leaked credentials, e.g., keys published in GitHub projects	4%

Time was of the essence in the compromise of the Google Cloud instances. **The shortest amount of time between deploying a vulnerable Cloud instance exposed to the Internet and its compromise was determined to be as little as 30 minutes.** In 40% of instances the time to compromise was under eight hours. This suggests that the public IP address space is routinely scanned for vulnerable Cloud instances. It will not be a matter of if a vulnerable Cloud instance is detected, but rather when.

Analysis of the systems used to perform unauthorized cryptocurrency mining, where timeline information was available, revealed that **in 58% of situations the cryptocurrency mining software was downloaded to the system within 22 seconds of being compromised** as shown in Figure 1. This suggests that the initial attacks and subsequent downloads were scripted events not requiring human intervention. The ability to manually intervene in these situations to prevent exploitation is nearly impossible. The best defense would be to not deploy a vulnerable system or have automated response mechanisms.

Figure 1: Time between initial compromise and download of cryptocurrency mining software



Strategic Significance

Google Cloud customers who stand up non-secure Cloud instances will likely be detected and attacked in a relatively short period of time. Given that most instances were used for cryptocurrency mining rather than exfiltration of data, Google analysts concluded the Google Cloud IP address range was scanned rather than particular Google Cloud customers being targeted. The amount of time from the launch of a vulnerable Google Cloud instance until compromise varied with the shortest amount of time being under 30 minutes.

Google Cloud Specific Mitigations

Aside from the [best practices](#) of ensuring accounts always have strong passwords, updating third-party software prior to a cloud instance being exposed to the web, and not publishing credentials in GitHub projects, Google customers have several different options to help mitigate risks.

Google Cloud customers can use [Container Analysis](#) to perform vulnerability scanning and metadata storage for containers and the [Web Security Scanner](#) in the [Security Command Center](#) to identify security vulnerabilities in their App Engine, Google Kubernetes Engine, and Compute Engine web applications. The scanner will crawl applications, following all links within the scope of the starting URL and attempt to exercise as many user inputs and event handlers as possible.

In addition to the Web Security Scanner, Google Cloud customers have additional resources including:

- A variety of [access control](#) options within Compute Engine including using [service accounts](#) to authenticate apps instead of using user credentials.
- [Policy Intelligence tools](#) to help understand and manage policies to proactively improve security configurations.

- Pre-defined configurations through [Assured Workloads](#) to reduce the risk of accidental misconfigurations by choosing from available platform security configurations, controls can be put in place.
- [Conditional alerts](#) in the Cloud Console to determine when resource consumption exceeds certain thresholds.
- [Enforcing and monitoring password requirements for their users](#) through the Google Admin console.
- [Recommendations](#) for designing online applications with a password-based authentication system.
- [Best practices](#) for configuring Cloud environments.

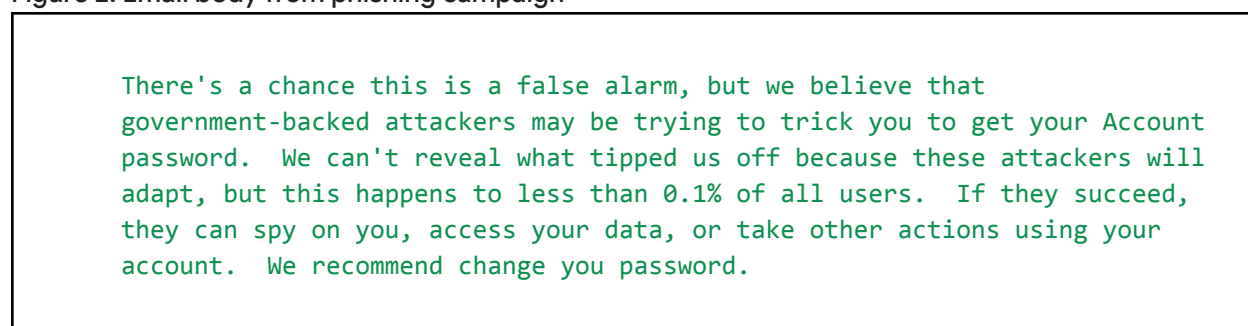
APT28/Fancy Bear launched Gmail phishing campaign

Threat Description / TTPs

Based on research from TAG, the Russian government-backed attackers APT28 / Fancy Bear, which more recently has typically targeted Yahoo! and Microsoft users, was observed at the end of September sending a large-scale attack to approximately 12K+ Gmail accounts in a credential phishing campaign. Google blocked these messages and no users were compromised.

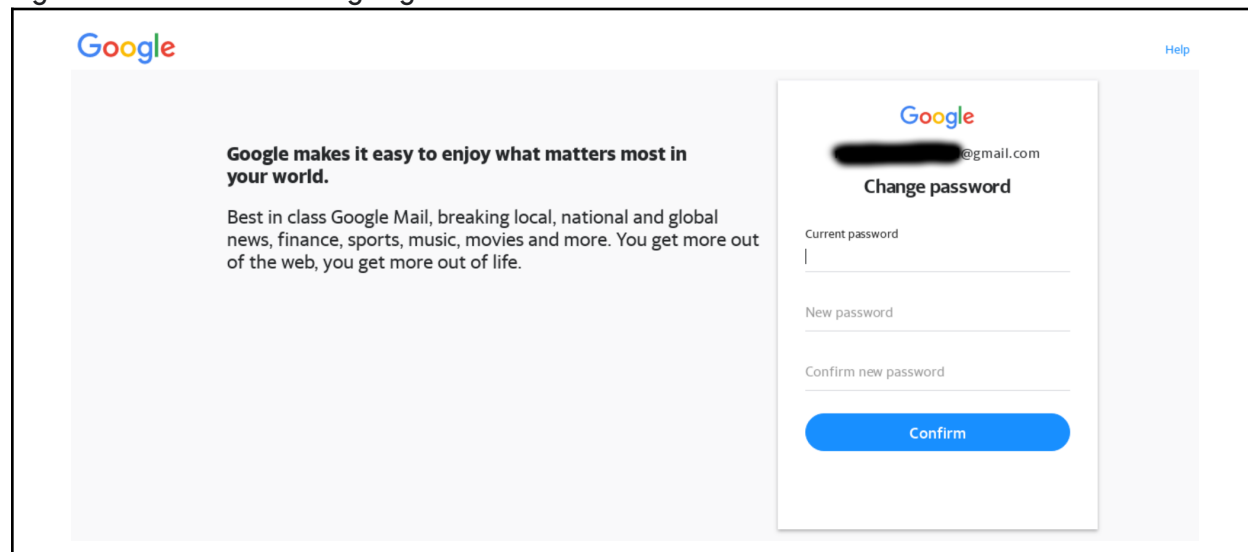
The attackers were using patterns similar to TAG's [government-backed attack alerts](#) to lure users to change their credentials on the attacker's controlled phishing page. The attackers kept changing the emails' subject line but attackers used a variation of **Critical security alert**. The email body contained the information shown in Figure 2.

Figure 2: Email body from phishing campaign



The credential phishing page, which is shown in Figure 3, appeared similar to a Gmail login page.

Figure 3: Credential Phishing Page



Upon closer inspection, TAG observed that the fonts in the phishing page did not match the fonts on the legitimate Google owned page. This was because the attackers tried to reuse their Yahoo! toolkit

and left various Yahoo! artifacts in the Gmail HTML login page, including a specific version of Yahoo's CSS which uses a different font making the page look slightly different. This is shown in Figure 4.

Figure 4: Artifacts from login page

URL
http://digital-email-software.great-site.net/? [REDACTED]
https://code.jquery.com/jquery-3.5.0.js
https://s.yimg.com/wm/mbr/6fc59edacb070b34e5c18360671fbbd732b90954/yahoo-main.css
https://s.yimg.com/wm/mbr/images/capslock-v0.0.2.svg
https://s.yimg.com/wm/modern/images/fuji-spinner-1.0.1.svg
https://s.yimg.com/wm/mbr/6fc59edacb070b34e5c18360671fbbd732b90954/bundle.js
https://www.gstatic.com/images/branding/googlelogo/svg/googlelogo_clr_74x24px.svg
https://s.yimg.com/cv/ae/sports/fonts/2017/Yahoo_Sans-Regular.woff2
https://s.yimg.com/cv/ae/sports/fonts/2017/Yahoo_Sans-Semibold.woff2
https://s.yimg.com/cv/ae/sports/fonts/2017/Yahoo_Sans-Medium.woff2
https://s.yimg.com/cv/ae/sports/fonts/2017/Yahoo_Sans-Bold.woff2

The phishing messages were sent from what appeared to be compromised mail servers. This was a change from the previous Yahoo! campaigns, which predominantly used some variant of spoofing to send emails. In the Gmail campaigns, the majority of messages passed SPF and even the ones that did not appeared to be misconfigured servers rather than the usual spoofing. Table 3 displays the SPF breakdown for attacker messages.

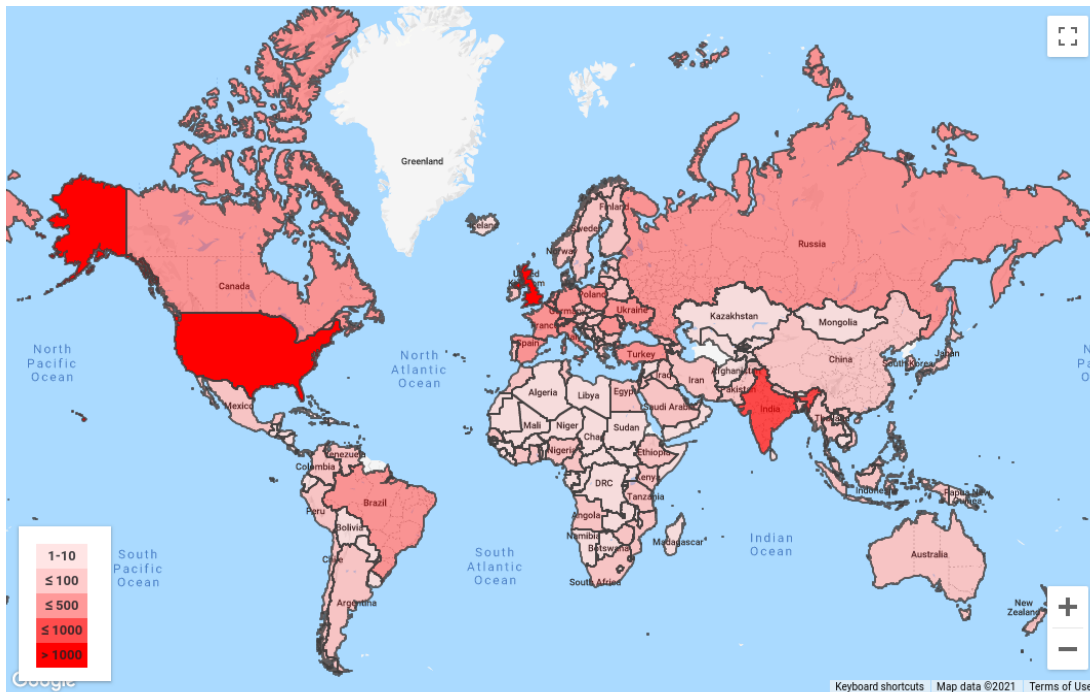
Table 3: SPF breakdown

SpfResultCode	SpfResultLayer	cnt
SPF_RESULT_PASS	2	232
SPF_RESULT_PASS	3	248
SPF_RESULT_PASS	0	8923
SPF_RESULT_FAIL	4	7
SPF_RESULT_FAIL	0	461
SPF_RESULT_SOFTFAIL	4	9
SPF_RESULT_SOFTFAIL	0	1192
SPF_RESULT_NEUTRAL	0	62
SPF_RESULT_NEUTRAL	3	2444
SPF_RESULT_UNKNOWN	4	2
SPF_RESULT_ERROR	0	12
SPF_RESULT_NONE	4	34
SPF_RESULT_PERMERROR	0	63
SPF_RESULT_NOCHECK	2	1

One significant difference between legitimate emails from the compromised mail servers and phishing messages was the domain part of Messageld which is different and unique for every email address domain.

Highly targeted regions for this particular campaign include the United States, United Kingdom, and India. Other noteworthy regions include Canada, Russia, Brazil, and members of the European Union. Figure 5 depicts the targets for this particular campaign.

Figure 5: Heatmap of targets for this particular campaign



Strategic Significance

Phishing and spear phishing campaigns continue to use login pages that impersonate legitimate Google login pages to steal credentials. While many fake pages may appear similar to Google's own page, a close inspection of URLs, certificates, fonts, and graphics will identify discrepancies.

Google Cloud Specific Mitigations

As with all phishing and spear phishing threats, Google customers should engage in [best practices](#). Google [Safe Browsing](#) will help secure many users; however, Workspace customers and Gmail users should validate that they are providing credentials to legitimate Google sites, [employ two-factor authentication](#), and enroll in the [Advanced Protection Program](#), whenever possible.

Customers may also use [Google's Work Safer](#), which provides companies with access to best-in-class security for email, meetings, messages, documents, and more. Work Safer brings together Cloud-native, zero-trust solutions of Google Workspace with [BeyondCorp Enterprise](#) for secure access with integrated threat and data protection.

Fraudsters employ new TTP to abuse Cloud resources

Based on research from TAG, a group of attackers were observed abusing Cloud resources to generate traffic to YouTube for view count manipulation. Attackers have been using various approaches to gain free Cloud credits, including using free trial projects, abusing start up credits with fake companies, and joining Google Developer Community for free projects. A group of attacker's TTP for exploiting payment was discovered recently by Google Cloud's abuse team. Attackers were able to use free credits by making small credit card payments and declining the payment afterward. Upon cloud abuse enforcement, the attacker quickly switched to Qwiklab projects and the Google Cloud abuse team pivoted to counter this offensive.

Strategic Significance

Attackers have continued to exploit Google Cloud projects where free credits were provided to engage in traffic pumping to YouTube, and there is a likelihood that attackers will continue to exploit Cloud instances for the same purpose. Attackers, who gain access to legitimate Cloud instances, will exploit the platforms for various forms of financial gain. Cloud customers who are not mindful of the consumption of their Cloud resources could be unwittingly abused.

Google Cloud Specific Mitigations

Aside from [best practices](#) of ensuring accounts always have strong passwords, updating third-party software prior to a Cloud instance being exposed to the web, and not publishing credentials in GitHub projects, Google Cloud customers have several different options to help mitigate risks.

Google Cloud customers can use [Container Analysis](#) to perform vulnerability scanning and metadata storage for containers and the [Web Security Scanner](#) in the [Security Command Center](#) to identify security vulnerabilities in their App Engine, Google Kubernetes Engine, and Compute Engine web applications. The scanner will crawl applications, following all links within the scope of the starting URL and attempt to exercise as many user input and event handlers as possible.

In addition to the Web Security Scanner, Google Cloud customers have additional resources including:

- A variety of [access control](#) options within Compute Engine including using [service accounts](#) to authenticate apps instead of using user credentials.
- [Policy Intelligence tools](#) to help understand and manage policies to proactively improve security configurations.
- Pre-defined configurations through [Assured Workloads](#) to reduce the risk of accidental misconfigurations by choosing from available platform security configurations—we'll help put the controls in place.
- [Conditional alerts](#) in the Cloud Console to determine when resource consumption exceeds certain thresholds.
- Tools to [enforce and monitor password requirements for their users](#) through the Google Admin console.
- These [recommendations](#) for designing online applications with a password-based authentication system.
- [Best practices](#) for configuring Cloud environments.

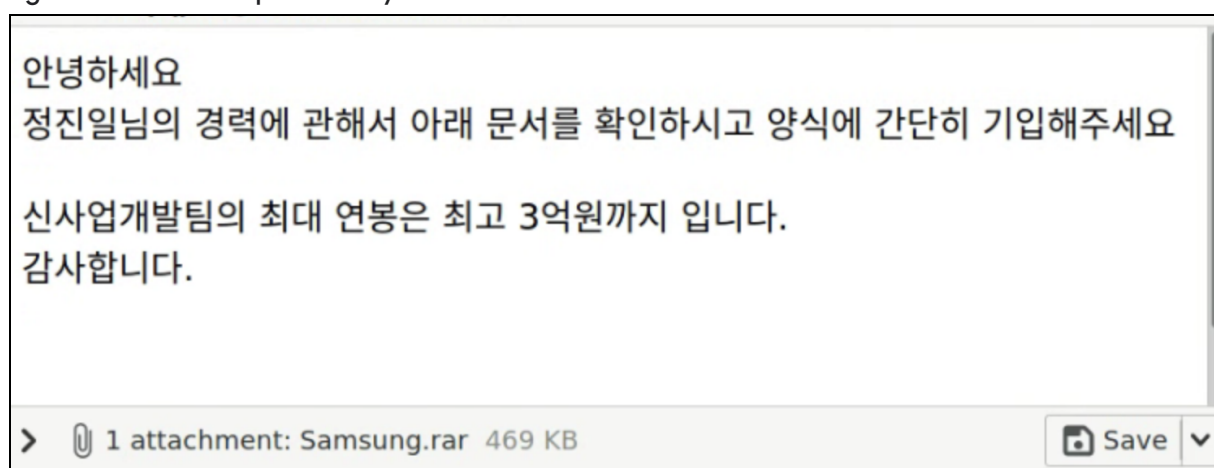
North Korea targets users by posing as employment recruiters

Threat Description / TTPs

TAG observed a North Korean government-backed attacker group that previously [targeted security researchers](#) posing as recruiters at Samsung and sending fake job opportunities to employees at multiple South Korean information security companies that sell anti-malware solutions.

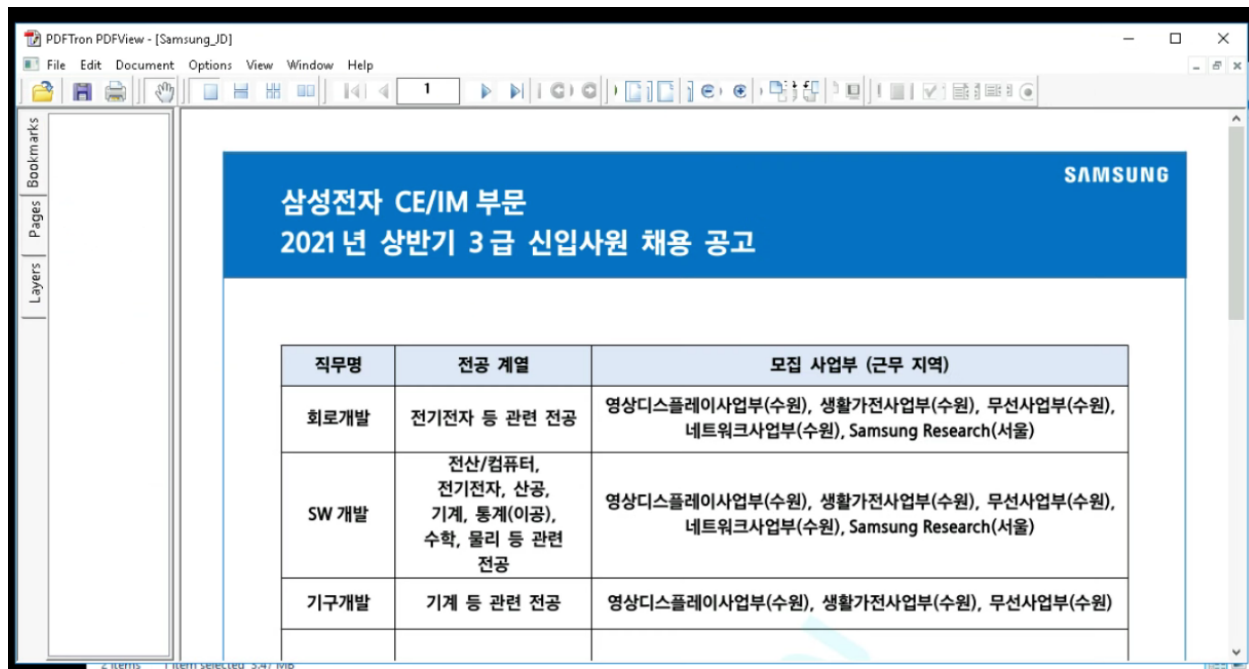
The emails included a PDF allegedly claiming to be of a job description for a role at Samsung; however, the PDFs were malformed and did not open in a standard PDF reader. When targets replied that they could not open the job description, attackers responded with a malicious link to malware purporting to be a “Secure PDF Reader” stored in Google Drive which has now been blocked. An example of the initial email can be seen in Figure 6.

Figure 6: Email example used by the attackers



The Secure PDF Reader was a modified version of [PDFTron](#) which attempted to decode an embedded Portable Executable (PE) and PDF from a supplied PDF. The PE was XOR encoded with a single-byte key and in this case drops an implant, which uses a legit but exploited South Korean website for Command and Control (C2) and affords the attackers various capabilities like being able to execute arbitrary commands and upload files. A screenshot of the modified PDFTron reader after opening one of the malformed PDFs can be seen in Figure 7.

Figure 7: Modified PDFTron reader used in campaign



This is not the first time this attacker has modified a PDF viewer. A modified version of SumatraPDF was used last year, which used the hash of the malformed PDF as a key to decrypt and drop an implant and legitimate PE that were embedded within the viewer itself. Other groups were recently seen using a similar technique of providing a malicious PDF viewer to view malformed PDFs.

Strategic Significance

PDFs and associated viewers remain an attack tactic used by various groups. Social media postings on sites such as LinkedIn continue to be a source of information for technical professionals. Phishing and spear phishing campaigns continue to use login pages that impersonate legitimate Google login pages to steal credentials. While many fake pages may appear similar to Google's own page, a close inspection of URLs, certificates, fonts, and graphics will identify discrepancies.

Google Cloud Specific Mitigations

As with all phishing and spear phishing threats, Google Cloud customers should engage in [best practices](#). Additionally, Workspace customers and Gmail users should validate that they are providing credentials to legitimate Google sites, [employ two-factor authentication](#), and enroll in the [Advanced Protection Program](#), whenever possible.

While [Google Safe Browsing](#) can provide one layer of defense, customers may also use [Google's Work Safer](#), which provides companies with access to best-in-class security for email, meetings, messages, documents, and more. Work Safer brings together Cloud-native, zero-trust solutions of Google Workspace with [BeyondCorp Enterprise](#) for secure access with integrated threat and data protection.

Black Matter ransomware rises out of DarkSide

Threat Description / TTPs

Based on research from Google Cloud Threat Intelligence for Chronicle, Black Matter, the successor of DarkSide, is one of many ransomware families currently being used to extort money from victims by locking their files using encryption. While Black Matter is considered a relatively new player in this space, evidence suggests it is merely the immediate offspring of DarkSide.

At its core, Black Matter is a configurable, whole-system and network share encryption tool capable of encrypting files on a victim's hard drive in a relatively short period of time by distributing the workload across multiple threads. Like most ransomware, it is less about the core encryption mechanism that defines the malware and more about the ancillary support code that makes the malware novel.

Black Matter is highly configurable and allows for a very targeted deployment. The configuration contains the ability to whitelist files, directories, extensions and even entire computers based on their names. It is possible for an attacker to supply known credentials to the malware in order to allow the malware to access higher privileges and additional network resources.

Despite the warning given to victims in the ransom note dropped by the malware, Black Matter did not have the ability to exfiltrate data at the time this research was performed. The malware was only capable of reporting statistics of its operations, such as total number of files found, number of files not encrypted, duration of operation, etc., in addition to the details about the victim's computer, e.g., hostname, language, operating system, drive details, etc. This lack of exfiltration capability could very well indicate that the deployment of Black Matter on a victim's computer occurred after the attacker(s) had already gained access to the victim's infrastructure. This claim is further supported by the fact that the malware can be configured with target specific credentials and ransom notes that include very specific target server information.

The malware does not provide a significant barrier to analysis as it does not employ sophisticated anti-analysis techniques such as packers, cryptors, or code flow manipulation. The malware does use hashing to frustrate analysis as strings are not present - only their hash values are.

On the whole, Black Matter is a formidable ransomware family. As with any other ransomware, the use of heavy encryption makes recovery of files nearly impossible without paying for the decryption tool. The use of multiple threads and asynchronous operations allows the malware to quickly compromise the victim's local computer and networked resources, if configured. The malware also deletes all files within the victim's Recycle Bins on all attached drives before deleting all Shadow Copy files. This is an effective means of destroying backups of files that are soon to be encrypted. After all encryption operations have been completed, and the reporting configuration option is enabled, the malware reports statistics to its Command and Control (C2) server.

Strategic Significance

The presence of Black Matter ransomware on a network is an indication that a network has been compromised through another means. Given that the malware does not have the capability to upload files to a C2 server, it is highly likely that the attacker(s) had prior access to the victim's infrastructure prior to deploying the ransomware. Incident response teams should look for additional

indicators of compromise. Google has received reports that the Black Matter ransomware group has announced it will shut down operations given outside pressure. Until this is confirmed, Black Matter still poses a risk.

Google Cloud Specific Mitigations

As with all ransomware threats, Google customers should engage in [best practices](#) with respect to common attack vectors, e.g., email and phishing campaigns. Detection of Black Matter is largely unnecessary as the malware will make its existence known to the victim; however, being able to detect compromised computers that might not be normally touched by users, e.g., will be beneficial. Google customers may upload the YARA-L2 rules that appear in Table 4 into [Chronicle](#).

Table 4: YL2 rules to detect Black Matter ransomware

```
rule UC_ttp_BlackMatter__RegKeys {
  meta:
    author = "Google Cloud Threat Intelligence"
    description = "Known registry keys used by Black Matter"

  events:
    // Modifying the privacy settings screen settings
    ($e.principal.registry.registry_key = /software\policies\microsoft\windows\oobe/ nocase
and
    $e.principal.registry.registry_value_name = "disableprivacyexperience" nocase) or

    // Storing the screen's resolution in the registry
    ($e.principal.registry.registry_key = /SOFTWARE\[A-Za-z0-9]{8}/ and
    ($e.principal.registry.registry_value_name = /hScreen/ or
    $e.principal.registry.registry_value_name = /vScreen/ )) or

    // RunOnce key
    ($e.principal.registry.registry_key =
/SOFTWARE\Microsoft\Windows\CurrentVersion\RunOnce/ nocase and
    $e.principal.registry.registry_value_name = /[A-Z]{3}[0-9]{3}[a-z]{3}/ )

  condition:
    $e
}
```

```
rule UC_ttp_BlackMatter__SafeBoot {
  meta:
    author = "Google Cloud Threat Intelligence"
    description = "Detects a machine's configuration being changed to safe boot"
    ext_description = "Known command line for Black Matter's SafeBoot"

  events:
    ($e.principal.process.file.full_path = /bootcfg/ nocase and
    ($e.principal.process.command_line = /\raw \/a \safeboot:network \/id 1/ or
    ($e.principal.process.command_line = /\raw \fastdetect \/id 1/)) or
    ($e.principal.process.file.full_path = /bcdedit/ nocase and
    ($e.principal.process.command_line = /\raw \set \{current\} safeboot network/ or
    $e.principal.process.command_line = /\raw \deletevalue \{current\} safeboot/))

  condition:
    $e
}
```


Recommendations

Google Cloud continues to operate within a "[shared fate](#)" model that exemplifies a true partnership with its customers. This partnership includes providing trends and lessons learned from recent incidents or close-calls in the wild with which Google assisted. The following is a summary of Google Cloud's recommendations based upon incidents that it helped address:

Audit published projects to ensure certs and credentials are not accidentally exposed. Certs and credentials are mistakenly included in projects published on GitHub and other repositories on a regular basis. An audit of published projects can ensure that this mistake does not happen.

Code downloaded by clients should undergo hashing authentication. It is a common practice for clients to download updates and code from cloud resources, raising concern that unauthorized code may be downloaded in the process. Meddler in the Middle (MITM) attacks may cause unauthorized source code to be pulled into production. By hashing and verifying all downloads, the [integrity of the software supply chain](#) can be preserved and an effective [chain of custody](#) can be established.

Use multiple layers of defense to combat theft of credentials and authentication cookies. Cloud-hosted resources have the benefit of high availability and "anywhere, anytime" access. While this streamlines workforce operations, bad actors can try to take advantage of the ubiquitous nature of the cloud to compromise cloud resources. Despite the growing public attention to cybersecurity, spear-phishing and social engineering tactics are frequently successful. As for other forms of IT security, defensive measures need to be robust and layered to protect cloud resources due to ubiquitous access. In addition to [two-factor authentication](#), Cloud administrators should strengthen their environment through [Context-Aware Access](#) and solutions such as [BeyondCorp Enterprise](#) and [Work Safer](#).

Table 5: Observed risks and countermeasures

Risk	Countermeasures
Exploiting vulnerable GCP instances	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.

	<p>Set up conditional alerts in the Cloud Console to send alerts upon high resource consumption.</p> <p>Enforce and monitor password requirements for users through the Google Admin console.</p>
Spear-phishing	<p>Engage in email best practices.</p> <p>Employ 2-Step Verification.</p> <p>Enroll in the Advanced Protection Program..</p> <p>Use Google's Work Safer and BeyondCorp Enterprise.</p> <p>Deploy Context-Aware Access.</p>
Downloading software updates	<p>Establish a strong chain of custody by hashing and verifying downloads.</p>
Using public code repositories	<p>Audit projects published on GitHub and other sites to ensure credentials and certificates were not included.</p>

