




Threat Horizons

January 2023 Threat Horizons Report

Table of contents

Mission statement	03
Letter from the editor	
2023 predictions: Planning for the unexpected in cloud threats	04
Summary	
Initial-access vectors diversify, point toward possible increase in automation of attacks	07
Threat trends	
Malware communicating and hiding interactions with cloud providers' IP addresses and open ports	09
APT10: Lessons learned from studying government-backed cloud targeting	14
Threat groups probably developing methods to threaten operational technology deployed in the cloud	17
Backups increasingly targeted by threat actors	22
Use of cloud infrastructure to conduct DDoS attacks	24



Mission statement

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

Letter from the Editor

2023 predictions: Planning for the unexpected in cloud threats

One of the most important activities I encourage threat intelligence teams to do as they mature and grow is to start making predictions about what threats they expect their organization to face in the future. It's an important step to move from a reactive intelligence team supporting ongoing investigations and incidents to a proactive one that helps senior leaders in their organization prevent threats, understand the risks their organization is already facing, and plan strategically for the future.

None of us have a crystal ball, but the very act of formalizing threat predictions carries with it benefits beyond identifying potential threats that might come to fruition. It forces the team to think carefully about one's own organization and what resources and operations are most important to it; the team should notice trends in common factors among their predictions, possibly identifying ways to achieve positive security outcomes that cut across identified threats. It offers an opportunity for creativity that some team members may excel at, giving managers an opportunity to benefit from a professionally diverse team and giving team members a different way to contribute than more detail-oriented technical work. It keeps the team cognizant of nontechnical factors that could contribute to cyber risk, such as the geopolitical position of the company, macroeconomic and security trends, and changes to the organization's public profile over time.

Making the best predictions will involve brainstorming from a base of existing problems, the team's own ideas, and new ideas gleaned from a variety of sources, giving the team a chance to think about which sources they have and are most important in their work, and which new sources may need to be acquired or built. It also often involves cross-functional insight from other teams, such as the security operations center (SOC), incident responders, senior leaders, business unit and sales managers, and individual contributors with geopolitical and other analytic skills, building ties to the rest of the organization.

As former US President Dwight D. Eisenhower once said about his life in the Army, "Plans are worthless, but planning is everything."

Even if none of your team's particular predictions come to fruition, the very act of gathering sources, teaming up, formalizing predictions, and tracking them throughout the following year can greatly improve efficacy, develop your organization's threat intelligence talent, keep the team mindful of potential blind spots, and ensure accountability. As former US President Dwight D. Eisenhower once said about his

(Letter from the editor, cont'd.)

life in the Army, “Plans are worthless, but planning is everything.”

It is in that spirit that we want to share a few of the intelligence-based Office of the CISO with during their brainstorm, which we will be tracking over the next year or more as we head into 2023:

- **Identity and trust** relationships in and between cloud environments will continue to get more complex, challenging visibility and enabling threat actors to have a wider and deeper impact on organizations. We anticipate an increase in targeting of identities that allow cross-platform authentication as threat actors recognize the value in compromising identities rather than endpoints. The Chinese Government group APT10’s [Cloud Hopper](#) campaign (see [page 14](#)) which pivoted from MSP access to [exploitation of VPN technology](#) and, more recently, the Russian Government group APT29’s compromises of [Microsoft 365](#) and similar cloud-hosted workplaces provide a template they and less sophisticated groups will follow. In 2023, we will be watching to see if there is at least one public incident of a threat actor gaining access to a customer environment at one Cloud Service Provider (CSP) and leveraging that into assets hosted on a different CSP due to a lack of identity verification controls, overly permissive trust architecture, or both.
- Threat actor use of one-off cloud-hosted instances will become increasingly harmful as threat actors generate more effective and potent uses of **short-term tenancy**. The top malware used by short-term infections will still be cryptominers in 2023, but other forms of monetization, such as phishing or ransoming customer environments, could grow as well.
- As cloud customer environments expand, **third-party assets**—software libraries, external data feeds, third-party tools, and so on—are being integrated within these environments. Given the cloud’s extended automation capabilities compared to on-premises settings, users can therefore benefit from such new capabilities faster. At the same time, such assets may be integrated faster than security teams can assess the risk to them, necessitating updates to risk-management processes—potentially with their own automation—to keep up. The risk of third-party dependencies will be an important issue in 2023. Given the high-profile success of public incidents like the one that affected SolarWinds, we predict at least one APT actor will use seemingly legitimate software updates to push malware to third-party systems in 2023, after having compromised a software provider.
- Organizations have increasingly **integrated OT systems with IT infrastructure**, including the use of cloud services, to scale production, develop efficiencies, and handle geographically distributed processes for critical infrastructure. Threat groups with experience targeting physical production networks are likely already planning how to compromise targets using cloud services. In 2023, we will likely see increasing discussion by threat actors on how to disrupt cloud services and resources that support OT production systems via denial-of-service (DoS) attacks—for example, moving across improperly segmented networks, deploying ransomware, or even developing customized malware. (See [page 09](#).)

(Letter from the editor, cont'd.)

- **Attacker tools and malware** are evolving to better target customer cloud environments specifically. As more companies move more things to the cloud and software-as-a-service (SaaS) providers and away from on-premises, more attacks will inevitably shift to target customers' cloud environments. Cloud providers continue to invest in defending themselves and partnering with customers to improve their defenses, but vigilance is needed to keep pace with evolving threats. We predict new and upgraded cloud-specific attack tools to start appearing in 2023. Cloud-focused malware will also be updated to more efficiently abuse cloud instances. We predict a ransomware strain that targets cloud-based backups, including revision history and cloud-stored backups in 2023.


(See [page 09](#).)

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


Summary

Initial-access vectors diversify, point toward possible increase in automation of attacks

Cloud Compromise Factors (Q3 2022)



Software Targeted in Cloud (Q3 2022)



In Q3 2022, analysis of data about Google Cloud customer compromises indicates that threat actors diversified their initial access vectors compared to what we saw earlier in the year. Weak passwords continued to be the most common factor at 41% of observed compromises. However, API key compromise played a role in nearly 20% of cases studied last quarter. In terms of which software was most targeted in Q3, we observed significant diversification. SSH was targeted in 26% of cases, but Jenkins and PostgreSQL were close behind at around 22% and 17%, respectively.

Increased diversification efforts by threat actors in targeting and access vectors highlights the constantly evolving threat landscape faced by organizations. In particular, the use of API compromise may suggest increased levels of automation by threat actors. While threat activity historically has dropped toward the end of Q4, the use of automation may allow actors to keep activity steady or even growing in volume into 2023.

¹ The following statistics were derived from the portion of compromises observed by our incident response teams, which will be skewed to the platform affected in these cases and may not be representative of all customer environments and verticals on Google Cloud.

Threat trends

Malware communicating and hiding interactions with cloud providers' IP addresses and open ports

Issue description

Examining new attack vectors against cloud environments, we searched VirusTotal (VT) for 2022 malware samples communicating with three geographic regions of large cloud service providers (CSPs). We found over 6,000 malware samples—dynamically analyzed by VT—communicating with the CSPs using many pre-specified or at times randomly selected IP addresses and TCP/IP ports. The malware also at times tried to hide its activities among legitimate services by communicating to CSPs using well-known ports, as well as by explicitly utilizing TLS. Cloud users should mitigate these types of malicious operations by monitoring and restricting inbound as well as internal Google Cloud network communications, using hardened VM images, and reviewing cloud instance audit events for unexpected administrative or user activities.

The malware samples we identified that communicated with CSPs utilized well-known ports—like 80, 23, and 443—more frequently compared to registered and

ephemeral ports (e.g., above 1023). This may be because CSPs and their customers more frequently open such ports to expose associated standard services, thereby providing open channels for misuse attempts. Malware, however, may target the registered or ephemeral ports when scanning for open TCP/IP ports. Malware may also target such ports to exploit various less-common cloud services. Malware, of course, may also simply be exploring if it has an internet connection.

We examined Q2-Q3 2022 VT data showing malware trying to contact either cloud services, customer-created workloads hosted by the three largest CSPs (Google Cloud, Azure, and AWS), or both. Using the CSPs' IP addresses representing the Canadian, German, and South American regions (chosen for geographic diversity and traffic differences, as well as the manageable number of malware samples to analyze), we looked at malware first submitted to VT

(Malware communicating and hiding, cont'd.)

CSP Protocol	CSP Protocol Total Overall Frequency		
	GCP	Azure	AWS
TCP	537	507	5670
UDP	1	0	24
ICMP	0	0	4
(unclassified protocol)	35	26	222

CSP Port	CSP Protocol	CSP Port/Protocol Frequency (in parenthesis, calculated as % of all Port Communications for given CSP)		
		GCP	Azure	AWS
80	TCP	309 (53.9%)	52 (9.8%)	1821 (30.8%)
443	TCP	86 (15.0%)	74 (13.9%)	86 (15.0%)
23	TCP	82 (14.3%)	267 (50.1%)	376 (6.4%)
15647	TCP	23 (4.0%)	(malware didn't contact this port in our data)	(malware didn't contact this port in our data)
445	TCP	15 (2.6%)	42 (7.9%)	156 (2.6%)
8080	TCP	2 (< 1%, not among the top contacted ports)	5 (1%, not among the top contacted ports)	203 (3.4%)
2323	TCP	1 (< 1%, not among the top contacted ports)	19 (3.6%)	15 (< 1%, not among the top contacted ports)
About 500 other ports in the combined GCP, Azure, and AWS port lists; with the bulk belonging to AWS. For all CSPs, most remaining ports greater than 1023	The remaining ports used mostly TCP across the CSPs; although some used UDP	Remaining GCP port/protocol frequencies ranged from 3-7 to 1-2; the latter being most prevalent	Remaining Azure port/protocol frequencies ranged from 3-14 to 1-2; the latter being most prevalent	Remaining AWS port/protocol frequencies ranged from 12-139 to 1-11; the latter being most prevalent

Illustration 1

over a four-month period, April through July 2022. Illustration 1 shows how well-known port numbers (below 1024) were targeted more frequently when compared to registered and ephemeral ports (above 1023) in our sample.

Additionally, malware was found attempting to communicate differently to the same well-known ports on different CSPs, as each CSP exposes different

services. For example, port 445 was relatively more popular for malware communicating with Azure than for malware communicating with the other CSPs. Port 445 is used for SMB communications for managing Windows machines and related file-sharing services. Windows-based cloud services probably make up a larger share of workloads on Microsoft-owned Azure than on the other CSPs.

(Malware communicating and hiding, cont'd.)

Threat actors may also be disguising their malware's activities. We discovered malware that was programmed to use well-known ports, otherwise leveraged for legitimate services, to "blend in."² Malware authors may also emphasize, or conversely not prioritize, the hiding of their malware's activities when trying to communicate with CSP IP addresses being utilized for malicious purposes. Often, when malware communicates with malicious IPs, it tries to protect its interactions from scrutiny using TLS. However, one-third of the time, other malware communicates with malicious IPs over the unprotected—and more "monitorable"—port 80. Such patterns may reflect the desire of some malware authors to protect communications and let their malware potentially persist, disguised, within cloud environments. Other authors, however, might be focused on immediate target compromise, rather than on long-term persistence and associated "detection evasion" techniques.³

- For example, consider well-known port 443, used for encrypted communications. From July through October 2022, Google observed some customer Google Cloud environments compromised by cryptominers via Google Cloud APIs. Customers accidentally leaked service account credentials to public code repos, like GitHub; after which, automated malicious scripts captured such credentials, authenticated to environment-managing Google Cloud APIs over port 443, and spun up new VMs, installing the cryptomining software within them.
- Our data shows that TLS is used to protect communications in almost half of cases when

malware is trying to communicate through common, well-known CSP ports to CSP IPs behind such ports.

- This overall analysis is also supported by other studies. A Sophos Q2 2021 report found that malware communicated with web and cloud services—such as GitHub, and similar cloud services—using TLS almost half the time.⁴ And this was a 100% increase over 2020, when only 23% of malware used TLS for such communications.

Malware attempts to communicate with registered and ephemeral ports too. Our data identified two live campaigns contacting the CSPs—one contacting all three CSPs via port 2323 and another contacting Google Cloud via port 15647. The campaign trying to contact the CSPs via port 2323 was targeting IPs that the malware dynamically generated. To understand malware communicating with the CSPs on port 2323, we selected from our dataset a sample of 4 out of 19 files communicating with Azure port 2323, the one file communicating with Google Cloud port 2323, and a sample of 3 out of 15 files communicating with AWS port 2323. To understand the malware communicating with Google Cloud port 15647, we examined a sample of 4 out of 23 files communicating with this port in our dataset. We also examined VT's large data corpus.

- The campaign contacting the three CSPs via port 2323 is the Mirai malware. This malware tries to compromise IoT devices to coordinate them for botnet attacks. The malware generates pseudo-random IPv4 addresses—contacting them via port 2323 to find potential, "compromisable"

² Attackers "blending in" by using well-known ports in externally facing on-premises, as well as general internet, environments is supported by other studies too—such as MITRE, "Commonly Used Port," MITRE ATT&CK, September 27, 2022, <https://attack.mitre.org/techniques/T0885/>, accessed December 9, 2022.

³ Note: It is possible that malware's common use of TLS and port 80 when communicating with malicious IPs may happen when the IPs represent general websites, rather than just CSP-specific IP ranges; and one reference, a 2021 Sophos study, describing some of these behaviors, is provided in the article. Nevertheless, we did not investigate the full range of such behaviors in our data.

⁴ Jai Vijayan, "Nearly Half of All Malware Is Concealed in TLS-Encrypted Communications," Dark Reading, April 22, 2021, <https://www.darkreading.com/vulnerabilities--threats/nearly-half-of-all-malware-is-concealed-in-tls-encrypted-communications-/d/d-id/1340792>, accessed November 8, 2022.

(Malware communicating and hiding, cont'd.)

IoT devices (for example, having known default passwords).⁵ Our research suggests that Mirai malware must have incidentally generated IPv4 addresses matching some addresses in the Azure, Google Cloud, and AWS public IP ranges, and the malware was exploring potential IoT communicability.

- The campaign contacting Google Cloud port 15647 is the Redline malware. This software steals information like passwords and saved credit card data from endpoint computers, such as from browsers, and other local machine information like machine memory size and similar data, and sends this data to a remote C2 server. Our data shows that the few Redline samples targeting Google Cloud port 15647 were all targeting the same IP address, which—per VT's analysis of the malware samples—was the C2 for this Redline campaign. Redline grew in popularity in 2022, and it communicated with a variety of IPs beyond Google Cloud during its campaign. There were under a hundred Redline submissions to VT for analysis at the beginning of 2022, but by the June-September period, there were roughly 7,000-8,000 submissions per month.

Suggested mitigations for Google Cloud customers

1. Use various Google Cloud Firewall capabilities to limit external access to and create appropriate restrictions within Google Cloud environments. Configure [Virtual Private Cloud \(VPC\) firewall rules](#) to restrict IP and port communications to the minimum required. Further, [Firewall Insights](#)
2. Ensure that only hardened VM images are used within Google Cloud instances. For example, check that operating systems (OSs) are kept appropriately patched, unneeded OS services are turned off, and any default OS service account credentials have been changed. We strongly recommend using Google Cloud's [Shielded VMs](#), which are VMs hardened by a set of security controls to prevent remote attacks, privilege escalation, and related security threats.
3. Use [VPC Service Controls](#) as part of a defense in depth strategy to prevent malicious external sources from accessing your cloud resources. VPC Service Controls can restrict access to cloud

helps you understand the effectiveness of your firewall rules by identifying misconfigurations and providing metrics that can alert on malicious behavior based on significant changes. In particular, Telnet (port 23) is one of the most popular ports targeted (Illustration 1). And since Telnet has a number of vulnerabilities, it's strongly recommended to deny access to Google Cloud resources via the port using firewall settings. Port 23's activities should also be monitored using Firewall Insights. Also, consider eventually using Google Cloud Threat Intelligence Objects (GCTIO)—a capability of [Cloud Firewall Standard](#), which is currently in preview, and will be available in 2023. GCTIO, among other features, can alert or block malicious external IPs and domains from communicating with your Google Cloud instance, based on the analysis of several threat intelligence feeds.

⁵ Forensicxs, "Mirai: the 'open source' Botnet," March 15, 2020, <https://www.forensicxs.com/mirai-the-open-source-botnet/>, accessed November 7, 2022.

(Malware communicating and hiding, cont'd.)

resources based on a requestor's IP address, identity, and trusted client devices, and it can log access denials in Cloud Logging for subsequent review.

4. Consider signing up for Security Command Center Premium to take advantage of [Event Threat Detection](#) to quickly detect Google Cloud threats based on logged Cloud events. Event Threat Detection monitors Cloud Logging and Google Workspace logs by analyzing administrator actions, authentication, and other key activities.
5. Consider enrolling in Security Command Center Premium to utilize the [Virtual Machine Threat Detection](#) (VM Threat Detection) capability, to detect cryptomining activities within VMs. VM Threat Detection compares a VM's memory, available to the Google Cloud hypervisor, against memory patterns created by cryptomining software, when it's executing in a VM. If signatures match, VM Threat Detection will place the corresponding findings into Security Command Center to be viewed in the service's dashboard.
6. Use appropriate authentication and authorization controls to restrict access to important Google Cloud and Google Workspace applications. Turn on MFA for critical applications and key users, such as administrators. Consider using hardware-based tokens such as security keys for the second factor of authentication. Offerings such as the Advanced Protection Program can also provide protection measures that you can also consider using for your own private accounts. For Google Cloud, use BeyondCorp Enterprise context-aware rules to restrict access to the Google Cloud console and the Google APIs based on requesting-user characteristics such as their OS or IP address. For Google Cloud-hosted web applications, use Identity-Aware Proxy (IAP). Configuring certain individuals or groups to access specific applications via Identity and Access Management (IAM) and IAP, and having the applications validate signed application headers in the HTTP requests to the applications, allows only the aforementioned identities to get access to the applications.
7. Use [VirusTotal](#) (VT) to examine if malware is contacting your Google Cloud IP addresses. VT will identify malware samples contacting specific IP addresses or ports. It also identifies if particular target IP addresses are hard coded within the malware samples themselves. Searching in VT for malware communicating with or embedding your Google Cloud IPs into its code can help lead to even more focused IP or port examinations, the shutdown of suspect assets, and the throttling of certain communications, if required.

APT10: Lessons learned from studying government-backed cloud targeting

[APT10](#), also known as [MenuPass](#), is a threat actor group sponsored by the People's Republic of China. The group has specialized in targeting cloud infrastructure and, between the discovery of the Cloud Hopper campaign in 2016 and the A41APT campaign of 2020, has evolved its techniques from basic cloud account hijacking to the targeting of VPN technologies. This elegant approach deliberately targets those organizations that have yet to adopt a full zero trust environment, preferring to use trusted VPN connections in a hybrid or transitioning cloud environment. APT10's ability to leverage both open source and custom tooling to target an organization's unique infrastructure composition make them highly adept at identifying the inevitable weak spots of hybrid enterprises.

The Cloud Hopper campaign

In 2016, a combined effort by PwC, BAE Systems, and the UK's National Cyber Security Centre (NCSC) (with support from other unnamed organizations) discovered that APT10 had compromised multiple IT service providers who were providing services to the

enterprise networks of various public and private sector organizations. The initial attack vectors varied, with the actor often using techniques that do not require significant technical know-how, such as spear-phishing, and a variety of initial-stage implants to establish an initial foothold. From there, the actor would use a series of open source tools such as Mimikatz, NBTScan, and TCPing to move laterally through the network and identify system administrator accounts that had access to the relevant "jump boxes" in their customers' cloud environments. This provided extensive access, limited only by the access of the compromised administrators, to the foundational infrastructure of any of the original IT service providers' customers. From there, the malicious actors would exfiltrate data either directly over the host's cloud infrastructure or via the IT service provider themselves.

When initially identified, reporting included analysis from both private and public sector researchers and victims, and indicated that activity could have started as early as 2014. Due to the nature of underlying

(APT10: Lessons learned, cont'd.)




Figure 1. A means to an end: APT10's Cloud Hopper campaign compromised several IT service providers to gain access to target cloud environments.

infrastructure compromise, the likelihood of persistence being established (for example, by creating cloud-specific service accounts to regain access at a later date) was also high, and subsequent discoveries continued for years to come. In 2018, the United States Department of Justice issued an [indictment](#) for two Chinese nationals believed to be behind the attacks. One of the key technological remediations recommended by the wider security community at the time was to increase the use of VPNs and similar technologies, which could add an extra layer of security to cloud environments; however, this would go on to become a further attack vector for unsuspecting targets.

Targeting of VPN clients

More recent APT10 attacks have identified targeting of VPN capabilities. APT10 used a custom malware loader, dubbed [Ecipekac](#) by researchers, which hijacked VPN sessions by exploiting known vulnerabilities in VPN software. The exploit was custom designed to run in

memory and target specific instances of the Pulse Connect Secure VPN software. The loader would be delivered through low-equity tooling, in keeping with previous techniques, though upgrading to the Cobalt Strike framework.

The targeting of VPN software is especially significant for enterprises that have not yet implemented zero trust environments and so rely on VPN setups and restrictive firewall policies to manage corporate network accesses. Such a setup is common in the transition toward cloud, as it provides a convenient interface for security engineers to manage both cloud access and wider-enterprise tooling, including SaaS, by limiting access to predetermined network ranges. This, too, provides convenience for the attacker who merely needs to compromise the VPN infrastructure, leveraging that trusted infrastructure to exploit subsequent targets—including, but not limited to, cloud infrastructure.

(APT10: Lessons learned, cont'd.)

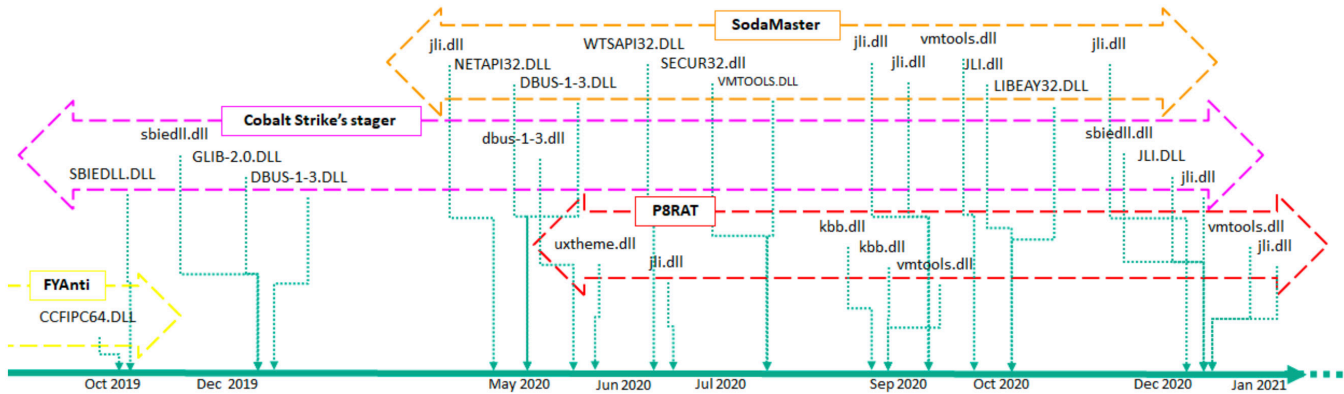


Figure 2. Timeline of the Ecipekac loader files and payloads [image: Kaspersky]

Conclusion

Among cyberthreat researchers, we often organize our research efforts according to different types of technologies. This allows analysts to focus and become experts on those technologies. The same is true of state-sponsored malicious cyber actors: APT10 may well be one such specialist team that concentrates on technologies that enable infrastructure access and control. Where such a technology is the ultimate goal of the attack, such as the Cloud Hopper campaign, we observe APT10 taking a firm lead. But their efforts are continuous and, arguably, more effective when allowed to track the technology developments of their targets and establish prepositioning for future attacks—perhaps even to enable other actor groups’ operations, handing off access from one to support another. Other state-sponsored groups are also upskilling on targeting cloud environments with increasingly refined techniques.

Threat groups probably developing methods to threaten operational technology deployed in the cloud

Historically, organizations using operational technology (OT) to support the production of goods and services have attempted to isolate production networks from external information technology (IT) services. This separation was meant to ensure the safety of their people, technology, and processes by impeding external actors from having access to cyber-physical infrastructure and critical data.

Organizations have increasingly integrated OT systems with IT infrastructure, including the implementation of cloud infrastructure, to scale production, develop efficiencies, and handle geographically distributed processes. Google Cloud's Mandiant cyber-physical threat intelligence analysts and incident responders are not aware of any high-impact cyberattacks against organizations that have implemented cloud services to support OT systems. Nevertheless, Google Cloud does assess that threat groups are poised to attempt to

carry out such attacks on customer deployments based on: our understanding of prior APT operations targeting physical production networks; the rise in cloud adoption generally leading adversaries to develop different methods to reach their targets; and the importance of reliability to operation of physical systems controlled or supervised by OT networks.

Drivers of OT cloud services

While cloud implementations remain more popular in corporate environments, organizations involved in physical production are increasingly moving in the same direction. The main use cases for cloud services supporting OT are as follows:

- **Data collection, monitoring, and analytics.** Development of infrastructure to facilitate large-scale data flows and analysis using artificial intelligence and machine learning capabilities.

(Threat groups, cont'd.)

Applications include developing data lakes to gather and process operational data, tracking large numbers of geographically dispersed assets, building disaster recovery and database backups, and measuring end-user operational data.

- **Predictive maintenance.** Third-party cloud providers process data collected from field assets to perform large-scale data analytics and deliver insights on production assets. The feedback enables users to learn about the asset's health and deliver timely maintenance. To provide an example, [Schneider Electric](#) in 2022 described how predictive maintenance for circuit breakers can help organizations to ensure safety and reliability, and to avoid the costs of downtime.
- **Remote asset control and operation.** Cloud-based solutions sometimes support remote asset control applications. This includes both cloud-based platforms to interact with industrial Internet of Things (IIoT) devices, and cloud-based supervisory control and data acquisition (SCADA) systems that can sometimes control widely dispersed systems distributed across large geographical areas, such as pipelines, rails, and energy transmission or distribution devices.
- **Collaborative workspaces.** While less common in industrial environments, some vendors offer cloud-based workspaces for product and code development. These services enable engineers to collaborate on projects and share them internally and externally.

(Threat groups, cont'd.)

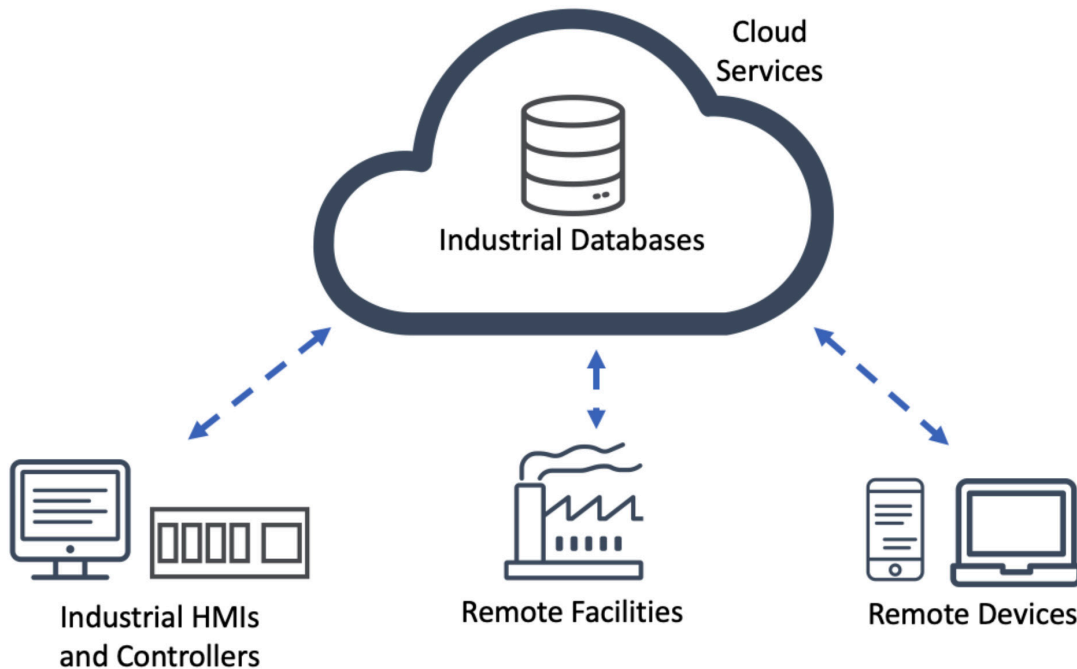


Figure 1. Overview of cloud infrastructure supporting OT.

Methods of probable OT attack against cloud deployments

As of today, there have been no publicly documented attacks impacting OT run via cloud services. This is not surprising given the small number of OT incidents observed in the wild and that cloud solutions for OT are not yet widely adopted across the industry. Yet in 2020, Mandiant [estimated](#) that 15% of their incident response investigations involved public cloud assets, demonstrating a shift in both enterprise planning and attacker operations against IT networks, which we expect to follow against OT networks.

- [Mandiant's Red Team](#) has also reproduced remote attacks against cyber-physical infrastructure. This was the case of an engagement in which the

red team accessed an endpoint meter control infrastructure for a state-wide smart grid environment from the internet and turned it off. Among other things, the team accessed a human machine interface (HMI) portal for meter control infrastructure and issued a disconnect command for a target endpoint meter in the smart grid environment.

Despite the lack of such documented attacks, cloud exposure can make previously isolated physical systems known to attackers, who could then remotely reach critical systems and data in production environments where there is no

(Threat groups, cont'd.)

proper segmentation or secure configurations. The implementation of cloud services modifies the attack surface of an organization by creating new paths for attackers to gather information or even interact with physical processes.

- Highly sophisticated actors targeted OT systems and infrastructure leveraging [intermediary systems](#), which are networked IT assets such as computers and servers. This was the case in the [TRITON](#) attack, where malicious actors traversed the victim's IT network and deployed a custom attack framework to manipulate industrial safety systems at a critical infrastructure facility and inadvertently caused a process shutdown. When incorporating cloud services to interact with or gain visibility into OT assets, these become another intermediary system that an actor can leverage to remotely reach physical production assets.
- In 2022, Mandiant disclosed [INCONTROLLER](#), a set of tools that can be leveraged by actors to target a range of industrial controllers from Schneider Electric and Omron. However, to deploy these tools, an actor requires remote access to the target's production networks or equipment. From an OT perspective, an actor may leverage cloud services to reach the target assets or to gather process information and determine additional tools they need to build.
- Mandiant has observed actors deleting or limiting access to cloud data for several reasons, such as covering their tracks or harming their victim. Due to the high availability requirements of OT systems—which often require real-time process data to support physical processes—loss of access to data stored in clouds could mean lost process visibility or even halted operations.
- » In 2019, steel producer [Norsk Hydro](#) was impacted by a ransomware attack. This resulted in suspended production resulting from loss of access to process data and inventories. While the case was not related to cloud services, it illustrates how curtailed access to such data flows can stop the production of goods and services.
- » If an asset owner is using cloud services to remotely control production systems, it is also possible that an interruption of the service may lead to loss of control over such assets. This can be prevented by establishing manual and logical redundancy mechanisms.
- During the last few years, other [actors with varying levels of skill](#) and resources have used common IT tools and techniques to gain access to and interact with OT systems that are exposed on the internet. This illustrates that organizations relying on IT or cloud services for remote interaction or visibility into OT assets and data cannot count on “security through obscurity,” and should meticulously establish segmentation and access boundaries to prevent external actors from finding and interacting with their assets.
- OT process data that is stored in infrastructure from cloud providers can also become a new target for attackers. By attempting to compromise the third-party infrastructure instead of the industrial organization itself, an attacker may seek to get access to process data from multiple customers at a single time. This is particularly relevant in the case of OT, where actors often perform reconnaissance campaigns

(Threat groups, cont'd.)

to gather information that is necessary to build tools and capabilities to target production systems.

Mitigations

As the adoption of cloud-based solutions to support OT production expands, defenders should focus on understanding how such implementation modifies their attack surface and on foreseeing potential future impacts on production. Defenders should concentrate on:

- **Design architecture.** Segment networks to establish boundaries to monitor and analyze traffic in the same way that it would happen in internal IT networks. Implement micro-segmentation by dividing the cloud services into different blocks and workload levels. Ensure that operational data—often based on less secure legacy communication protocols previously available only on-premises—that is monitored, stored, or analyzed in cloud services is encapsulated and encrypted before transitioning out of the OT demilitarized zone (DMZ). Limit remote connectivity to OT networks and assets to only what is necessary for operation.
 - **Risk assessment.** Understand associated risks and legal responsibilities of cloud implementations interfacing with OT assets. Analyze implications for incident response procedures, broader OT cybersecurity regulatory requirements, data ownership and availability, and security standards.
 - **Redundancy.** Establish redundancy mechanisms to maintain access to process data and asset control if communication flows are interrupted due to a technical failure or an attack against the cloud provider. Establish data backups in case data in the cloud is either lost or corrupted.
- **Incident response.** Adapt organizational procedures to respond to quicker change management processes and train security personnel who understand the nuances of both cloud security and OT. Some examples of relevant knowledge include:
 - » A comprehensive understanding of what cloud services are critical to maintaining local control over infrastructure in the event of an enterprise-side compromise that forces the organization to disconnect OT from enterprise and cloud networks.
 - » Familiarity with fail-safe modes that allow control and safety systems to be controlled locally if cloud SCADA systems are disrupted.
 - » Familiarity with the specific tooling required for data collection and forensics required to interact with OT assets. To support this process we suggest following [Mandiant's Digital Forensics and Incident Response Framework for Embedded OT Systems](#).

Backups increasingly targeted by threat actors

Mandiant research indicates that threat actors are increasingly targeting backups to inhibit reconstitution after an attack. In addition, targeting and, in some cases, creating backups allows threat actors to engage in reconnaissance of affected organizations, to escalate privileges, and to gather intelligence. These actions may include disabling and deleting backups, deleting virtual machines, disabling security software, and stopping processes and services that may interfere with file encryption. Examples of this type of activity Mandiant has observed include:

- Operators and users of high-profile Conti ransomware have deployed malware capable of deleting shadow copies, backups, virtual machines and snapshots, and evidence of their activity.
 - » Criminal actor group FIN12 (Conti-affiliated) used living-off-the-land techniques to manually delete volume shadow copies.
 - » Forensic evidence and leaked chat logs from the Conti group indicate that threat actors sought to identify and prioritize systems such as domain controllers, network-attached storage and file servers, virtualized environments, and backups in order to maintain persistence and maximize impact.
 - Numerous other ransomware families also contain functionality to automatically delete volume shadow copies and stop services related to backup solutions, including LockBit, Ryuk, and Babuk.
- Off-site backups have also been targeted by threat actors in multiple cases:
- In August 2019, the DDS Safe cloud-based backup system used by hundreds of American dental offices to safeguard patients' medical records was hit with Sodinokibi ransomware. The attackers reportedly exploited remote management software used by DDS Safe to back up client data, resulting in the encryption of hundreds of files that contained patient information.
 - In November 2020, NetGain Technologies, a cloud-hosting provider and MSP, was forced to take some of its data centers offline after falling victim to a ransomware attack. Because of the attack, some 200,000 patients' personally identifiable information (PII) may have been compromised.

(Backups, cont'd.)

Mitigation recommendations

- Have a cloud-specific backup strategy in place that is tested at least twice annually. This backup strategy should also include configurations and templates of stored assets, not solely backups of data or machine state.
- Create IAM permissions that segment the access and roles needed for creation, deletion, and changes to backups, thereby ensuring that account compromises do not create a direct pathway to move to the backups.
- Consider using technologies such [WORM](#) (Write Once Read Many) or the [Bucket Lock](#) feature on Google Cloud to provide immutable and policy-compliant backup storage.
- Consider implementing resilient architecture such as multi-region cloud use and backup mirroring to reduce risk of data loss or inaccessibility.
- Encrypt all backups and, as an extra measure, use customer-managed encryption keys (CMEK) and segregate key access roles, which would prevent attackers from being able to read the backups.

Use of cloud infrastructure to conduct DDoS attacks

The low barrier of entry for cloud computing in the form of trials or free tiers offered by cloud service providers (CSPs), the ability to instantly create and scale resources, and the readily available tooling has unlocked new opportunities for bad actors. The implications of these challenges to customers range from incurring additional cloud usage costs to reputational damage. In Q3 2022, Google's Trust and Safety team observed an increase in outbound layer 7 distributed denial-of-service (DDoS) abuse on Google Cloud while around the same time period in late Q2 2022 the Google Cloud Armor team blocked the [world's largest layer 7 DDoS attack](#) from external sources. Attackers have shifted from relying on compromised computers in residential environments and a single tactic to leveraging cloud resources in data centers and combining tactics to achieve their goals.

In Q3 2022, Trust and Safety systems flagged free tier or trial accounts abusing Google Cloud resources by conducting outbound DDoS attacks. During this period, we observed the attackers creating cost-optimized Compute Engine instances. Within two hours of creation, 50% of these flagged projects triggered

DDoS alerts. And within four hours, 87% triggered alerts. This suggested that attackers were creating instances for this purpose and launching attacks the same day. Upon triggering an alert, our systems will notify the project owner, giving them an opportunity to investigate and resolve the alert—after which, if no action is taken, the instance will be shut down. Furthermore, our systems are continually updated to adapt to the latest DDoS attacks observed.

- Throughout 2022, attackers attempting to evade volumetric traffic detection have been observed tunneling their DDoS traffic through open proxies hosted across the internet, including cloud providers. Attackers scan for open proxies, compile a list of available servers, and focus their traffic on their intended targets.
- During one of these attacks, Google observed malicious actors gaining access to a customer's proxy servers intended for their business customers. Due to the scaling abilities of the cloud, the customer's load balancer autoscaled and spun up additional instances of the services that were also used in a DDoS attack, further exacerbating

(Use of cloud infrastructure, cont'd.)

the situation. Our systems notified the customer of the observed DDoS behavior and our teams assisted the customer in returning to normal operations.

Mitigation recommendations

The Event Threat Detection service within Security Command Center can utilize VPC flow logs to detect [outgoing DoS](#). When considering defenses for incoming DDoS attacks, customers should review the best practice guide for [DDoS protection and mitigation](#) and leverage [Google Cloud Armor](#), which includes a web application firewall with rules to mitigate the [OWASP Top 10 risks](#) and a [24/7 DDoS response team](#).

Google Cloud