telelink business services

Advanced Security Operations Center
Telelink Business Services
www.tbs.tech

Monthly Security Bulletin

August 2023

This security bulletin is powered by Telelink Business Services' Advanced Security Operations Center

The modern cybersecurity threat landscape is constantly evolving. New vulnerabilities and zero-day attacks are discovered every day. The old vulnerabilities still exist. The tools to exploit these vulnerabilities are applying more complex techniques. But are getting easier to use.

Mitigating modern cyber threats require solutions for continuous monitoring, correlation, and behavior analysis that are expensive and require significant amount of time to be implemented. Moreover, many organizations struggle to hire and retain the expensive security experts needed to operate those solutions and provide value by defending the organizations.

The ASOC by Telelink allows organizations get visibility, control,

Why Advanced Security Operations Center (ASOC) by Telelink?

- Delivered as a service, which guarantees fast implementation, clear responsibility in the Supplier and ability to cancel the contract on a monthly basis.
- Built utilizing state of the art leading vendor's solutions.
- Can be sized to fit small, medium, and large business needs.
- No investment in infrastructure, team, trainings or required technology.
- Flexible packages and add-ons that allow pay what you need approach.
- Provided at a fraction of the cost of operating your own SOC.

LITE Plan

425 EUR/mo

- Gain visibility on the security posture of all your company's IT infrastructure
- Analysis of up to 2 GB/day log data
- Optional emergency response team (ERT) and user and endpoint behavior analytics (UEBA)

Get visibility on the cyber threats targeting your company!

PROFESSIONAL Plan

1225 EUR/mo

- Gain visibility on your company's security posture and recommendations on how to deal with security threats, risks, and actors
- Analysis of up to 5 GB/day log data and 100 GB/day network data
- Optional ERT and UEBA

Start to mitigate cyber threats and minimize the risk!

ADVANCED Plan

2 575 EUR/mo

- Gain complete visibility, deep analysis, recommendations, and security awareness trainings for your employees
- Analysis of up to 10 GB/day log data and 200 GB/day network data
- Included ERT and optional UEBA

Complete visibility, deep analysis, and cyber threat mitigation!

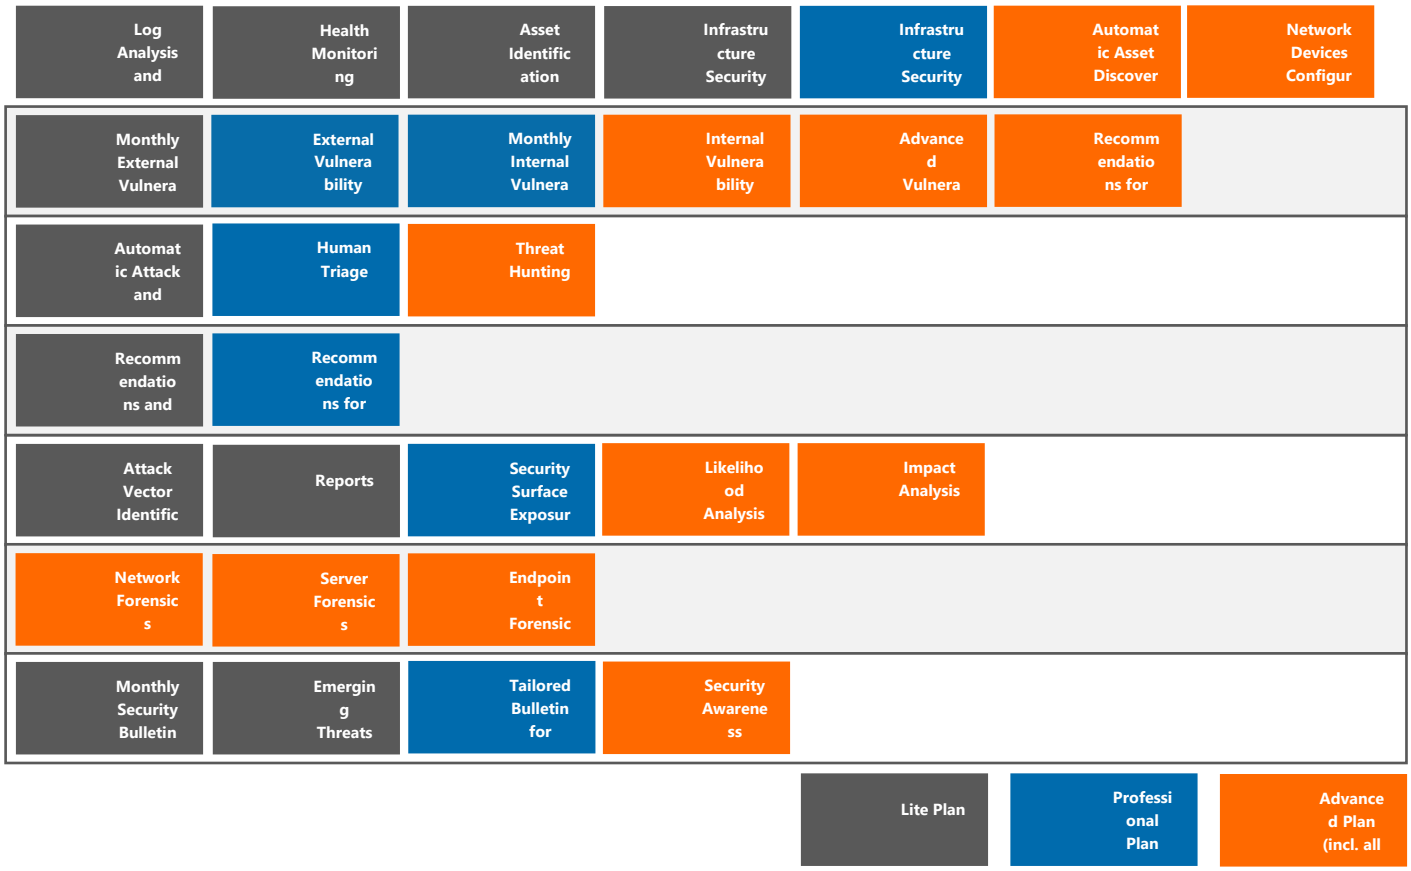

What is inside:

- Infrastructure Security Monitoring – the essential minimum to cybersecurity and to detect anomalies is to monitor your infrastructure 24x7x365
- Vulnerability Management – get visibility on the risks new or old vulnerabilities are posing to your IT infrastructure and get recommendations on how to reduce or mitigate those risks
- Attack Detection – get data from state-of-the-art cybersecurity tools, detect attacks and breaches, and involve our ASOC Analytics Team to perform human triage and threat hunting to precisely define the risks of the attack
- Reports and Recommendations – get detailed tailored reports with structured recommendations on how to prevent malicious activities (attacks) and take preventive measures
- Advanced Attack Analysis – get information on the attack vector, the attack surface, potential threat actors, and their objectives and motives for the attack
- Forensic Analysis – in case of severe cybercrimes the ASOC team can perform forensic analysis and/or support the authorities
- Bulletins, Training and Awareness – be timely informed on critical vulnerabilities with tailored and emerging threats bulletins and security awareness trainings to stop people being the weakest link

Table of Contents

1. New tool exploits Microsoft Teams bug to send malware to users

A member of U.S. Navy's red team has published a tool called TeamsPhisher that leverages an unresolved security issue in Microsoft Teams to bypass restrictions for incoming files from users outside of a targeted organization, the so-called external tenants.

The tool exploits a problem highlighted last month by Max Corbridge and Tom Ellson of UK-based security services company Jumpsec, who explained how an attacker could easily go around Microsoft Teams' file-sending restraints to deliver malware from an external account.

The feat is possible because the application has client-side protections that can be tricked into treating an external user as an internal one just by changing the ID in the POST request of a message.

Streamlining attacks on Teams

'TeamsPhisher' is a Python-based tool that provides a fully automated attack. It integrates the attack idea of Jumpsec's researchers, techniques developed by Andrea Santese, and authentication and helper functions from Bastian Kanbach's 'TeamsEnum' tool.

"Give TeamsPhisher an attachment, a message, and a list of target Teams users. It will upload the attachment to the sender's Sharepoint, and then iterate through the list of targets," reads the description from Alex Reid, the developer of the red team utility.

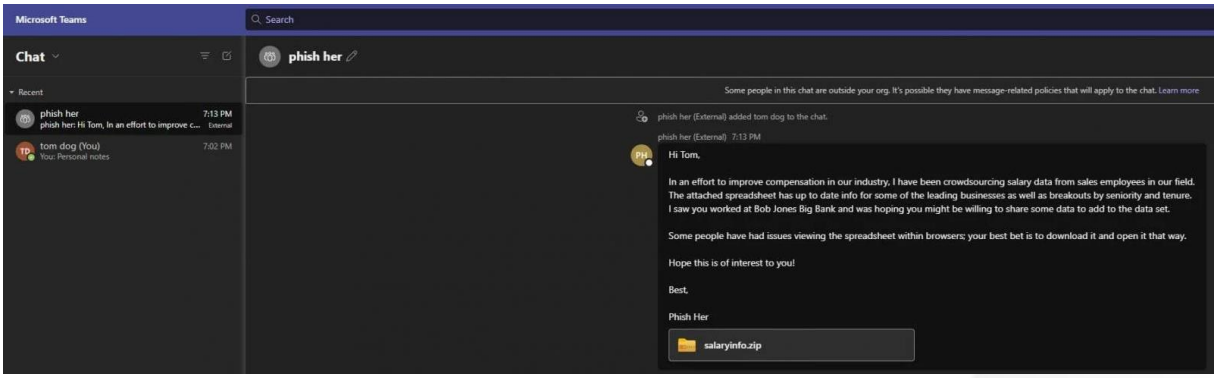

Phishing message as seen by the recipient (github.com/Octoberfest7)

TeamsPhisher first verifies the existence of the target user and their ability to receive external messages, which is a prerequisite for the attack to work.

It then creates a new thread with the target, sends them a message with a Sharepoint attachment link. The thread appears in the sender's Teams interface for (potential) manual interaction.

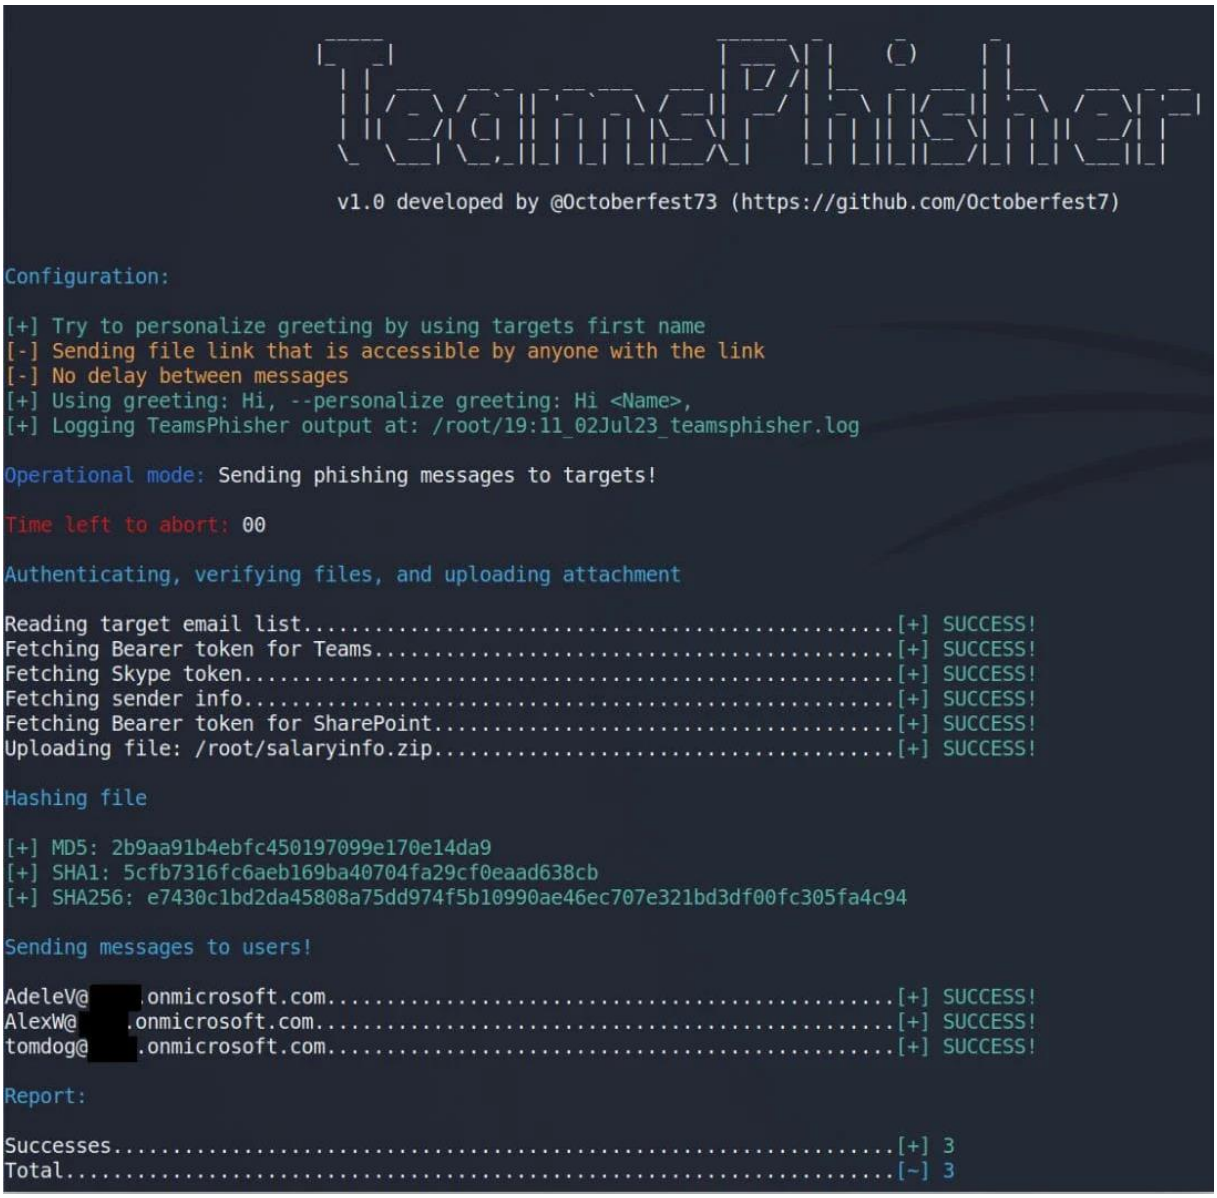

TeamsPhisher output (github.com/Octoberfest7)

TeamsPhisher requires users to have a Microsoft Business account (MFA is supported) with a valid Teams and Sharepoint license, which is common for many major companies.

The tool also offers a "preview mode" to help users verify the set target lists and to check the appearance of messages from the recipient's perspective.

Other features and optional arguments in TeamsPhisher could refine the attack. These include sending secure file links that can only be viewed by the intended recipient, specifying a delay between message transmissions to bypass rate limiting, and writing outputs to a log file.

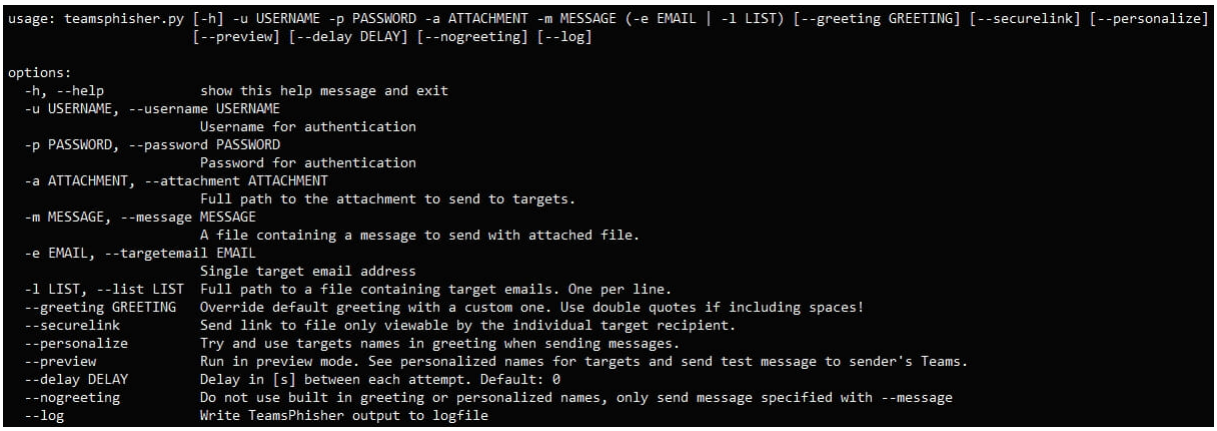

All options and arguments supported by the tool (github.com/Octoberfest7)

Unsolved problem

The issue that TeamsPhisher exploits is still present and Microsoft told Jumpsec researchers that it did not meet the bar for immediate servicing.

BleepingComputer also reached out to the company last month for a comment about plans to fix the problem but did not receive a response. We reiterated our request for comment from Microsoft but did not receive a reply at publishing time.

Although TeamPhisher was created for authorized red team operations, threat actors can also leverage it to deliver malware to target organizations without setting off alarms.

Until Microsoft decides to take action about this, organizations are strongly advised to disable communications with external tenants if not needed. They can also create an allow-list with trusted domains, which would limit the risk of exploitation.

Update 7/6 - A Microsoft spokesperson has sent BleepingComputer the following comment:

> *We're aware of this report and have determined that it relies on social engineering to be successful.*
>
> *We encourage customers to practice good computing habits online, including exercising caution when clicking on links to web pages, opening unknown files, or accepting file transfers.*

Source: *https://www.bleepingcomputer.com/news/security/new-tool-exploits-microsoft-teams-bug-to-send-malware-to-users/*

2. New StackRot Linux kernel flaw allows privilege escalation

Technical information has emerged for a serious vulnerability affecting multiple Linux kernel versions that could be triggered with "minimal capabilities." The security issue is being referred to as StackRot (CVE-2023-3269) and can be used to compromise the kernel and elevate privileges.

A patch is available for the affected stable kernels since July 1st and full details about the issue along with a complete exploit code are expected by the end of the month.

Security researcher Ruihan Li discovered and reported the vulnerability. He explains in a post today that it affects the kernel's memory management subsystem, a component in charge with implementing the virtual memory and demand paging, memory allocation for the kernel's needs and the user space programs, as well as mapping files into the processes' address space.

StackRot impacts all kernel configurations on Linux versions 6.1 through 6.4.

Although Li sent the vulnerability report on June 15th, creating a fix took almost two weeks due to its complexity, and Linus Torvalds led the effort.

"On June 28th, during the merge window for Linux kernel 5.5, the fix was merged into Linus' tree. Linus provided a comprehensive merge message to elucidate the patch series from a technical perspective. These patches were subsequently backported to stable kernels (6.1.37, 6.3.11, and 6.4.1), effectively resolving the "Stack Rot" bug on July 1st," the researcher clarified.

StackRot details

StackRot arises from the Linux kernel's handling of stack expansion within its memory management subsystem, tied to managing virtual memory areas (VMAs).

Specifically, the weak spot is in "maple tree," a new data structure system for VMAs introduced in Linux kernel 6.1 that replaced the "red-black trees" and relied on the read-copy-update (RCU) mechanism.

The vulnerability is a use-after-free (UAF) problem stemming from the way stack expansion was handled, because the maple tree could replace a node without obtaining the memory management (MM) write lock.

As the Linux kernel expands the stack and removes the gap between VMAs, a new node is created in the "maple tree," and the old one is marked for deletion after current reads finish due to the maple tree's RCU safety.

However, during the RCU grace period, a use-after-free issue may occur when a process accesses the old node, thus creating an exploitable context for elevating privileges.

```
- CPU 0 -                                          - CPU 1 -

mm_read_lock()                                     mm_read_lock()
expand_stack()                                     find_vma_prev()
  expand_downwards()                                 mas_walk()
    mas_store_prealloc()                               mas_state_walk()
      mas_wr_story_entry()                               mas_start()
        mas_wr_modify()                                    mas_root()
          mas_wr_store_node()                                node = rcu_dereference_check()
            mas_replace()                                    [ The node pointer is recorded ]
              mas_free()
                ma_free_rcu()
                  call_rcu(&mt_free_rcu)
                  [ The node is dead ]
mm_read_unlock()

[ Wait for the next RCU grace period.. ]
rcu_do_batch()                                       mas_prev()
  mt_free_rcu()                                        mas_prev_entry()
    kmem_cache_free()                                    mas_prev_nentry()
    [ The node is freed ]                                  mas_slot()
                                                             mt_slot()
                                                               rcu_dereference_check(node->..)
                                                               [ UAF occurs here ]
                                                   mm_read_unlock()
```

Race condition in a multi-CPU system that results in use-after-free flaw (github.com/lrh2000)

Exploit coming

Ruihan Li notes that exploiting StackRot is a challenging task and that CVE-2023-3269 may be the first example of a theoretically exploitable use-after-free-by-RCU (UAFBR) vulnerability.

However, the researcher announced plans to disclose the complete technical details about StackRot and a proof-of-concept (PoC) exploit by the end of July.

Linux kernel 6.1 has been approved as the long-term support (LTS) version since February. However, not all major Linux distributions have adopted it.

For instance, Ubuntu 22.04.2 LTS (Jammy Jellyfish), whose standard support ends in April 2027, ships with Linux kernel version 5.19. On the other hand, Debian 12 (Bookworm) comes with Linux kernel 6.1.

A complete list of Linux distributions using kernel version 6.1 or higher is available from DistroWatch.

Users should check the kernel version their Linux distro runs on and choose one that is not affected by StackRot or an updated release that contains the fix.

Source: *https://www.bleepingcomputer.com/news/security/new-stackrot-linux-kernel-flaw-allows-privilege-escalation/*

3. New ‘Big Head’ ransomware displays fake Windows update alert

Security researchers have dissected a recently emerged ransomware strain named ‘Big Head’ that may be spreading through malvertising that promotes fake Windows updates and Microsoft Word installers.

Two samples of the malware have been analyzed before by cybersecurity company Fortinet, who looked at the infection vector and how the malware executes.

Today, Trend Micro published a technical report on Big Head that claiming that both variants and a third they sampled originate from a single operator who is likely experimenting with different approaches to optimize their attacks.

Faking a Windows update

‘Big Head’ ransomware is a .NET binary that installs three AES-encrypted files on the target system: one is used to propagate the malware, another is for Telegram bot communication, and the third encrypts files and can also show the user a fake Windows update.

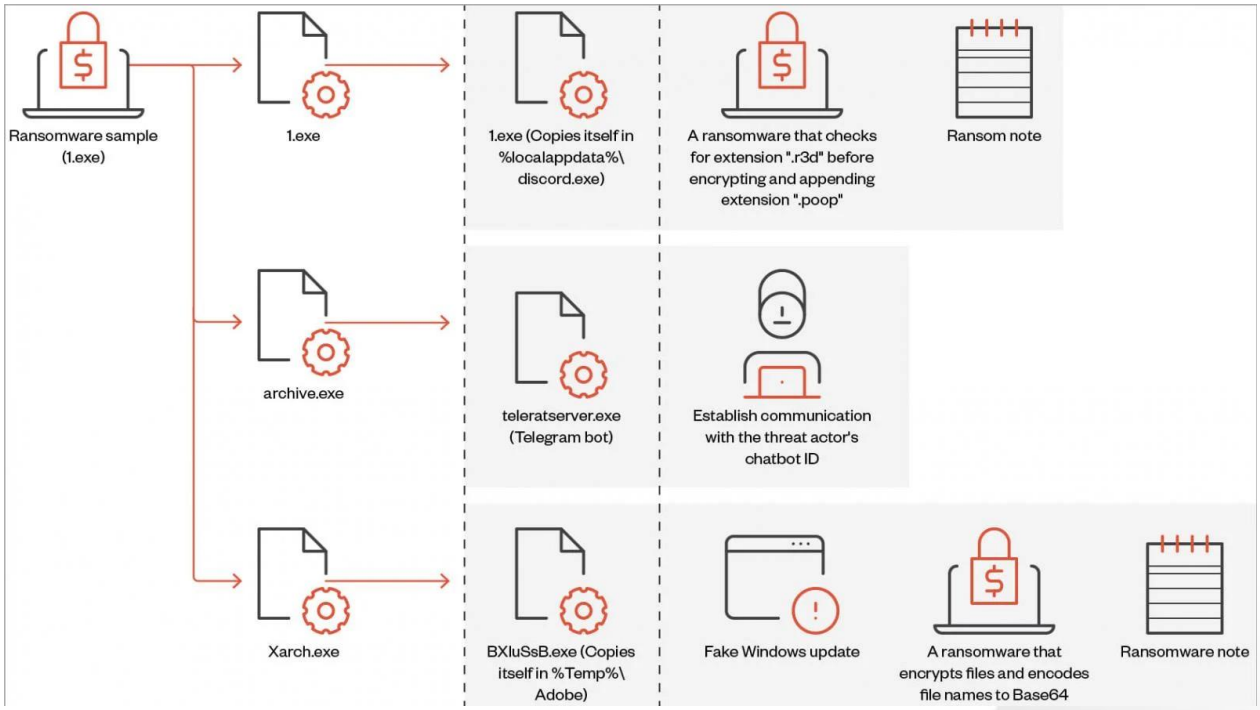

Big Head's infection routine (Trend Micro)

On execution, the ransomware also performs actions such as creating a registry autorun key, overwriting existing files if needed, setting system file attributes, and disabling the Task Manager.

```
public void Autorun()
{
    string text = Path.GetTempPath() + "Adobe//";
    try
    {
        if (!Directory.Exists(this.pathbackup))
        {
            DirectoryInfo directoryInfo = Directory.CreateDirectory(text);
            directoryInfo.Attributes = (FileAttributes.Hidden | FileAttributes.Directory);
        }
    }
    catch
    {
    }
    string location = Assembly.GetExecutingAssembly().Location;
    string fileName = Path.GetFileName(location);
    try
    {
        File.Copy(location, Path.Combine(text, fileName), false);
    }
    catch
    {
    }
    RegistryKey registryKey = Registry.CurrentUser.CreateSubKey("Software\\Microsoft\\Windows\\CurrentVersion\\RunOnce\\");
    string str = Path.GetTempPath() + "Adobe";
    registryKey.SetValue(fileName, str + "\\" + fileName);
    registryKey.Close();
}
```

Creating the Registry Autorun (Trend Micro)

Each victim is assigned a unique ID that's either retrieved from the %appdata%\ID directory or it is generated using a random 40-character string.

The ransomware deletes shadow copies to prevent easy system restoration before encrypting the targeted files and appending a ".poop" extension to their filenames.

```
".mdf", ".db", ".mdb", ".sql", ".pdb", ".pdb", ".pdb", ".dsk", ".fp3", ".fdb",
".accdb", ".dbf", ".crd", ".db3", ".dbk", ".nsf", ".gdb", ".abs", ".sdb", ".sdb",
".sdb", ".sqlitedb", ".edb", ".sdf", ".sqlite", ".dbs", ".cdb", ".cdb", ".cdb", ".bib",
".dbc", ".usr", ".dbt", ".rsd", ".myd", ".pdm", ".ndf", ".ask", ".udb", ".ns2", ".kdb",
".ddl", ".sqlite3", ".odb", ".ib", ".db2", ".rdb", ".wdb", ".tcx", ".emd", ".sbf",
".accdr", ".dta", ".rpd", ".btr", ".vdb", ".daf", ".dbv", ".fcd", ".accde", ".mrg",
".nv2", ".pan", ".dnc", ".dxl", ".tdt", ".accdc", ".eco", ".fmp", ".vpd", ".his",
".fid"
```

File types targeted by Big Head (Trend Micro)

Also, Big Head will terminate the following processes to prevent tampering with the encryption process and to free up data that the malware should lock.

```
"taskmgr", "sqlagent", "winword", "sqlbrowser", "sqlservr", "sqlwriter", "oracle",
"ocssd", "dbsnmp", "synctime", "mydesktopqos", "agntsvc.exeisqlplussvc", "xfssvccon",
"mydesktopservice", "ocautoupds", "agntsvc.exeagntsvc", "agntsvc.exeencsvc",
"firefoxconfig", "tbirdconfig", "ocomm", "mysqld", "sql", "mysqld-nt", "mysqld-opt",
"dbeng50", "sqbcoreservice"
```

Processes terminated before encryption (Trend Micro)

The Windows, Recycle Bin, Program Files, Temp, Program Data, Microsoft, and App Data directories are skipped from encryption to avoid rendering the system unusable.

Trend Micro has found that the ransomware checks if it runs on a virtual box, looks for the system language, and only proceeds to the encryption if it's not set on that of a country member of the Commonwealth of Independent States (former Soviet states).

```
"ar-SA", "ar-AE", "nl-BE", "nl-NL", "en-GB", "en-US", "en-CA", "en-AU", "en-NZ", "fr-
BE", "fr-CH", "fr-FR", "fr-CA", "fr-LU", "de-AT", "de-DE", "de-CH", "it-CH", "it-IT",
"ko-KR", "pt-PT", "es-ES", "sv-FI", "sv-SE", "bg-BG", "ca-ES", "cs-CZ", "da-DK", "el-
GR", "en-IE", "et-EE", "eu-ES", "fi-FI", "hu-HU", "ja-JP", "lt-LT", "nn-NO", "pl-PL",
"ro-RO", "se-FI", "se-NO", "se-SE", "sk-SK", "sl-SI", "sv-FI", "sv-SE", "tr-TR"
```

System languages valid for encryption (Trend Micro)

During the encryption, the ransomware displays a screen that purports to be a legitimate Windows update.

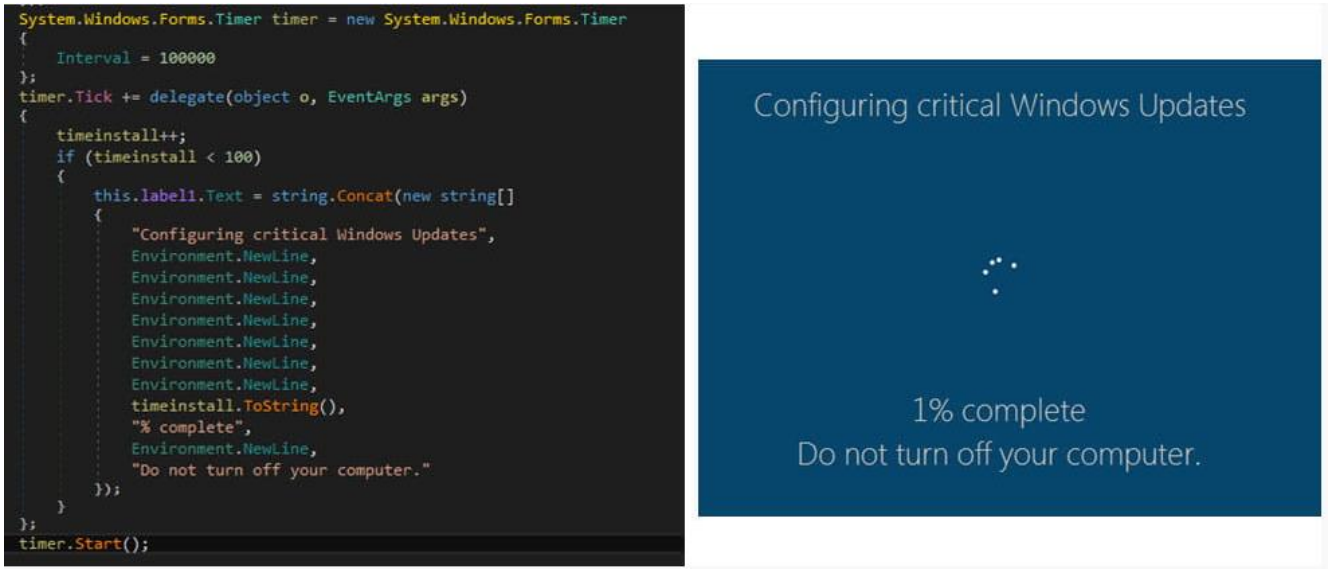

Fake Windows update masking the file encryption (Trend Micro)

After the encryption process completes, the following ransom is dropped on multiple directories, and the victim's wallpaper is also changed to alert of the infection.

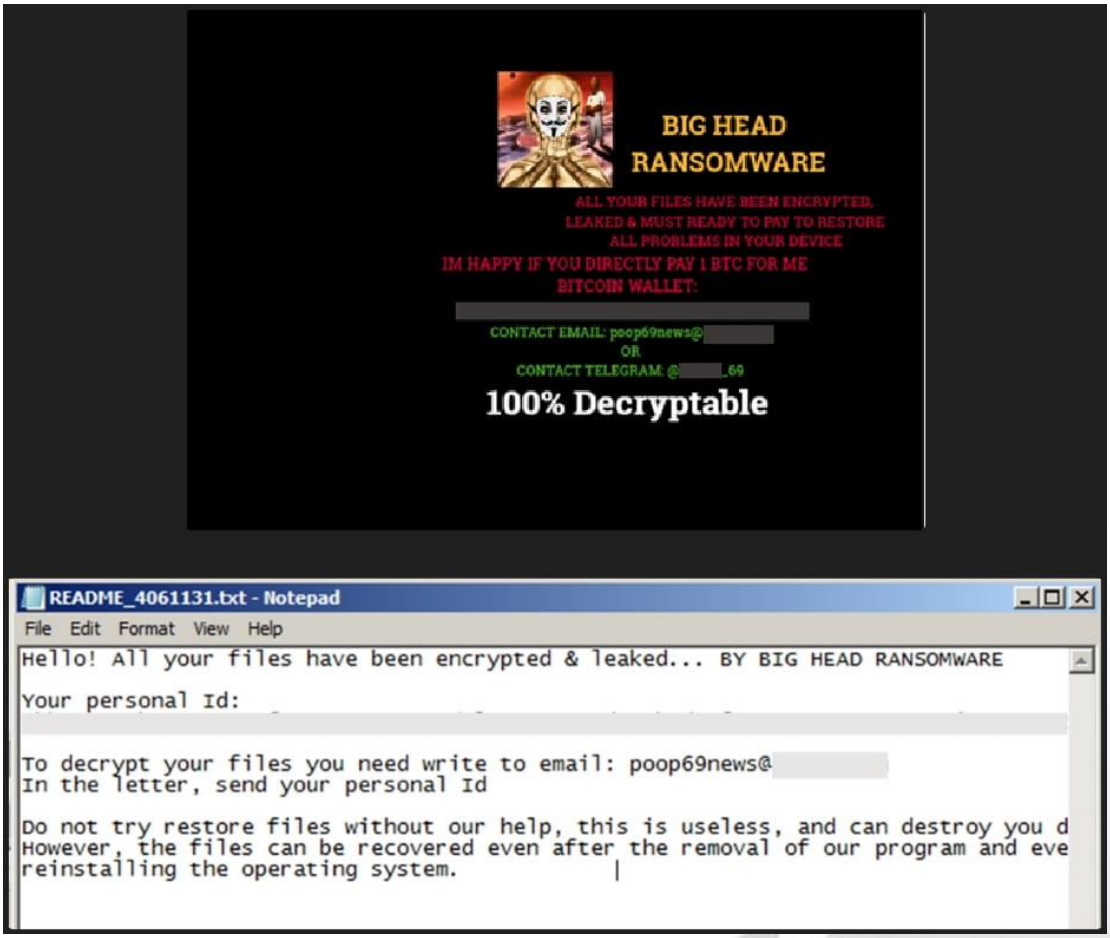

Wallpaper and ransom note (Trend Micro)

Other variants

Trend Micro also analyzed two more Big Head variants, highlighting some key differences compared to the standard version of the ransomware.

The second variant maintains ransomware capabilities but also incorporates stealer behavior with functions to collect and exfiltrate sensitive data from the victim system.

The data that this version of Big Head can steal include browsing history, list of directories, installed drivers, running processes, product key, and active networks, and it can also capture screenshots.

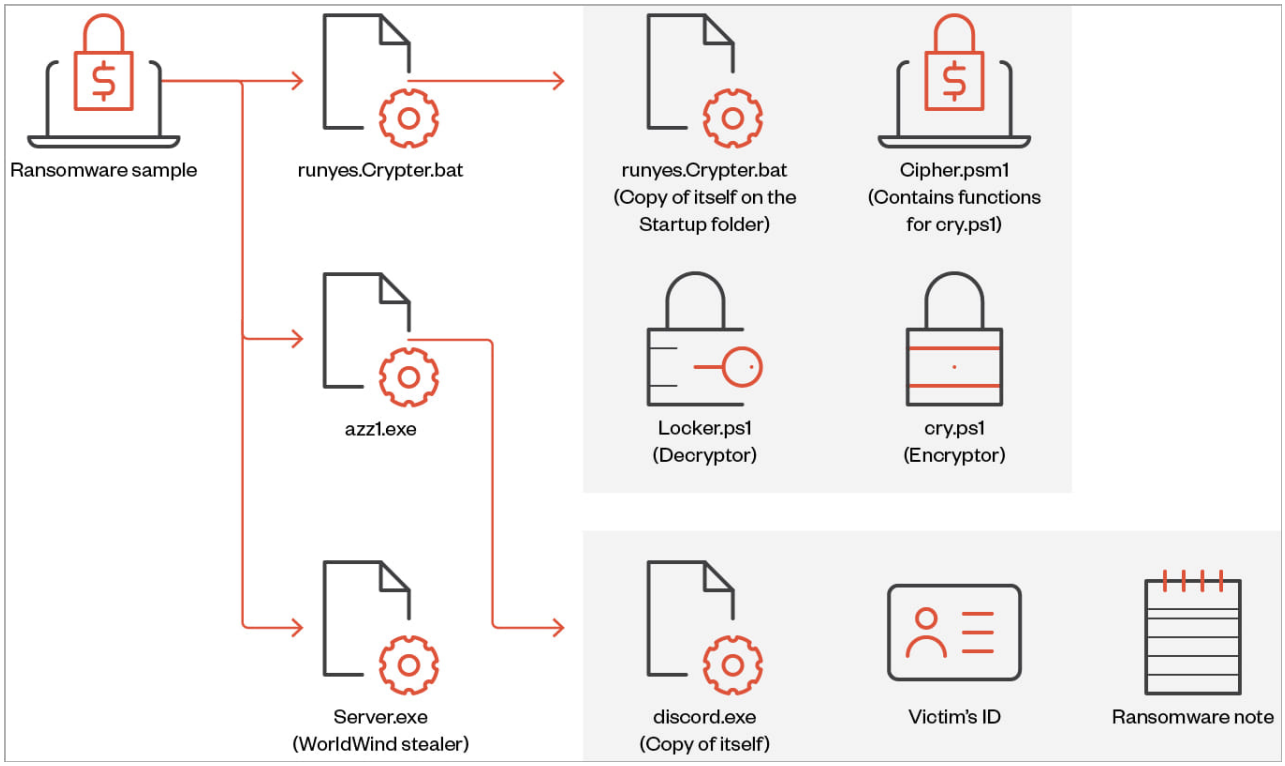

Second variant infection routine (Trend Micro)

The third variant, discovered by Trend Micro, features a file infector identified as "Neshta," which inserts malicious code into executables on the breached system.

Although the exact purpose of this is unclear, Trend Micro's analysts speculate that it could be to evade detection that relies on signature-based mechanisms.

Notably, this variant uses a different ransom note and wallpaper from the other two, yet it is still tied to the same threat actor.

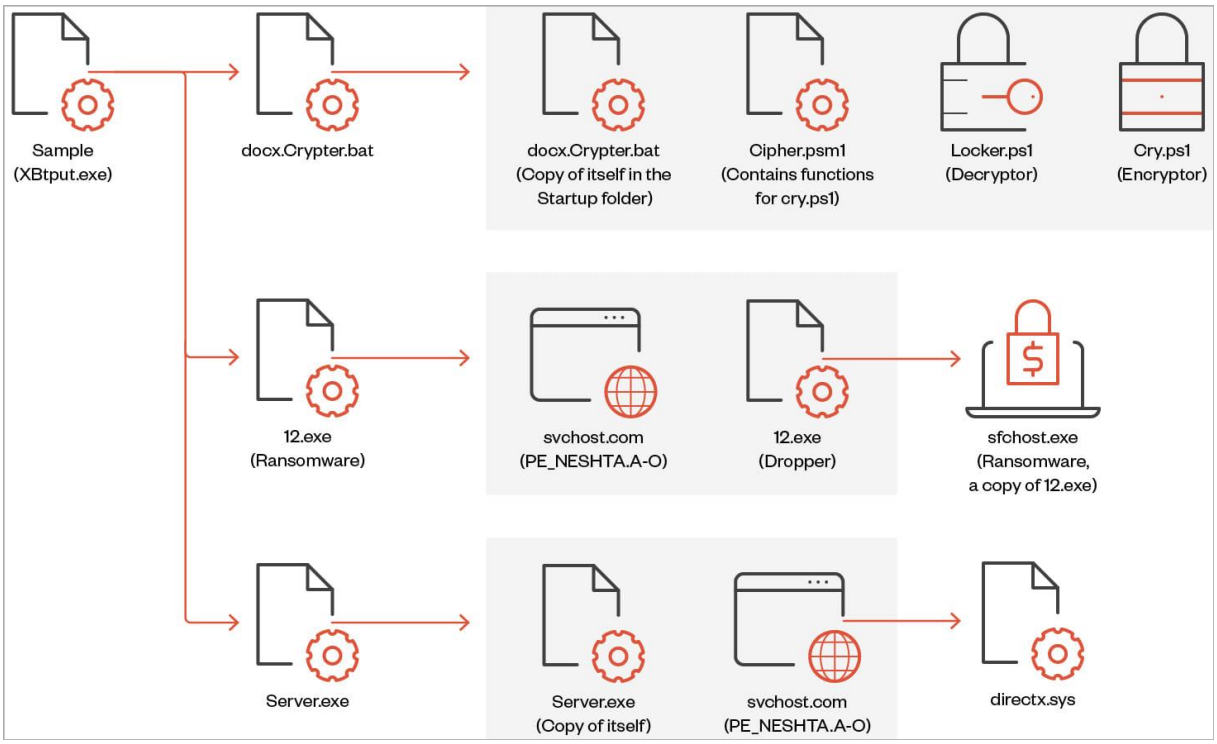

Third variant infection routine (Trend Micro)

Conclusion

Trend Micro comments that Big Head is not a sophisticated ransomware strain, its encryption methods are pretty standard, and its evasion techniques are easy to detect.

Nevertheless, it appears to focus on consumers who can be fooled with easy tricks (e.g. fake Windows update) or they have difficulty understanding the safeguards necessary to steer away from cybersecurity risks.

The multiple variants in circulation suggest that the creators of Big Head are continuously developing and refining the malware, experimenting with various approaches to see what works best.

Update 7/10/23 - Cyber-intelligence firm KELA shared additional information with BleepingComputer, indicating that Big Head's main author is likely of Indonesian origin.

KELA's analysts have discovered a user on Telegram using the same names and avatars as those found in Big Head's ransom note, claiming to be a "ransomware expert" on posts published on "IndoGhostsec."

Big Head's author on Telegram (KELA)

The user switched the group's name from 'BIG HEAD HACKER!' to 'BLACKHAT HACKER INDONESIA' in June 2022, while in March 2023, he started seeking the help of other members in his effort to create a ransomware builder and other relevant tools.

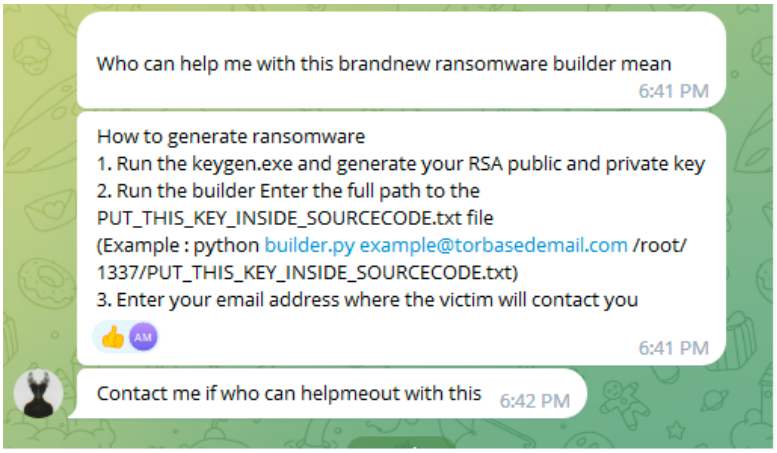

Threat actor's post on Telegram (KELA)

Source: https://www.bleepingcomputer.com/news/security/new-big-head-ransomware-displays-fake-windows-update-alert/

4. Charming Kitten hackers use new 'NokNok' malware for macOS

Security researchers observed a new campaign they attribute to the Charming Kitten APT group where hackers used new NokNok malware that targets macOS systems.

The campaign started in May and relies on a different infection chain than previously observed, with LNK files deploying the payloads instead of the typical malicious Word documents seen in past attacks from the group.

Charming Kitten is also known as APT42 or Phosphorus and has launched at least 30 operations in 14 countries since 2015, according to according to Mandiant.

Google has linked the threat actor to the Iranian state, more specifically, the Islamic Revolutionary Guard Corps (IRGC).

In September 2022, the U.S. government managed to identify and charge members of the threat group.

Proofpoint reports that the threat actor has now abandoned the macro-based infection methods involving laced Word documents and instead deploys LNK files to load their payloads.

Regarding the phishing lures and social engineering methods seen in the campaign, the hackers posed as nuclear experts from the U.S. and approached targets with an offer to review drafts on foreign policy topics.

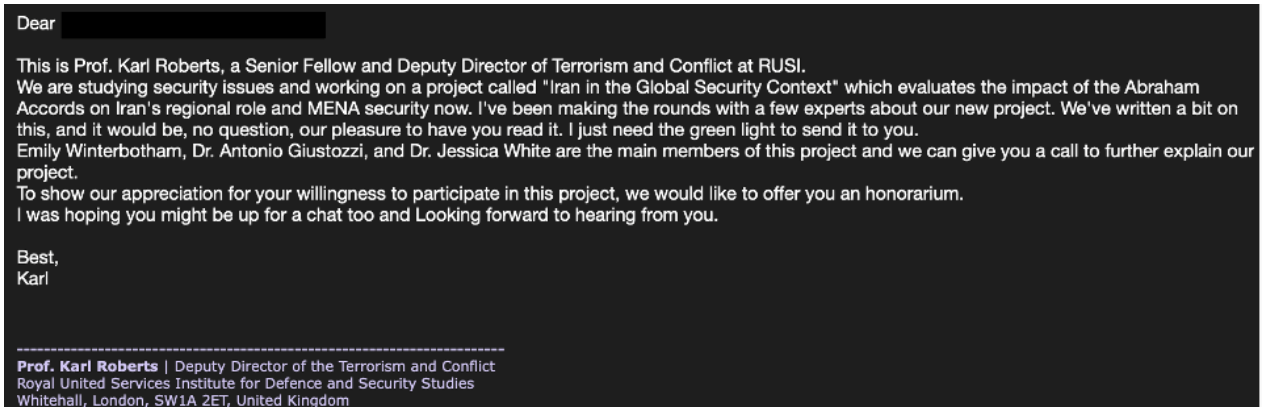
Dear [redacted]

This is Prof. Karl Roberts, a Senior Fellow and Deputy Director of Terrorism and Conflict at RUSI.
We are studying security issues and working on a project called "Iran in the Global Security Context" which evaluates the impact of the Abraham Accords on Iran's regional role and MENA security now. I've been making the rounds with a few experts about our new project. We've written a bit on this, and it would be, no question, our pleasure to have you read it. I just need the green light to send it to you.
Emily Winterbotham, Dr. Antonio Giustozzi, and Dr. Jessica White are the main members of this project and we can give you a call to further explain our project.
To show our appreciation for your willingness to participate in this project, we would like to offer you an honorarium.
I was hoping you might be up for a chat too and Looking forward to hearing from you.

Best,
Karl

Prof. Karl Roberts | Deputy Director of the Terrorism and Conflict
Royal United Services Institute for Defence and Security Studies
Whitehall, London, SW1A 2ET, United Kingdom

Email sampled from the latest Charming Kitten campaign (Proofpoint)

In many cases, the attackers insert other personas in the conversation to add a sense of legitimacy and establish a rapport with the target.


Follow-Up: Iran in the Global Security Context Project

Close Fullscreen

From: Emily Winterbotham

6/6/2023 at 4:34 AM

Hi [redacted] - This is Emily Winterbotham, a director of the Terrorism and Conflict at the RUSI Center. Yesterday I was talking about the project with Prof. Karl, who informed me that he sent you the completed parts of the project and how to access the shared folder, but unfortunately, he has not received any response from you yet.
We are very interested to benefit from your opinions and expertise in our project.
Looking forward to hearing from you.

Best
Emily
--
Emily Winterbotham | Director of the Terrorism and Conflict
Royal United Services Institute for Defence and Security Studies
Whitehall, London, SW1A 2ET, United Kingdom
T:+44 (0)207 747 2600 | E: Emswinterbotham@gmail.com

FOLLOW: www.facebook.com/RUSI.org | https://twitter.com/RUSI_org
Royal United Services Institute accepts no liability for the content of this email, or for the consequences of any actions taken on the basis of the information provided, unless that information is subsequently confirmed in writing. Any views or opinions presented in this email are solely those of the author and do not necessarily represent those of the company.

Second email from another fake persona (Proofpoint)

Charming Kitten's impersonation or fake persona assumption in phishing attacks has been documented, and so has its use of 'sock puppets' to create realistic conversation threads.

Attacks on Windows

After gaining the target's trust, Charming Kitten sends a malicious link that contains a Google Script macro, redirecting the victim to a Dropbox URL.

This external source hosts a password-protected RAR archive with a malware dropper that leverages PowerShell code and an LNK file to stage the malware from a cloud hosting provider.

The final payload is GorjolEcho, a simple backdoor that accepts and executes commands from its remote operators.

To avoid raising suspicion, GorjolEcho will open a PDF with a topic relevant to the discussion the attackers had with the target previously.

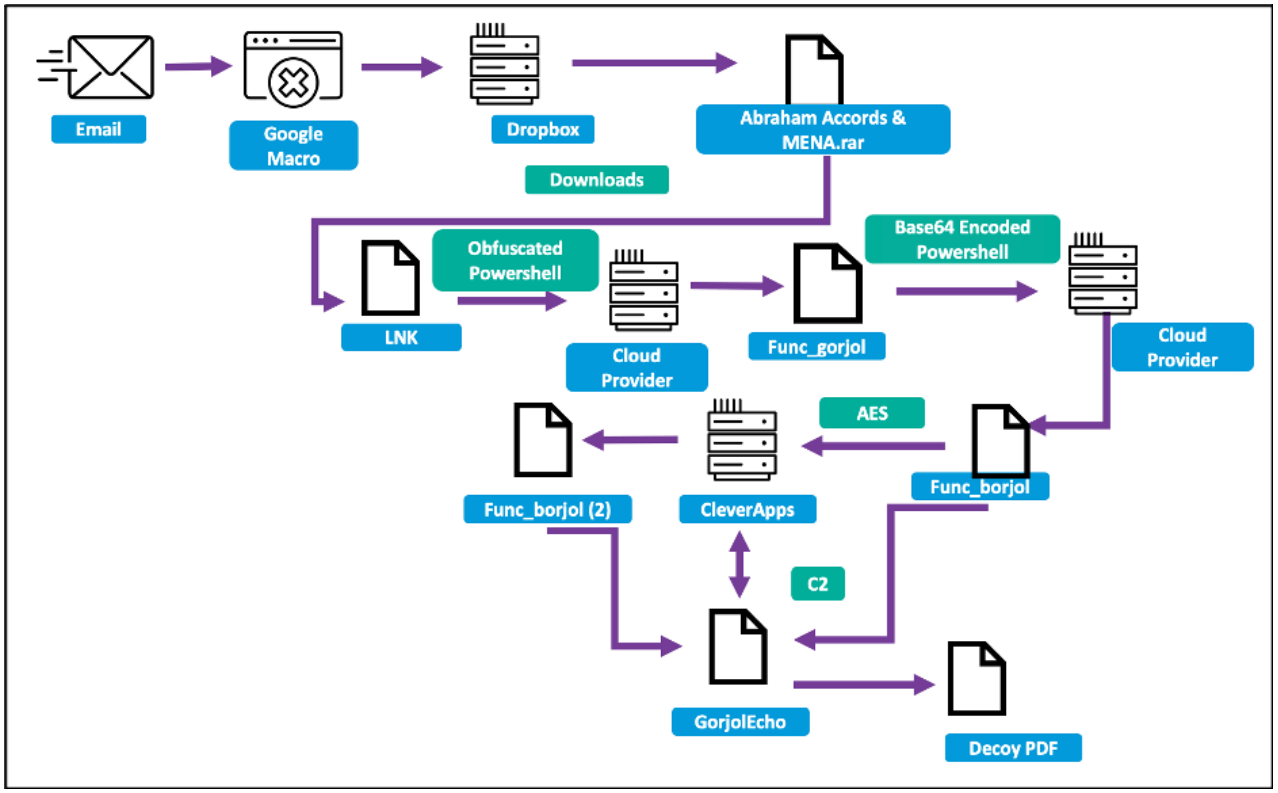

GorjolEcho infection chain (Proofpoint)

Attacks on macOS

If the victim uses macOS, which the hackers typically realize after they fail to infect them with the Windows payload, they send a new link to "library-store[.]camdvr[.]org" that hosts a ZIP file masquerading as a RUSI (Royal United Services Institute) VPN app.

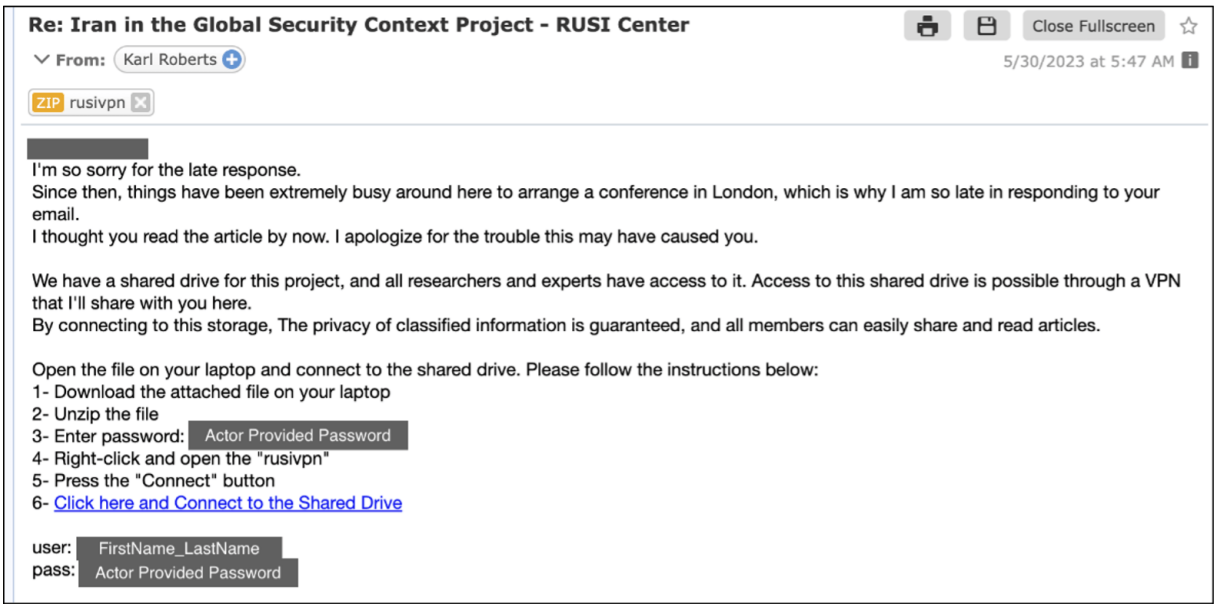

Follow-up email sent to macOS users (Proofpoint)

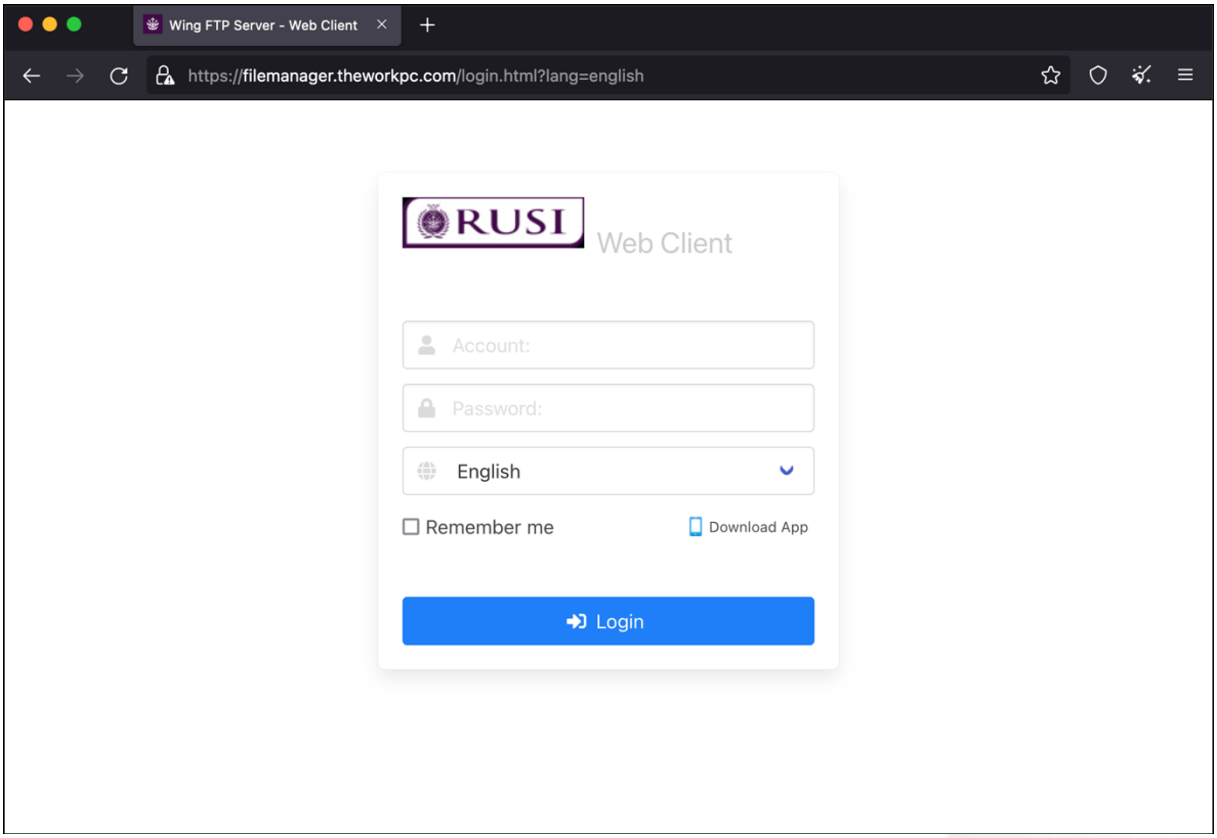

Fake RUSI VPN site dropping the NokNok malware (Proofpoint)

When executing the Apple script file in the archive, a curl command fetches the NokNok payload and establishes a backdoor onto the victim's system.

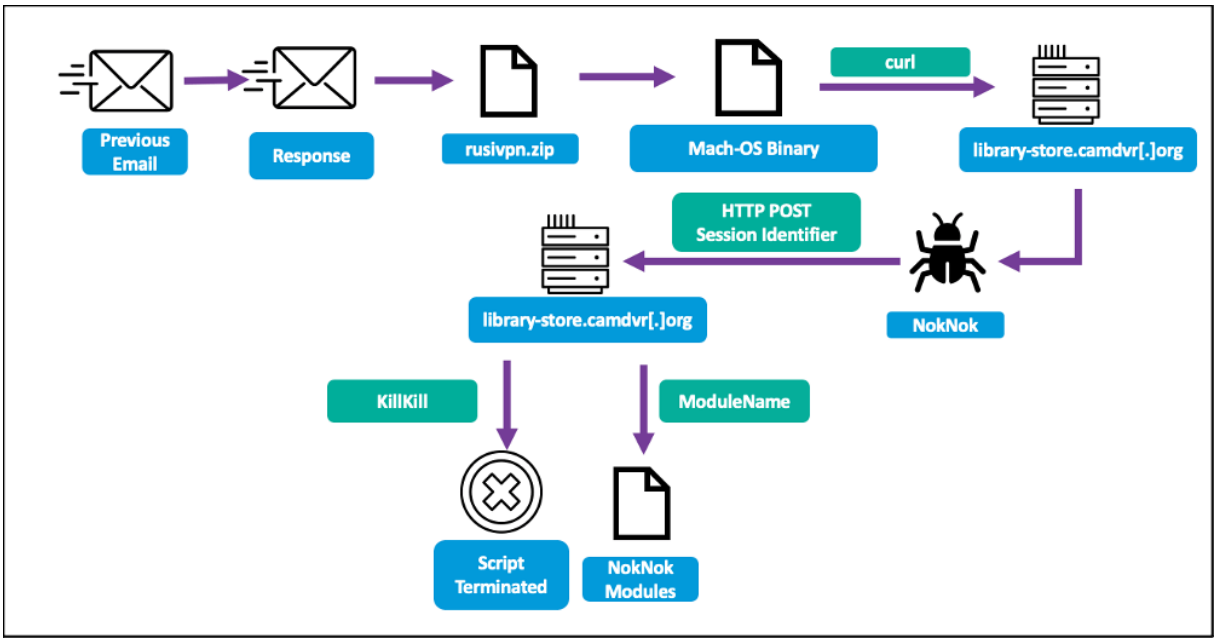

NokNok infection chain (Proofpoint)

NokNok generates a system identifier and then uses four bash script modules to set persistence, establish communication with the command and control (C2) server, and then starts exfiltrating data to it.

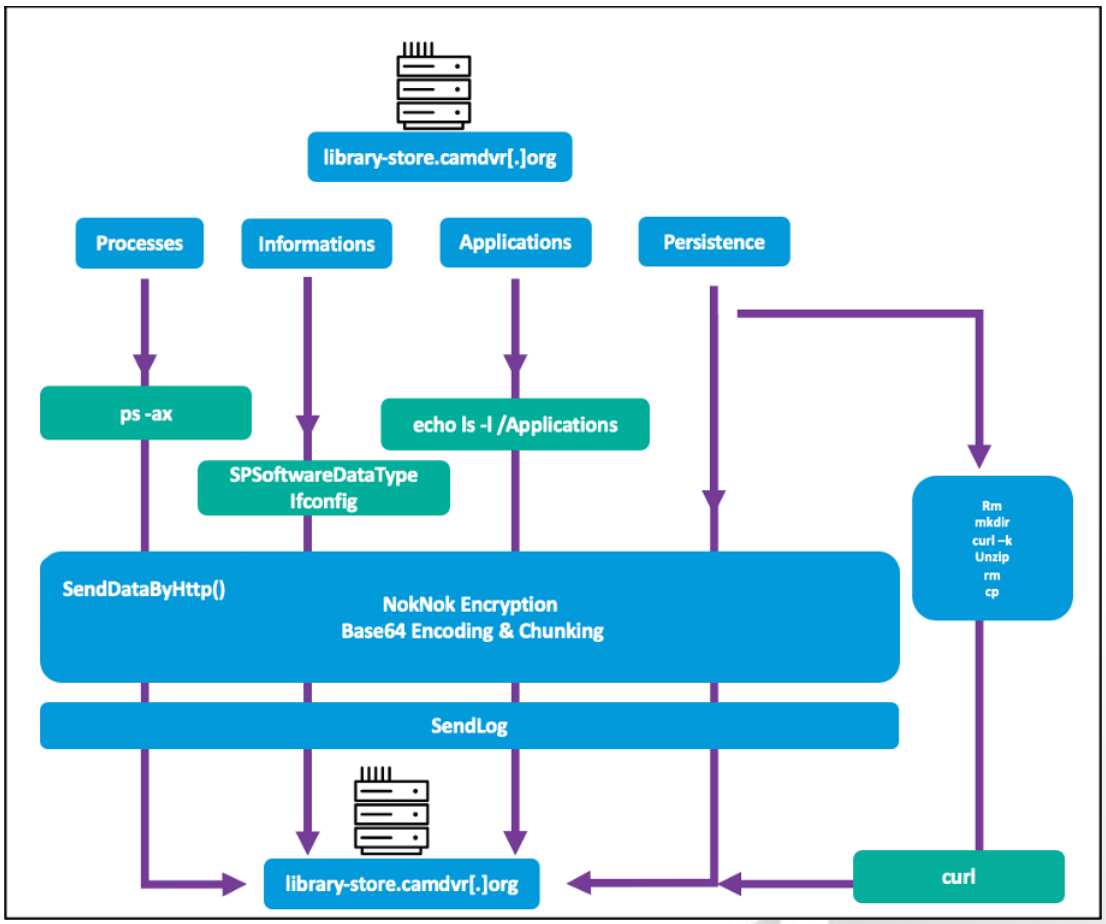

NokNok modules (Proofpoint)

The NokNok malware gathers system information that includes the version of the OS, running processes, and installed applications.

NokNok encrypts all collected data, encodes it in the base64 format, and exfiltrates it.

Proofpoint also mentions that NokNok might feature more specific espionage-related functionality through other unseen modules.

The suspicion is based on code similarities to GhostEcho, previously analyzed by Check Point.

That backdoor featured modules that allowed taking screenshots, command execution, and cleaning the infection trail. It is likely that NokNok has these functions too.

Overall, this campaign shows that Charming Kitten has a high degree of adaptability, is capable of targeting macOS systems when necessary, and highlights the growing threat of sophisticated malware campaigns to macOS users.

Source: https://www.bleepingcomputer.com/news/security/charming-kitten-hackers-use-new-noknok-malware-for-macos/

5. VMware warns of exploit available for critical vRealize RCE bug

VMware warned customers today that exploit code is now available for a critical vulnerability in the VMware Aria Operations for Logs analysis tool, which helps admins manage terabytes worth of app and infrastructure logs in large-scale environments.

The flaw (CVE-2023-20864) is a deserialization weakness patched in April, and it allows unauthenticated attackers to gain remote execution on unpatched appliances.

Successful exploitation enables threat actors to run arbitrary code as root following low-complexity attacks that don't require user interaction.

"VMware has confirmed that exploit code for CVE-2023-20864 has been published," the company noted in an update to the initial security advisory.

"CVE-2023-20864 is a critical issue and should be patched immediately as per the instructions in the advisory."

In April, VMware also issued security updates to address a less severe command injection vulnerability (CVE-2023-20865) that would let remote attackers with administrative privileges execute arbitrary commands as root on vulnerable appliances.

Both flaws have been fixed with the release of VMware Aria Operations for Logs 8.12. Fortunately, there is currently no evidence to suggest exploitation in attacks.

VMware Aria Operations flaws under attack

Recently, VMware issued another alert about a now-patched critical bug (CVE-2023-20887) in VMware Aria Operations for Networks (formerly vRealize Network Insight), allowing remote command execution as the root user and being actively exploited in attacks.

CISA also added the flaw to its list of known exploited vulnerabilities and ordered U.S. federal agencies to apply security updates by July 13th.

In light of this, admins are strongly advised to promptly apply CVE-2023-20864 patches as a precaution against potentially incoming attacks.

Although the number of online-exposed VMware vRealize instances is relatively low, it aligns with the intended design of these appliances, which primarily focus on internal network access within organizations.

Nonetheless, it's important to note that attackers often take advantage of vulnerabilities present in devices within compromised networks.

Therefore, even properly configured VMware appliances that remain vulnerable can become tempting targets within the internal infrastructure of targeted organizations.

Source: *https://www.bleepingcomputer.com/news/security/vmware-warns-of-exploit-available-for-critical-vrealize-rce-bug/*

6. Microsoft: Chinese hackers breached US govt Exchange email accounts

A Chinese hacking group has breached the email accounts of more than two dozen organizations worldwide, including U.S. and Western European government agencies, according to Microsoft.

The attacks have been pinned on a threat group tracked as Storm-0558, believed to be a cyber-espionage outfit focused on collecting sensitive information by breaching email systems.

Microsoft started investigating these attacks on June 16, 2023, following customer reports regarding unusual Office 365 mail activity.

The company discovered that starting from May 15, 2023, Storm-0558 threat actors managed to access Outlook accounts belonging to roughly 25 organizations (reportedly including the U.S. State and Commerce Departments) and some consumer accounts likely connected to them.

However, Microsoft did not share what organizations, government agencies, or countries were affected by these email breaches.

To do that, the attackers used authentication tokens forged with the help of a stolen Microsoft account (MSA) consumer signing key.

"Microsoft investigations determined that Storm-0558 gained access to customer email accounts using Outlook Web Access in Exchange Online (OWA) and Outlook.com by forging authentication tokens to access user email," Microsoft said in a blog post published late Tuesday evening.

"The actor used an acquired MSA key to forge tokens to access OWA and Outlook.com. MSA (consumer) keys and Azure AD (enterprise) keys are issued and managed from separate systems and should only be valid for their respective systems. The actor exploited a token validation issue to impersonate Azure AD users and gain access to enterprise mail."

Microsoft added that it found no evidence indicating any additional unauthorized access after it "completed mitigation of this attack."

Discovered and reported by the U.S. government

The incident was reported to Microsoft by U.S. government officials last month after the discovery of unauthorized access to Microsoft cloud-based email services.

This was confirmed by National Security Council spokesperson Adam Hodge in a statement shared with CNN.

"Last month, US government safeguards identified an intrusion in Microsoft's cloud security, which affected unclassified systems," Hodge told CNN.

"Officials immediately contacted Microsoft to find the source and vulnerability in their cloud service. We continue to hold the procurement providers of the US Government to a high security threshold."

On Tuesday, Microsoft also revealed that the RomCom Russian-based cybercriminal group exploited an unpatched Office zero-day in recent spear-phishing attacks targeting organizations attending the NATO Summit in Vilnius, Lithuania.

Source: *https://www.bleepingcomputer.com/news/security/microsoft-chinese-hackers-breached-us-govt-exchange-email-accounts/*

7. New PyLoose Linux malware mines crypto directly from memory

A new fileless malware named PyLoose has been targeting cloud workloads to hijack their computational resources for Monero cryptocurrency mining.

PyLoose is a relatively simple Python script with a precompiled, base64-encoded XMRig miner, a widely abused open-source tool that uses CPU power to solve complex algorithms required for cryptomining.

According to researchers at Wiz, PyLoose's direct execution from memory makes it incredibly stealthy and challenging to detect by security tools.

Fileless malware leaves no physical footprint on the system's drives, so it's less vulnerable to signature-based detection and typically utilizes legitimate system tools (living off the land) to inject malicious code into legitimate processes.

Wiz's security researchers first detected PyLoose attacks in the wild on June 22nd, 2023, and have since confirmed at least 200 cases of compromise by the novel malware.

"As far as we know, this is the first publicly documented Python-based fileless attack targeting cloud workloads in the wild, and our evidence shows close to 200 instances where this attack was used for cryptomining," explains the new Wiz report.

PyLoose attack chain

Wiz observed attacks that began by gaining initial access to devices through publicly accessible Jupyter Notebook services, which failed to restrict system commands.

The attacker uses an HTTPS GET request to fetch the fileless payload (PyLoose) from a Pastebin-like site, "paste.c-net.org," and load it straight into Python's runtime memory.

The PyLoose script is decoded and decompressed, loading a precompiled XMRig miner directly into the instance's memory using the "memfd" Linux utility, a known fileless malware technique in Linux.

```
1  import ctypes, os, base64, zlib
2  l = ctypes.CDLL(None)
3  s = l.syscall
4  c = base64.b64decode(b'eNrsvXlcVOX3OH4HGBZFZ3CLzI
5  e = zlib.decompress(c)
6  f = s(319, '', 1)
7  os.write(f, e)
8  p = '/proc/self/fd/%d' % f
9  os.execle(p, 'smd', {})
```

The PyLoose script (Wiz)

"The memory file descriptor, memfd, is a Linux feature that allows the creation of anonymous memory-backed file objects that can be used for various purposes, such as inter-process communication or temporary storage," explains Wiz in the report.

"Once the payload is placed within a memory section created via memfd, attackers can invoke one of the exec syscalls on that memory content, treating it as if it were a regular file on disk, and thereby launch a new process."

This enables attackers to perform payload execution straight from memory, evading most traditional security solutions.

The XMRig miner loaded into the compromised cloud instance's memory is a fairly recent version (v6.19.3) that uses the 'MoneroOcean' mining pool to mine for Monero.

Unknown threat actors

Wiz could not attribute the PyLoose attacks to any particular threat actor, as the attacker left no useful evidence behind.

The researchers comment that the adversary behind PyLoose appears highly sophisticated and stands out from the typical threat actors engaging in cloud workload attacks.

Cloud instance administrators are recommended to avoid the public exposure of services susceptible to code execution, use strong passwords and multi-factor authentication to protect access to those services, and place system command execution restrictions.

Source: *https://www.bleepingcomputer.com/news/security/new-pyloose-linux-malware-mines-crypto-directly-from-memory/*

8. Cyberattacks through Browser Extensions – the Importance of MFA

There are many avenues of attack that a threat actor can take. One that has been increasing in recent years is user-centric applications. Instead of focusing on a more highly-protected administrator account, attackers target applications or extensions that a user can easily install without IT involvement.

One example is the proliferation of extensions, such as those in Chromium-based browsers. Though not the first of its kind, a recent example is the Rilide malware strain. Trustwave identified this malware, which disguised itself as a Google Drive extension.

After installation, the extension enabled threat actors to monitor browser history, take screenshots, and inject malicious scripts that targeted cryptocurrency exchanges.

Also, cybersecurity giant, Kaspersky, recently identified 34 malicious Chrome extensions with over 87 million downloads. Multiple malicious extensions target user installations, leading to a real danger of data exfiltration and system compromise.

The Danger of Unchecked User Control

A significant shift occurred years ago when users transitioned from primarily running as administrators to running with least privileges. This change reduced the attack surface that malicious applications or attacks could exploit. In the event of an attack, the damage may be limited to that user's profile and the data they could access.

Although a compromised user account can lead to a compromised administrative account, separating the two provides significantly enhanced security. But, because users felt relatively safe with this separation and may feel annoyed about asking IT to install software, there has been a proliferation of user applications and extensions.

Examples include extensions in Chromium-based browsers or development tools like Visual Studio Code. Users may need to scrutinize the installation process more closely since these extensions are downloaded from traditionally trusted sources, such as Google and Microsoft's Visual Studio Code extension repository.

Due to this lack of attention, more and more attacks are occurring via extensions or user-profile installations of tools. Other examples include supply-side attacks from PyPi packages or malicious NPM package installations.

Previously legitimate extensions or packages may be sold to an unscrupulous group that allows a user to be silently compromised from a formerly trusted source.

Preventing Damaging User-Profile Extensions and Packages

What can an IT department and a user do to protect themselves? One approach is to vet extensions and packages and use allow-lists to proactively limit what a user can install. This way, both users and administrators can feel safer, and it ensures that only safe packages are used.

- Chrome: Allow or block apps and extensions
- Edge: Use group policies to manage extensions
- Chromium: Extension Settings
- Visual Studio Code: Not currently implemented, but a GitHub issue exists to track this

IT administrators should monitor extensions and packages that are allowed for changes in ownership and files that may signal danger, especially if done by a third party. Since

an extension may attempt to read data that a user profile can see, including files, it is especially important for users who have stored a password in a file to exercise caution.

This underscores the importance of Multi-Factor Authentication (MFA) in preventing further breaches, as a password alone would not be enough to access a sensitive system.

Of course, in the event of an attack, it is crucial to quickly clean and reset a user's account. The user-profile attack can leverage data contained within, which means that crafted phishing emails or emails sent from a legitimate user's account could be used to further propagate an attack. [...]

Source: https://www.bleepingcomputer.com/news/security/cyberattacks-through-browser-extensions-the-importance-of-mfa/

9. AVrecon malware infects 70,000 Linux routers to build botnet

Since at least May 2021, stealthy Linux malware called AVrecon was used to infect over 70,000 Linux-based small office/home office (SOHO) routers and add them to a botnet designed to steal bandwidth and provide a hidden residential proxy service.

This allows its operators to hide a wide spectrum of malicious activities, from digital advertising fraud to password spraying.

According to Lumen's Black Lotus Labs threat research team, while the AVrecon remote access trojan (RAT) compromised over 70,000 devices, only 40,000 were added to the botnet after gaining persistence.

The malware has largely managed to evade detection since it was first spotted in May 2021 when it was targeting Netgear routers. Since then, it went undetected for over two years, slowly ensnaring new bots and growing into one of the largest SOHO router-targeting botnets discovered in recent years.

"We suspect the threat actor focused on the type of SOHO devices users would be less likely to patch against common vulnerabilities and exposures (CVEs)," Black Lotus Labs said.

"Instead of using this botnet for a quick payout, the operators maintained a more temperate approach and were able to operate undetected for more than two years. Due to the surreptitious nature of the malware, owners of infected machines rarely notice any service disruption or loss of bandwidth."

Once infected, the malware sends the compromised router's info to an embedded command-and-control (C2) server. After contact making contact, the hacked machine is

instructed to establish communication with an independent group of servers, known as second-stage C2 servers.

The security researchers found 15 such second-stage control servers, which have been operational since at least October 2021, based on x.509 certificate information.

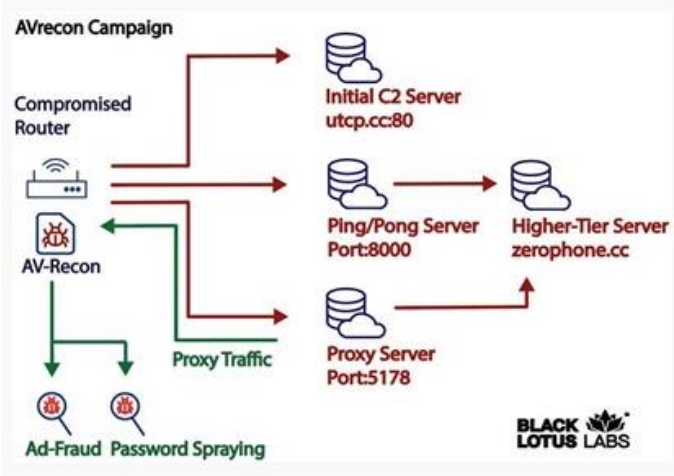

AVrecon attacks (Black Lotus Labs)

Lumen's Black Lotus security team also addressed the AVrecon threat by null-routing the botnet's command-and-control (C2) server across their backbone network.

This effectively severed the connection between the malicious botnet and its central control server, significantly impeding its capacity to execute harmful activities.

"The use of encryption prevents us from commenting on the results of successful password spraying attempts; however, we have null-routed the command and control (C2) nodes and impeded traffic through the proxy servers, which rendered the botnet inert across the Lumen backbone," Black Lotus Labs said.

In a recently issued binding operational directive (BOD) published last month, CISA ordered U.S. federal agencies to secure Internet-exposed networking equipment (including SOHO routers) within 14 days of discovery to block potential breach attempts.

Successful compromise of such devices would enable the threat actors to add the hacked routers to their attack infrastructure and provide them with a launchpad for lateral movement into their internal networks, as CISA warned.

The severity of this threat stems from the fact that SOHO routers typically reside beyond the confines of the conventional security perimeter, greatly diminishing defenders' ability to detect malicious activities.

The Volt Typhoon Chinese cyberespionage group used a similar tactic to build a covert proxy network out of hacked ASUS, Cisco, D-Link, Netgear, FatPipe, and Zyxel SOHO network equipment to hide their malicious activity within legitimate network traffic, according to a joint advisory published by Five Eyes cybersecurity agencies (including the FBI, NSA, and CISA) in May.

The covert proxy network was used by the Chinese state hackers to target critical infrastructure organizations across the United States since at least mid-2021.

"Threat actors are using AVrecon to proxy traffic and to engage in malicious activity like password spraying. This is different from the direct network targeting we saw with our other router-based malware discoveries," said Michelle Lee, threat intelligence director of Lumen Black Lotus Labs.

"Defenders should be aware that such malicious activity can originate from what appears to be a residential IP address in a country other than the actual origin, and traffic from compromised IP addresses will bypass firewall rules such as geofencing and ASN-based blocking."

Source: https://www.bleepingcomputer.com/news/security/avrecon-malware-infects-70-000-linux-routers-to-build-botnet/

10. Spotify reportedly makes users' private playlists public

In what is shaping up to be a widespread privacy controversy, Spotify has come under scrutiny following allegations by users that the music streaming service made their private playlists public without their consent.

This situation is reminiscent of a similar issue flagged back in March, raising concerns over a possible pattern of an ongoing privacy issue.

The controversy began when users reported this unexpected change to Twitter and Spotify's community forums.

"Apparently @SpotifyUSA silently made all of my private playlists public without my consent. The same happened to my wife too," tweeted Microsoft Edge Project Manager William Devereux.

"That's an absolutely unacceptable privacy violation. Anyone else noticed this happen recently? I haven't changed any privacy settings."

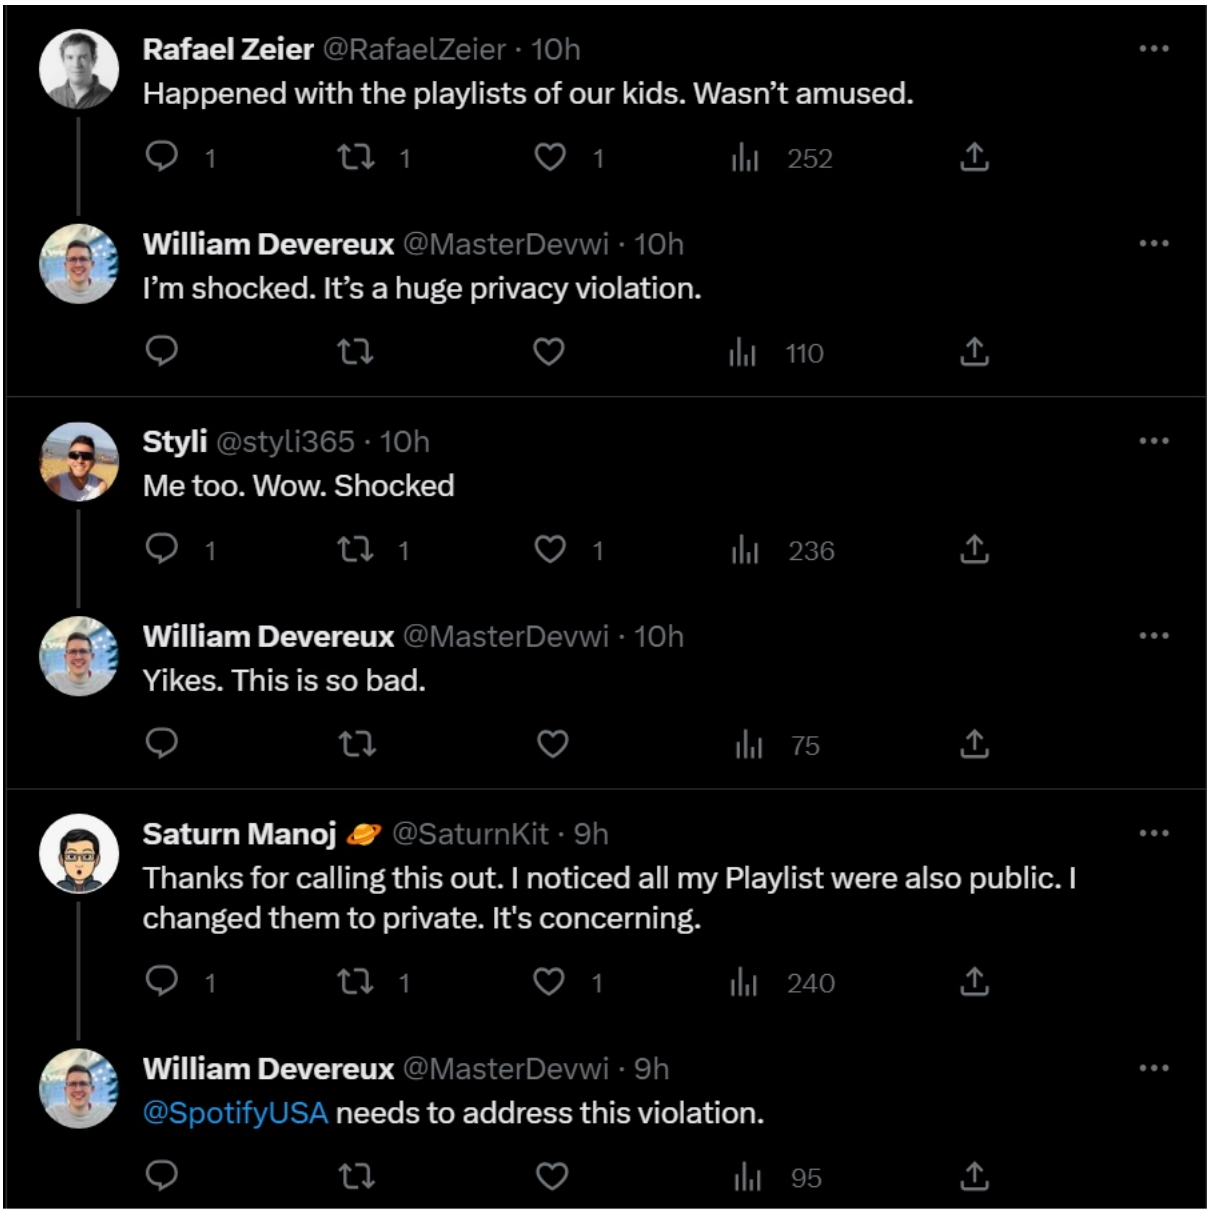

Other Spotify users find their private playlists public

There are similar reports on Spotify's forum in March, with one of the affected users being a music curator who uses Spotify professionally.

"I have revisited some lists made a month or so ago and they are all public now. Looking at more and they are now public as well!," wrote the user on Spotify's forums.

"Why has this happened? is there a way to make bulk lists private? I don't want to spend days of my life changing them one by one, there are over 1400 lists and I cant invoice for that time so it will take away from may wages."

Back in March, a user proposed a theory stating, "The actual settings of our playlists haven't changed. What was formerly known as 'private' and 'public' playlists are now all called 'public', since they weren't actually private previously, as they could be shared through a link."

The theory further suggested a new level of truly private playlists that could not be accessed by others even with a link and only playlists marked as 'on profile' could be found via search or in the 'Discovered on' section on artist pages.

Despite the theory, Spotify users insist their recent experiences indicate a different issue. They affirm that their playlists were initially marked as private upon creation and were inexplicably made public without their knowledge or permission.

In response to the reports in March, a Spotify moderator stated, "Spotify doesn't make such bulk changes and will not mess around with the settings of your collection/personal account unless you have requested this explicitly...".

However, this has done little to alleviate users' concerns, and it remains uncertain if the two issues are linked or entirely separate incidents.

We have reached out to Spotify about these reports but did not receive a reply at the time of this publishing

Source: *https://www.bleepingcomputer.com/news/technology/spotify-reportedly-makes-users-private-playlists-public/*

11. Thousands of images on Docker Hub leak auth secrets, private keys

Researchers at the RWTH Aachen University in Germany published a study revealing that tens of thousands of container images hosted on Docker Hub contain confidential secrets, exposing software, online platforms, and users to a massive attack surface.

Docker Hub is a cloud-based repository for the Docker community to store, share, and distribute Docker images. These container-creation templates include all of the necessary software code, runtime, libraries, environment variables, and configuration files to easily deploy an application in Docker.

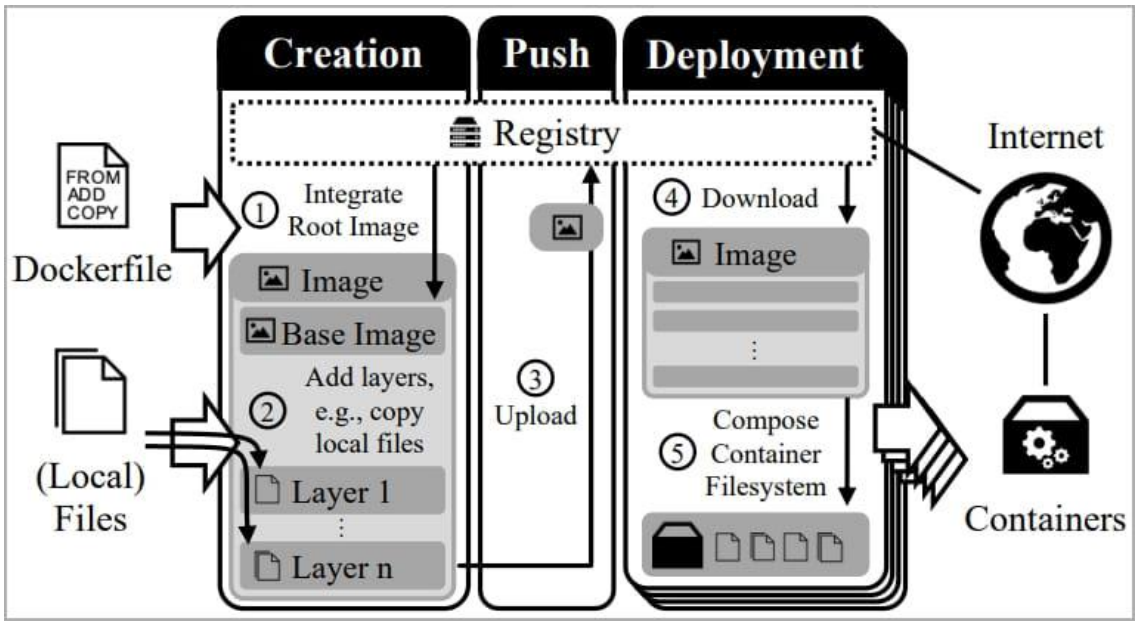

Docker image creation diagram (arxiv.org)

The German researchers analyzed 337,171 images from Docker Hub and thousands of private registries and found that roughly 8.5% contain sensitive data such as private keys and API secrets.

The paper further shows that many of the exposed keys are actively used, undermining the security of elements that depend on them, like hundreds of certificates.

(Inadvertently) exposing secrets

The study assembled a massive dataset of 1,647,300 layers from 337,171 Docker images, sourcing the latest image versions from each repository when possible.

Data analysis using regular expressions to search for specific secrets revealed the exposure of 52,107 valid private keys and 3,158 distinct API secrets in 28,621 Docker images.

The above figures were validated by the researchers excluding test keys, example API secrets, and invalid matches.

		Regular Expressions *(Section 5.1.1 / Appendix C)*	**(Distinct) Matches** *(Sec. 5.1.2)*		**Valid Secrets** *(Section 5.1.3)*		
	Domain	**Potential Threat / *(Service) Type***	**Images**	**Variables**	**Images**	**Variables**	**Total**
	Private Key	Perform man-in-the-middle attacks, fake identity, …	1,377,336	2	52,107	0	52,107
		PEM Private Key, PEM Private Key Block, PEM PKCS7, XML Private Key	(62,282)	(1)			
API	Cloud	Manage services, create new API keys, reconfigure DNS, access emails / SMS, control voice calls, read / alter private repositories, … *Alibaba*[76], *Amazon AWS*[76], *Azure*[76], *DigitalOcean*[76], *Github*[76], *Gitlab (v1, v2)*[76], *Google Cloud*[76], *Google Services*[58], *Heroku*[76], *IBM Cloud Identity Service*[76], *Login Radius*[76], *MailChimp*[58], *MailGun*[58], *Microsoft Teams*[76], *Netlify*[76], *Twilio*[58]	6,208,995 (74,460)	416 (84)	2,880	67	2,920
API	Financial	List / perform payments, inspect / alter invoices, … *Amazon MWS*[58], *Bitfinex*[76], *Coinbase*[76], *Currency Cloud*[76], *Paydirt*[76], *Paymo*[76], *Paymongo*[76], *Paypal Braintree*[58], *Picatic*[58], *Stripe*[58], *Square*[58], *Ticketmaster*[76], *WePay*[76]	42,901 (543)	4 (2)	23	2	25
API	Social Media	Tweet, access direct messages, retrieve relationships, … *Facebook*[76], [58], *Twitter*[58]	6,365,854 (439,822)	14 (8)	209	4	213
API	IoT	Retrieve (privacy-sensitive) IoT data, e.g., track cars, … *Accuweather*[76], *Adafruit IO*[76], *OpenUV*[76], *Tomtom*[76]	297 (117)	0 (0)	0	0	0

Final secrets findings (arxiv.org)

Most of the exposed secrets, 95% for private keys and 90% for API secrets, resided in single-user images, indicating that they were likely unintentionally leaked.

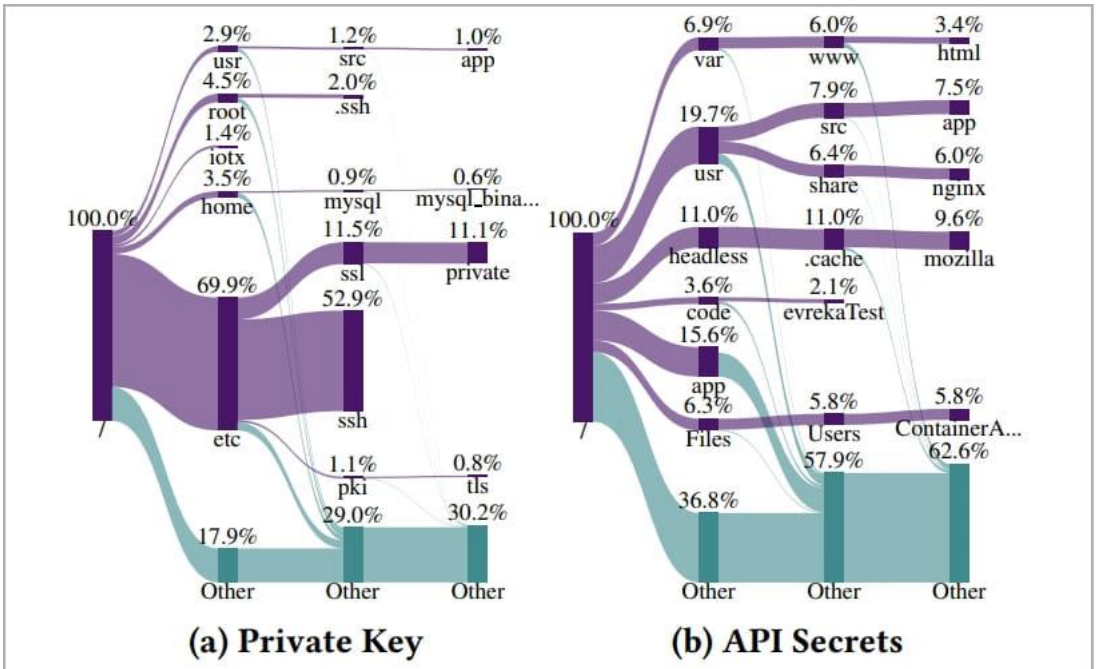

Figure 5: Most frequent file paths where we found secrets. While the major location of found private keys focuses on a few file paths, i.e., most private keys are stored where SSH host keys reside, API secrets are spread further.

The highest impact was on Docker Hub, which had a percentage of secret exposure of 9.0%, while images sourced from private registries exposed secrets at a rate of 6.3%.

This difference may indicate that Docker Hub users typically have a poorer understanding of container security than those setting up private repositories.

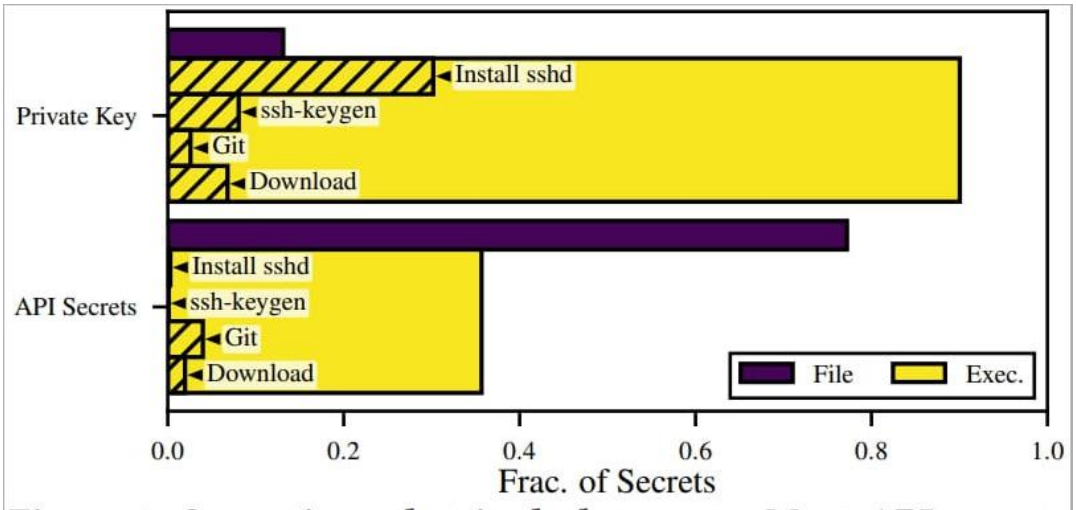

Figure 4: Operations that include secrets. Most API secrets are inserted by file and private keys by execution commands.

Use of exposed keys

Next, the researchers needed to determine the actual use of the exposed secrets to appreciate the attack surface size.

Alarmingly, 22,082 compromised certificates relying on the exposed private keys were found, including 7,546 private CA-signed and 1,060 public CA-signed certificates.

The thousand CA-signed certificates are of particular concern, as these certificates are typically used by a large number of users and are universally accepted.

At the time of the study, 141 CA-signed certificates were still valid, somewhat lessening the risk.

To further determine the use of the exposed secrets in the wild, the researchers used 15-month worth of internet-wide measurements provided by the Censys database and found 275,269 hosts that rely on the compromised keys.

These include:

- **8,674 MQTT** and **19 AMQP hosts** that potentially transfer privacy-sensitive Internet of Things (IoT) data.
- **6,672 FTP**, **426 PostgreSQL**, **3 Elasticsearch**, and **3 MySQL instances** that serve potentially confidential data.
- **216 SIP hosts** used for telephony.
- **8,165 SMTP**, **1,516 POP3**, and **1,798 IMAP servers** used for email.
- **240 SSH servers** and **24 Kubernetes instances** that use leaked keys which can lead to remote-shell access, extension of botnets, or further data access.

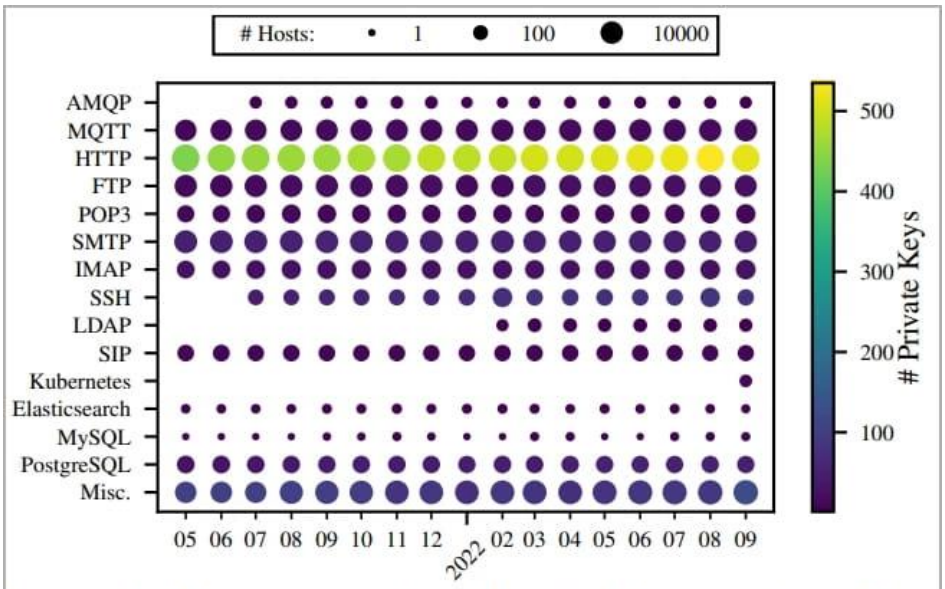

Figure 6: Alarming number of hosts (dot size) relying their authenticity on several compromised keys (dot color) over time (x-axis). Used protocols (y-axis) imply sensitive services.

This level of exposure highlights a massive problem in container security and an carelessness on the creation of images without first sanitizing them of secrets.

Regarding the API exposure, the analysis found that most of the containers (2,920) belong to cloud providers like Amazon AWS, but some pertained to financial services such as Stripe.

However, the researchers cited ethical limitations in validating exposed API secrets against their service endpoints, so their use in the wild is unknown.

Source: *https://www.bleepingcomputer.com/news/security/thousands-of-images-on-docker-hub-leak-auth-secrets-private-keys/*

12. Tracking Down a Suspect through Cell Phone Records

Interesting forensics in connection with a serial killer arrest:

> Investigators went through phone records collected from both midtown Manhattan and the Massapequa Park area of Long Island—two areas connected to a "burner phone" they had tied to the killings. (In court, prosecutors later said the burner phone was identified via an email account used to "solicit and arrange for sexual activity." The victims had all been Craigslist escorts, according to officials.)
>
> They then narrowed records collected by cell towers to thousands, then to hundreds, and finally down to a handful of people who could match a suspect in the killings.
>
> From there, authorities focused on people who lived in the area of the cell tower and also matched a physical description given by a witness who had seen the suspected killer.
>
> In that narrowed pool, they searched for a connection to a green pickup truck that a witness had seen the suspect driving, the sources said.
>
> Investigators eventually landed on Heuermann, who they say matched a witness' physical description, lived close to the Long Island cell site and worked near the New York City cell sites that captured the other calls.
>
> They also learned he had often driven a green pickup truck, registered to his brother, officials said. But they needed more than just circumstantial evidence.
>
> Investigators were able to obtain DNA from an immediate family member and send it to a specialized lab, sources said. According to the lab report,

Heuermann's family member was shown to be related to a person who left DNA on a burlap sack containing one of the buried victims.

There's nothing groundbreaking here; it's casting a wide net with cell phone geolocation data and then winnowing it down using other evidence and investigative techniques. And right now, those are expensive and time consuming, so only used in major crimes like murder (or, in this case, murders).

What's interesting to think about is what happens when this kind of thing becomes cheap and easy: when it can all be done through easily accessible databases, or even when an AI can do the sorting and make the inferences automatically. Cheaper digital forensics means more digital forensics, and we'll start seeing this kind of thing for even routine crimes. That's going to change things.

Source: *https://www.schneier.com/blog/archives/2023/07/tracking-down-a-suspect-through-cell-phone-records.html*

13. Cybersecurity firm Sophos impersonated by new SophosEncrypt ransomware

Cybersecurity vendor Sophos is being impersonated by a new ransomware-as-a-service called SophosEncrypt, with the threat actors using the company name for their operation.

Discovered yesterday by MalwareHunterTeam, the ransomware was initially thought to be part of a red team exercise by Sophos.

However, the Sophos X-Ops team tweeted that they did not create the encryptor and that they are investigating its launch.

"We found this on VT earlier and have been investigating. Our preliminary findings shows Sophos InterceptX protects against these ransomware samples," tweeted Sophos.

Furthermore, ID Ransomware shows one submission from infected victims, indicating that this Ransomware-as-a-Service operation is active.

While little is known about the RaaS operation and how it is being promoted, a sample of the encryptor was found by MalwareHunterTeam, allowing us to get a quick look at how it operates.

The SophosEncrypt ransomware

The ransomware encryptor is written in Rust and uses the 'C:\Users\Dubinin\' path for its crates. Internally, the ransomware is named 'sophos_encrypt,' so it has been dubbed SophosEncrypt, with detections already added to ID Ransomware.

When executed, the encryptor prompts the affiliate to enter a token associated with the victim that is likely first retrieved from the ransomware management panel.

When a token is entered, the encryptor will connect to 179.43.154.137:21119 and verify if the token is valid. Ransomware expert Michael Gillespie found it possible to bypass this verification by disabling your network cards, effectively running the encryptor offline.

When a valid token is entered, the encryptor will prompt the ransomware affiliate for additional information to be used when encrypting the device.

This information includes a contact email, jabber address, and a 32-character password, which Gillespie says is used as part of the encryption algorithm.

The encryptor will then prompt the affiliate to encrypt one file or encrypt the entire device, as shown below.

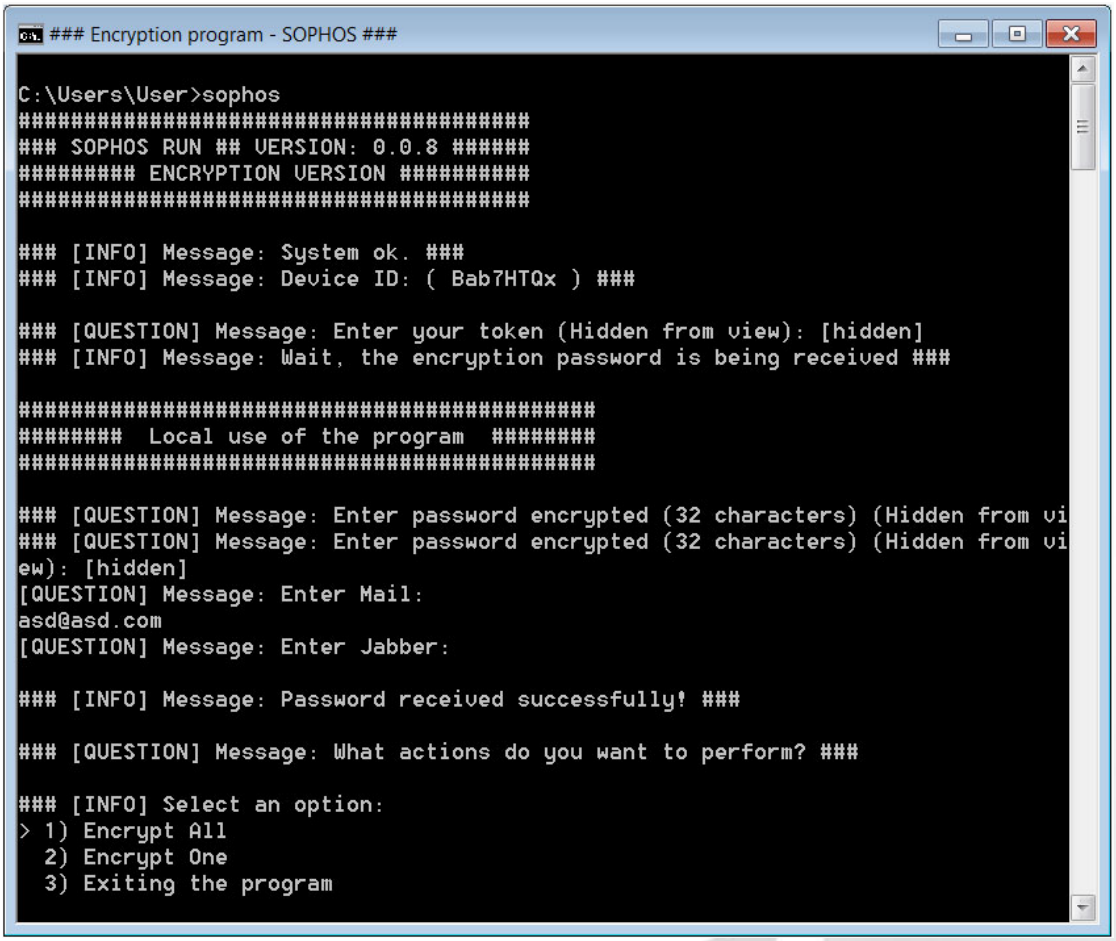

Encryptor prompting for information before encrypting
Source: BleepingComputer

When encrypting files, Gillespie told BleepingComputer that it uses AES256-CBC encryption with PKCS#7 padding.

Each encrypted file will have the entered token, the entered email, and the **sophos** extension appended to a file's name in the format :.[[]].[[]].sophos. This is illustrated below in a test encryption by BleepingComputer.

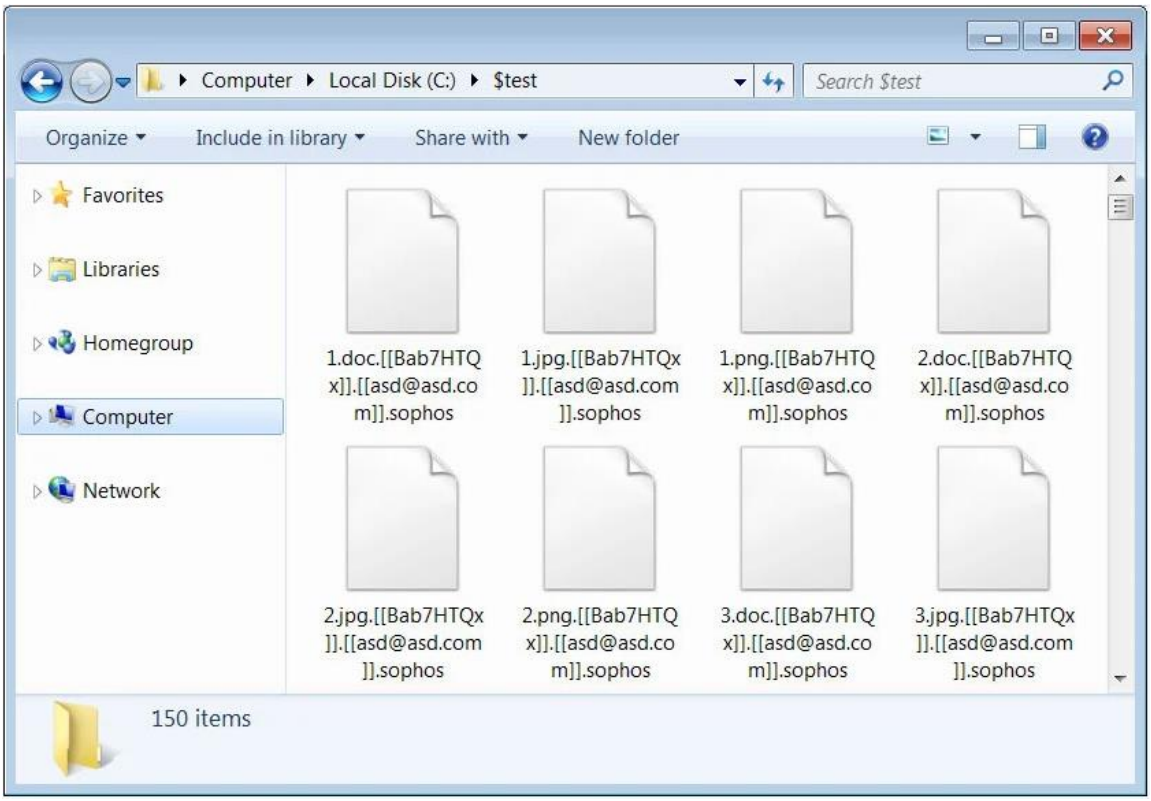

Files encrypted by the SophosEncrypt
Source: BleepingComputer

In each folder that a file is encrypted, the ransomware will create a ransom note named information.hta, which is automatically launched when the encryption is finished.

This ransom note contains information on what happened to a victim's files and the contact information entered by the affiliate before encrypting the device.

SophosEncrypt ransom note
Source: BleepingComputer

The ransomware also has the capability to change the Windows desktop wallpaper, with the current wallpaper boldly displaying the 'Sophos' brand that it is impersonating.

To be clear, this wallpaper was created by the threat actors and has no association with the legitimate Sophos cybersecurity company.

SophosEncrypt wallpaper
Source: BleepingComputer

The encryptor contains numerous references to a Tor site located at http://xnfz2jv5fk6dbvrsxxf3dloi6by3agwtur2fauydd3hwdk4vmm27k7ad.onion.

This Tor site is not a negotiation or data leak site but rather what appears to be the affiliate panel for the ransomware-as-a-service operation.

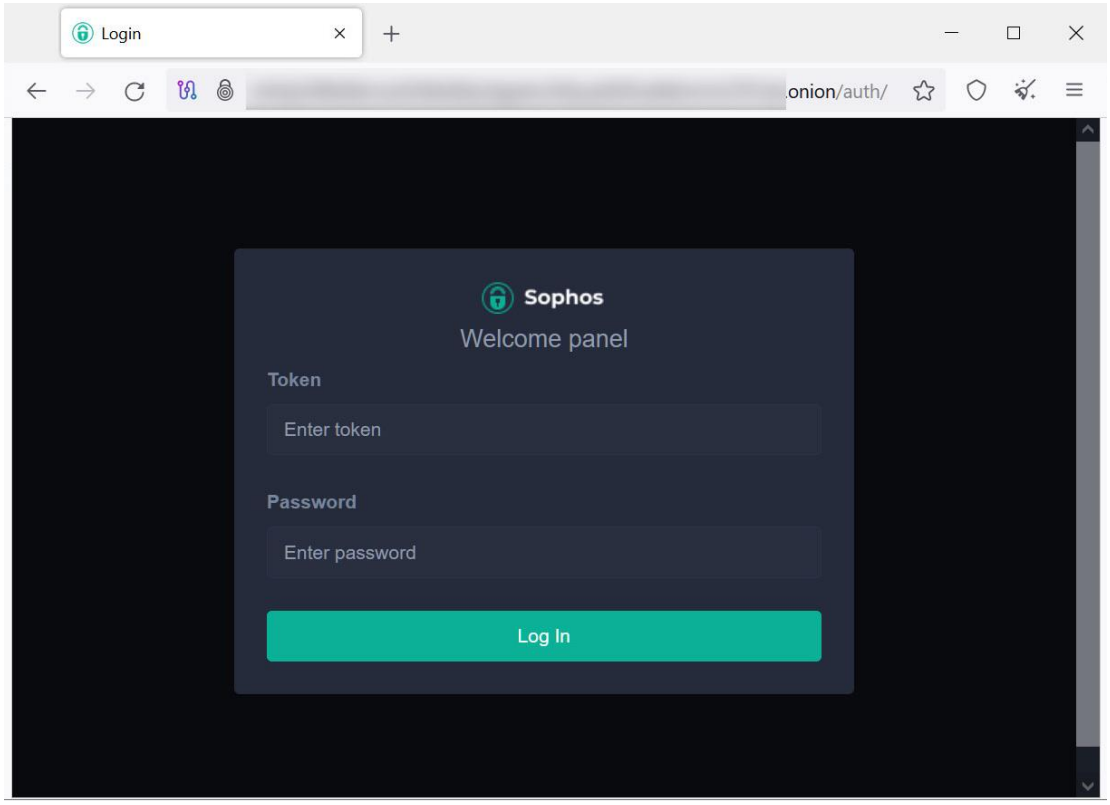

Ransomware-as-a-Service affiliate panel
Source: BleepingComputer

Researchers are still analyzing the SophosEncrypt to see if any weaknesses could allow the recovery of files for free.

If any weaknesses, or encryption issues, are found, we will publish an update to this article.

Update 7/18/23: After publication of our story, Sophos also released a report on the new SophosEncrypt ransomware.

According to their report, the ransomware gang's command and control server at 179.43.154.137 is also linked to Cobalt Strike C2 servers used in previous attacks.

"In addition, both samples contain a hardcoded IP address (one we did see the samples connect to)," explains Sophos' report.

"The address has been associated for more than a year with both Cobalt Strike command-and-control and automated attacks that attempt to infect internet-facing computers with cryptomining software."

Source: https://www.bleepingcomputer.com/news/security/cybersecurity-firm-sophos-impersonated-by-new-sophosencrypt-ransomware/

14. Stolen Azure AD key offered widespread access to Microsoft cloud services

The Microsoft consumer signing key stolen by Storm-0558 Chinese hackers provided them with access far beyond the Exchange Online and Outlook.com accounts that Redmond said were compromised, according to Wiz security researchers.

Redmond revealed on July 12th that the attackers had breached the Exchange Online and Azure Active Directory (AD) accounts of around two dozen organizations. This was achieved by exploiting a now-patched zero-day validation issue in the GetAccessTokenForResourceAPI, allowing them to forge signed access tokens and impersonate accounts within the targeted organizations.

The affected entities included government agencies in the U.S. and Western European regions, with the U.S. State and Commerce Departments among them.

On Friday, Wiz security researcher Shir Tamari said that the impact extended to all Azure AD applications operating with Microsoft's OpenID v2.0. This was due to the stolen key's ability to sign any OpenID v2.0 access token for personal accounts (e.g., Xbox, Skype) and multi-tenant AAD apps.

Microsoft clarified after the publishing of this article that it only impacted those that accepted personal accounts and had the validation error.

While Microsoft said that only Exchange Online and Outlook were impacted, Wiz says the threat actors could use the compromised Microsoft consumer signing key to impersonate any account within any impacted customer or cloud-based Microsoft application.

"This includes managed Microsoft applications, such as Outlook, SharePoint, OneDrive, and Teams, as well as customers' applications that support Microsoft Account authentication, including those who allow the 'Login with Microsoft' functionality," Tamari said.

"Everything in the world of Microsoft leverages Azure Active Directory auth tokens for access," Wiz CTO and Cofounder Ami Luttwak also told BleepingComputer.

"An attacker with an AAD signing key is the most powerful attacker you can imagine, because they can access almost any app – as any user. This is the ultimate cyber intelligence' shape shifter' superpower."

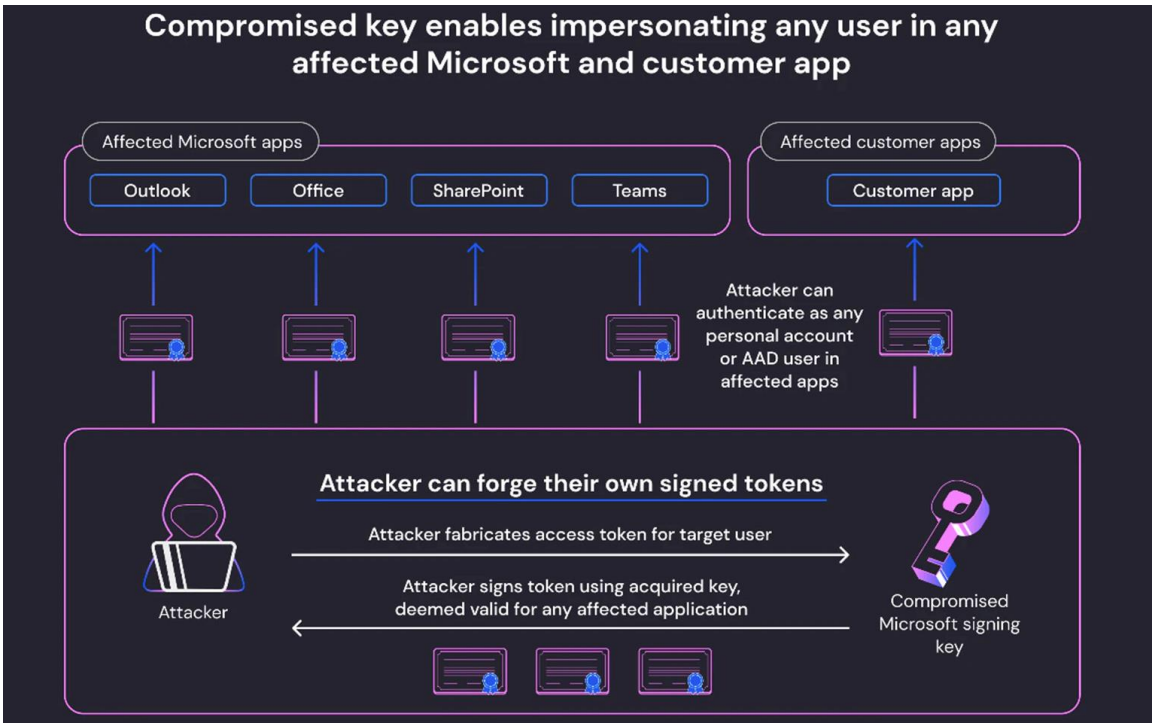

Compromised Microsoft signing key impact (Wiz)

In response to the security breach, Microsoft revoked all valid MSA signing keys to ensure that the threat actors didn't have access to other compromised keys.

This measure also thwarted any attempts to generate new access tokens. Further, Redmond relocated the newly generated access tokens to the key store for the company's enterprise systems.

After invalidating the stolen signing key, Microsoft found no further evidence suggesting additional unauthorized access to its customers' accounts using the same auth token forging technique.

Additionally, Microsoft reported observing a shift in Storm-0558 tactics, showing that the threat actors no longer had access to any signing keys.

Last but not least, the company revealed last Friday that it still doesn't know how the Chinese hackers stole the Microsoft consumer signing key. However, after pressure from CISA, they agreed to expand access to cloud logging data for free to help defenders detect similar breach attempts in the future.

Before this, these logging capabilities were only available to Microsoft customers who paid for Purview Audit (Premium) logging license. As a result, Microsoft faced considerable criticism for impeding organizations from promptly detecting Storm-0558 attacks.

"At this stage, it is hard to determine the full extent of the incident as there were millions of applications that were potentially vulnerable, both Microsoft apps and customer apps, and the majority of them lack the sufficient logs to determine if they were compromised or not," Tamari concluded today.

Update 7/22/23: Updated article with clarifications from Microsoft.

Source: *https://www.bleepingcomputer.com/news/security/stolen-azure-ad-key-offered-widespread-access-to-microsoft-cloud-services/*

15. Lazarus hackers hijack Microsoft IIS servers to spread malware

The North Korean state-sponsored Lazarus hacking group is breaching Windows Internet Information Service (IIS) web servers to hijack them for malware distribution.

IIS is Microsoft's web server solution used to host websites or application services, such as Microsoft Exchange's Outlook on the Web.

South Korean security analysts at ASEC previously reported that Lazarus was targeting IIS servers for initial access to corporate networks. Today, the cybersecurity company says that the threat group leverages poorly protected IIS services for malware distribution too.

The main advantage of this technique is the ease of infecting visitors of websites or users of services hosted on breached IIS servers owned by trustworthy organizations.

Attacks on South Korea

In the recent attacks observed by ASEC's analysts, Lazarus compromised legitimate South Korean websites to perform 'Watering Hole' attacks on visitors using a vulnerable version of the INISAFE CrossWeb EX V6 software.

Many public and private organizations in South Korea use this particular software for electronic financial transactions, security certification, internet banking, etc.

The INISAFE vulnerability was previously documented by both Symantec and ASEC in 2022, explaining that it was exploited using HTML email attachments at the time.

"A typical attack begins when a malicious HTM file is received, likely as a malicious link in an email or downloaded from the web. The HTM file is copied to a DLL file called scskapplink.dll and injected into the legitimate system management software INISAFE Web EX Client," explains the 2022 report by Symantec.

Exploiting the flaw fetches a malicious 'SCSKAppLink.dll' payload from an IIS web server already compromised before the attack for use as a malware distribution server.

"The download URL for 'SCSKAppLink.dll' was identified as being the aforementioned IIS web server," explains ASEC's new report.

"This signifies that the threat actor attacked and gained control over IIS web servers before using these as servers for distributing malware."

ASEC did not analyze the particular payload but says it is likely a malware downloader seen in other recent Lazarus campaigns.

Next, Lazarus uses the 'JuicyPotato' privilege escalation malware ('usopriv.exe') to gain higher-level access to the compromised system.

```
C:₩ProgramData>usopriv.exe


        JuicyPotatoNG
        by decoder_it & splinter_code



        JuicyPotatoNG
        by decoder_it & splinter_code


Mandatory args:
-t createprocess call: <t> CreateProcessWithTokenW, <u> CreateProcessAsUser, <*> try both
-p <program>: program to launch


Optional args:
-l <port>: COM server listen port (Default 10247)
-a <argument>: command line argument to pass to program (default NULL)
-c <CLSID>: (Default {854A20FB-2D44-457D-992F-EF13785D2B51})
-i : Interactive Console (valid only with CreateProcessAsUser)


Additional modes:
-b : Bruteforce all CLSIDs. !ALERT: USE ONLY FOR TESTING. About 1000 processes will be spawned!
-s : Seek for a suitable COM port not filtered by Windows Defender Firewall
```

JuicyPotato in action (ASEC)

JuicyPotato is used for executing a second malware loader ('usoshared.dat') that decrypts downloaded data files and executes them into memory for AV evasion.

```
kernel32 = fn_getModule(L"Kernel32.dll");
HeapAlloc = fn_getProc(kernel32, "HeapAlloc");
kernel32_1 = fn_getModule(L"Kernel32.dll");
GetProcessHeap = fn_getProc(kernel32_1, "GetProcessHeap");
user32 = fn_getModule(L"User32.dll");
fn_getProc(user32, "wsprintfW");
hHeap = GetProcessHeap();
result = HeapAlloc(hHeap, 8i64, 32i64);
mem_newAlloc = result;
if ( result )
{
  *result = 0;
  if ( a1 == 3 )
  {
    result[4] = data_sizeOfPE;
    *(result + 1) = data_decodedPE;
  }
  else
  {
    data_decodedPE = *(result + 1);
  }
  if ( fn_checkPE(data_decodedPE) && fn_allocMem(mem_newAlloc) && fn_resolveAPI(mem_newAlloc) )
  {
    if ( fn_runMem(mem_newAlloc, data_config) )
      *mem_newAlloc = 1;
```

Loading the decrypted executable in memory (ASEC)

ASEC recommends that NISAFE CrossWeb EX V6 users update the software to its latest version, as Lazarus' exploitation of known vulnerabilities in the product has been underway since at least April 2022.

The security company advises users to upgrade to version 3.3.2.41 or later and points to remediation instructions it posted four months ago, highlighting the Lazarus threat.

Microsoft application servers are becoming a popular target for hackers to use in malware distribution, likely due to their trusted nature.

Just last week, CERT-UA and Microsoft reported that Russian Turla hackers were using compromised Microsoft Exchange servers to deliver backdoors to their targets.

Source: https://www.bleepingcomputer.com/news/security/lazarus-hackers-hijack-microsoft-iis-servers-to-spread-malware/

16. Super Admin elevation bug puts 900,000 MikroTik devices at risk

A critical severity 'Super Admin' privilege elevation flaw puts over 900,000 MikroTik RouterOS routers at risk, potentially enabling attackers to take full control over a device and remain undetected.

The flaw, CVE-2023-30799, allows remote attackers with an existing admin account to elevate their privileges to "super-admin" via the device's Winbox or HTTP interface.

A VulnCheck report published today explains that while CVE-2023-30799 requires an existing admin account to exploit, this is not a high bar to clear.

This is because the Mikrotik RouterOS operating system does not prevent password brute-force attacks and comes with a well-known default "admin" user.

"'En masse' exploitation is going to be more difficult since valid credentials are required. However, as I outlined in the blog, the routers lack basic protections against password guessing," VulnCheck researcher Jacob Baines told BleepingComputer.

"We intentionally didn't release a proof-of-concept exploit, but if we had, I have no doubt that the exploit would have been successfully used in the wild quickly after the blog was released."

A large-scale problem

The Mikrotik CVE-2023-30799 vulnerability was first disclosed without an identifier in June 2022, and MikroTik fixed the issue in October 2022 for RouterOS stable (v6.49.7) and on July 19, 2023, for RouterOS Long-term (v6.49.8).

VulnCheck reports that a patch for the Long-term branch was made available only after they contacted the vendor and shared new exploits that targeted MikroTik hardware.

The researchers used Shodan to determine the flaw's impact and found that 474,000 devices were vulnerable as they remotely exposed the web-based management page.

However, as this vulnerability is also exploitable over Winbox, a Mikrotek management client, Baines found that 926,000 devices were exposing this management port, making the impact far larger.

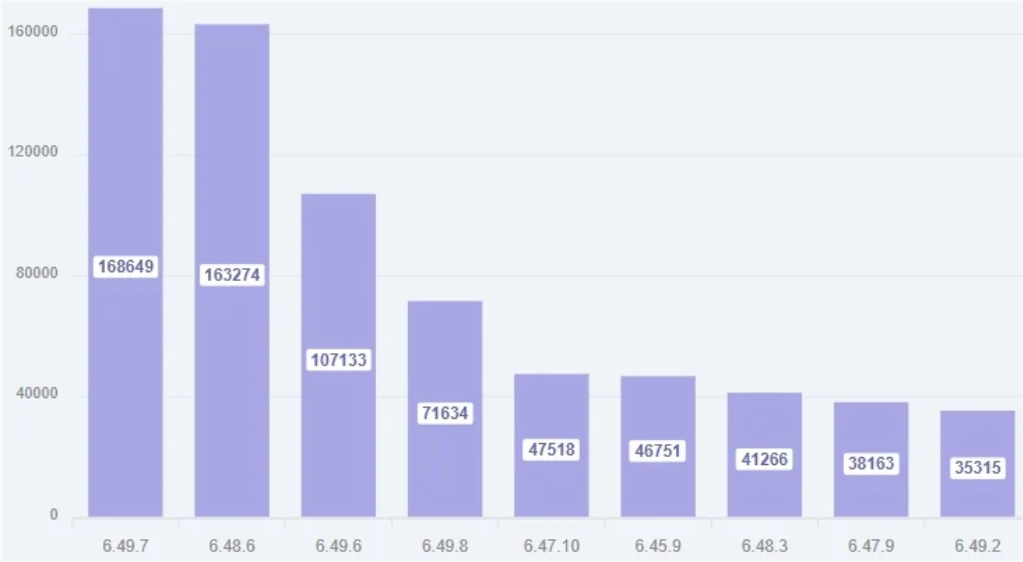

Detected RouterOS versions (VulnCheck)

The CVE-2023-30799 vulnerability

While exploiting this vulnerability requires an existing admin account, it elevates you to a higher privilege level called "Super Admin."

Unlike the admin account, which offers restricted elevated privileges, Super Admin gives full access to the RouteOS operating system.

"By escalating to super admin, the attacker can reach a code path that allows them to control the address of a function call," Baines told BleepingComputer.

"Super admin is not a privilege given to normal administrators, it's a privilege that is supposed to be given to certain parts of the underlying software (specifically, in this case, to load libraries for the web interface), and not to actual users.

This makes the vulnerability valuable to threat actors wishing to "jailbreak" the RouterOS device to make significant changes to the underlying operating system or hide their activities from detection.

To develop an exploit for CVE-2023-30799 that obtains a root shell on MIPS-based MikroTik devices, VulnCheck's analysts used Margin Research's FOISted remote RouterOS jailbreak exploit.

The new exploit developed by VulnCheck bypasses the requirement for FTP interface exposure and is not impacted by blocking or filtering of bindshells, as it uses the RouterOS web interface to upload files.

Finally, VulnCheck identified a simplified ROP chain that manipulates the stack pointer and the first argument register and calls dlopen, the instructions for which are present in three functions across different RouterOS versions, ensuring broad applicability.

The exploit still requires authentication as "admin," however, VulnCheck explains that RouterOS ships with a fully functional admin user by default, which nearly 60% of MikroTik devices still use despite the vendor's hardening guidance suggesting its deletion.

Moreover, the default admin password was an empty string until October 2021, when this issue was fixed with the release of RouterOS 6.49.

Finally, RouterOS does not impose admin password strengthening requirements, so users may set anything they like, which makes them susceptible to brute-forcing attacks, for which MikroTik does not offer any protection except on the SSH interface.

"All of this is to say, RouterOS suffers from a variety of issues that make guessing administrative credentials easier than it should be," comments VulnCheck

"We believe CVE-2023-30799 is much easier to exploit than the CVSS vector indicates."

Patch your devices

MikroTik devices have been targeted by malware many times and inadvertently helped build record-breaking DDoS swarms like the Mēris botnet.

Users need to move quickly to patch the flaw by applying the latest update for RouterOS, as attempts to exploit the flaw are bound to increase soon.

Mitigation advice includes removing administrative interfaces from the internet, restricting login IP addresses to a defined allow-list, disabling Winbox and only use SSH, and configuring SSH to use public/private keys instead of passwords.

Source: *https://www.bleepingcomputer.com/news/security/super-admin-elevation-bug-puts-900-000-mikrotik-devices-at-risk/*

17. Almost 40% of Ubuntu users vulnerable to new privilege elevation flaws

Two Linux vulnerabilities introduced recently into the Ubuntu kernel create the potential for unprivileged local users to gain elevated privileges on a massive number of devices.

Ubuntu is one of the most widely used Linux distributions, especially popular in the U.S., having an approximate user base of over 40 million.

Two recent flaws tracked as CVE-2023-32629 and CVE-2023-2640 discovered by Wiz's researchers S. Tzadik and S. Tamari were recently introduced into the operating system, impacting roughly 40% of Ubuntu's userbase.

CVE-2023-2640 is a high-severity (CVSS v3 score: 7.8) vulnerability in the Ubuntu Linux kernel caused by inadequate permission checks allowing a local attacker to gain elevated privileges.

CVE-2023-32629 is a medium-severity (CVSS v3 score: 5.4) flaw in the Linux kernel memory management subsystem, where a race condition when accessing VMAs may lead to use-after-free, allowing a local attacker to perform arbitrary code execution.

The two analysts found the problems after discovering discrepancies in implementing the OverlayFS module onto the Linux kernel.

OverlayFS is a union mount filesystem implementation targeted by threat actors many times in the past due to allowing unprivileged access via user namespaces and being plagued by easily exploitable bugs.

Ubuntu, as one of the distributions using OverlayFS, had implemented custom changes to its OverlayFS module in 2018, which were generally safe.

However, in 2019 and 2022, the Linux kernel project made its own modifications to the module, which conflicted with Ubuntu's changes.

The widespread distribution adopted the code containing these changes recently, and the conflicts caused the introduction of the two flaws.

Unfortunately, the risk of exploitation is imminent, as PoCs for the two flaws have been publicly available for a long time.

"Both vulnerabilities are unique to Ubuntu kernels since they stemmed from Ubuntu's individual changes to the OverlayFS module," warned the Wiz researchers.

"Weaponized exploits for these vulnerabilities are already publicly available given old exploits for past OverlayFS vulnerabilities work out of the box without any changes."

It should be noted that the two highlighted flaws only impact Ubuntu, and any other Linux distribution, including Ubuntu forks, not using custom modifications of the OverlayFS module should be safe.

Ubuntu has released a security bulletin about the issues and six more vulnerabilities addressed in the latest version of the Ubuntu Linux kernel and has made fixing updates available.

Users who don't know how to reinstall and activate third-party kernel modules are recommended to perform the update via their package manager, which should take care of all dependencies and post-install configurations.

A reboot is required after installing the updates for the Linux kernel update to take effect on Ubuntu.

Source: https://www.bleepingcomputer.com/news/security/almost-40-percent-of-ubuntu-users-vulnerable-to-new-privilege-elevation-flaws/

18. AI in the Wild: Malicious Applications of Mainstream AI Tools

It's not all funny limericks, bizarre portraits, and hilarious viral skits. ChatGPT, Bard, DALL-E, Craiyon, Voice.ai, and a whole host of other mainstream artificial intelligence tools are great for whiling away an afternoon or helping you with your latest school or work assignment; however, cybercriminals are bending AI tools like these to aid in their schemes, adding a whole new dimension to phishing, vishing, malware, and social engineering.

Here are some recent reports of AI's use in scams plus a few pointers that might tip you off should any of these happen to you.

1. AI Voice Scams

Vishing – or phishing over the phone – is not a new scheme; however, AI voice mimickers are making these scamming phone calls more believable than ever. In Arizona, a fake kidnapping phone call caused several minutes of panic for one family, as a mother received a demand for ransom to release her alleged kidnapped daughter. On the phone, the mother heard a voice that sounded exactly like her child's, but it turned out to be an AI-generated facsimile.

In reality, the daughter was not kidnapped. She was safe and sound. The family didn't lose any money because they did the right thing: They contacted law enforcement and kept the scammer on the phone while they located the daughter.1

Imposter scams accounted for a loss of $2.6 billion in the U.S. in 2022. Emerging AI scams could increase that staggering total. Globally, about 25% of people have either experienced an AI voice scam or know someone who has, according to McAfee's Beware the Artificial Imposter report. Additionally, the study discovered that 77% of voice scam targets lost money as a result.

How to hear the difference

No doubt about it, it's frightening to hear a loved one in distress, but try to stay as calm as possible if you receive a phone call claiming to be someone in trouble. Do your best to really listen to the "voice" of your loved one. AI voice technology is incredible, but

there are still some kinks in the technology. For example, does the voice have unnatural hitches? Do words cut off just a little too early? Does the tone of certain words not quite match your loved one's accent? To pick up on these small details, a level head is necessary.

What you can do as a family today to avoid falling for an AI vishing scam is to agree on a family password. This can be an obscure word or phrase that is meaningful to you. Keep this password to yourselves and never post about it on social media. This way, if a scammer ever calls you claiming to have or be a family member, this password could determine a fake emergency from a real one.

2. Deepfake Ransom and Fake Advertisements

Deepfake, or the digital manipulation of an authentic image, video, or audio clip, is an AI capability that unsettles a lot of people. It challenges the long-held axiom that "seeing is believing." If you can't quite believe what you see, then what's real? What's not?

The FBI is warning the public against a new scheme where cybercriminals are editing explicit footage and then blackmailing innocent people into sending money or gift cards in exchange for not posting the compromising content.2

Deepfake technology was also at the center of an incident involving a fake ad. A scammer created a fake ad depicting Martin Lewis, a trusted finance expert, advocating for an investment venture. The Facebook ad attempted to add legitimacy to its nefarious endeavor by including the deepfaked Lewis.3

How to respond to ransom demands and questionable online ads

No response is the best response to a ransom demand. You're dealing with a criminal. Who's to say they won't release their fake documents even if you give in to the ransom? Involve law enforcement as soon as a scammer approaches you, and they can help you resolve the issue.

Just because a reputable social media platform hosts an advertisement doesn't mean that the advertiser is a legitimate business. Before buying anything or investing your money with a business you found through an advertisement, conduct your own background research on the company. All it takes is five minutes to look up its Better Business Bureau rating and other online reviews to determine if the company is reputable.

To identify a deepfake video or image, check for inconsistent shadows and lighting, face distortions, and people's hands. That's where you'll most likely spot small details that aren't quite right. Like AI voices, deepfake technology is often accurate, but it's not perfect.

3. AI-generated Malware and Phishing Emails

Content generation tools have some safeguards in place to prevent them from creating text that could be used illegally; however, some cybercriminals have found ways around those rules and are using ChatGPT and Bard to assist in their malware and phishing operations. For example, if a criminal asked ChatGPT to write a key-logging malware, it would refuse. But if they rephrased and asked it to compose code that captures keystrokes, it may comply with that request. One researcher demonstrated that even someone with little knowledge of coding could use ChatGPT, thus making malware creation simpler and more available than ever.4 Similarly, AI text generation tools can create convincing phishing emails and create them quickly. In theory, this could speed up a phisher's operation and widen their reach.

How to avoid AI-written malware and phishing attempts

You can avoid AI-generated malware and phishing correspondences the same way you deal with the human-written variety: Be careful and distrust anything that seems suspicious. To steer clear of malware, stick to websites you know you can trust. A safe browsing tool like McAfee web protection – which is included in McAfee+ – can doublecheck that you stay off of sketchy websites.

As for phishing, when you see emails or texts that demand a quick response or seem out of the ordinary, be on alert. Traditional phishing correspondences are usually riddled with typos, misspellings, and poor grammar. AI-written lures are often written well and rarely contain errors. This means that you must be diligent in vetting every message in your inbox.

Slow Down, Keep Calm, and Be Confident

While the debate about regulating AI heats up, the best thing you can do is to use AI responsibly. Be transparent when you use it. And if you suspect you're encountering a malicious use of AI, slow down and try your best to evaluate the situation with a clear mind. AI can create some convincing content, but trust your instincts and follow the above best practices to keep your money and personal information out of the hands of cybercriminals.

Source: https://www.mcafee.com/blogs/internet-security/ai-in-the-wild-malicious-applications-of-mainstream-ai-tools/

19. Linux version of Abyss Locker ransomware targets VMware ESXi servers

The Abyss Locker operation is the latest to develop a Linux encryptor to target VMware's ESXi virtual machines platform in attacks on the enterprise.

As the enterprise shifts from individual servers to virtual machines for better resource management, performance, and disaster recovery, ransomware gangs create encryptors focused on targeting the platform.

With VMware ESXi being one of the most popular virtual machine platforms, almost every ransomware gang has begun to release Linux encryptors to encrypt all virtual servers on a device.

Other ransomware operations that utilize Linux ransomware encryptors, with most targeting VMware ESXi, include Akira, Royal, Black
Basta, LockBit, BlackMatter, AvosLocker, REvil, HelloKitty, RansomEXX, and Hive.

The Abyss Locker

Abyss Locker is a relatively new ransomware operation that is believed to have launched in March 2023, when it began to target companies in attacks.

Like other ransomware operations, the Abyss Locker threat actors will breach corporate networks, steal data for double-extortion, and encrypt devices on the network.

The stolen data is then used as leverage by threatening to leak files if a ransom is not paid. To leak the stolen files, the threat actors created a Tor data leak site named 'Abyss-data' that currently lists fourteen victims.

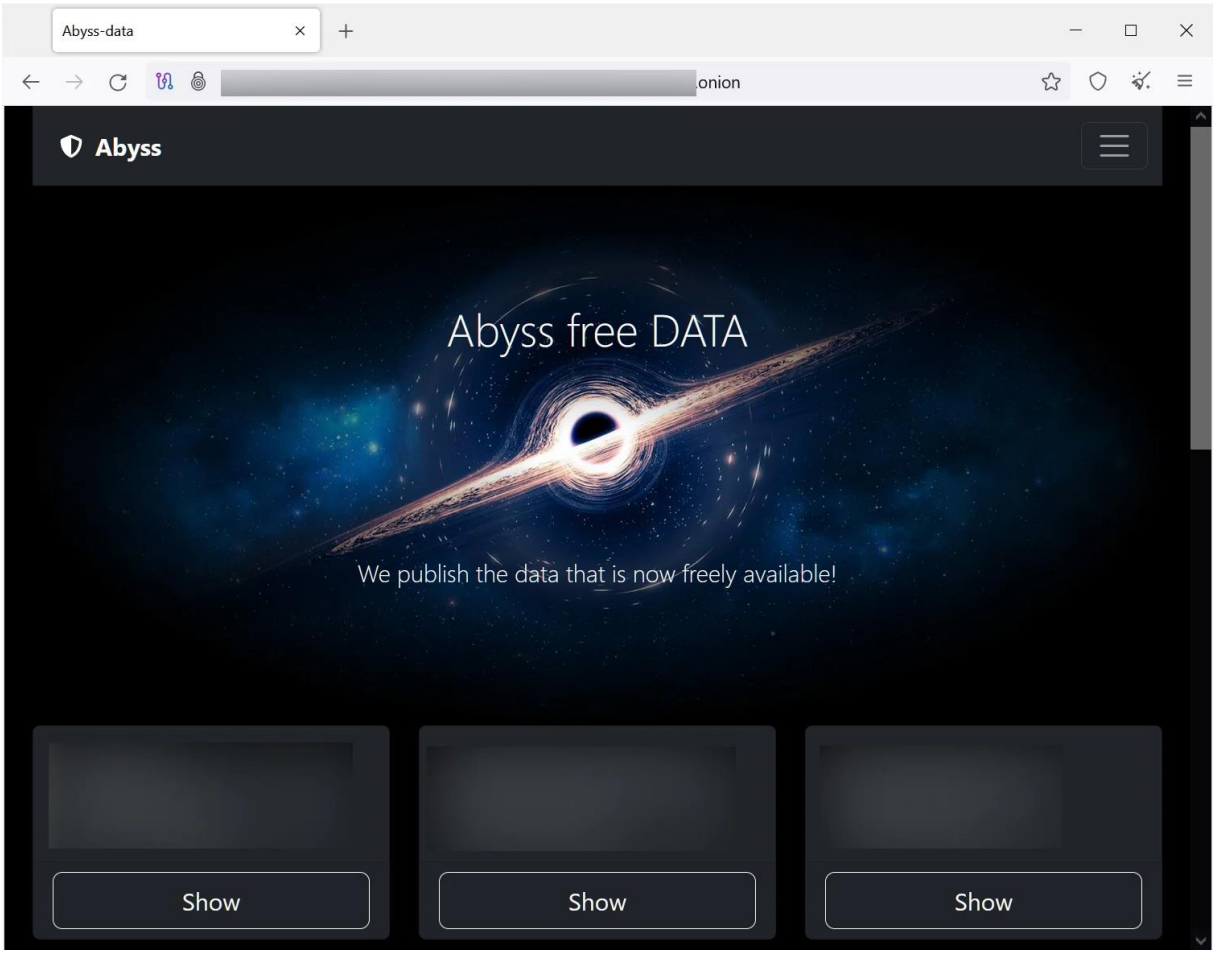

Abyss Locker data leak site
Source: BleepingComputer

The threat actors claim to have stolen anywhere between 35 GB of data from one company to as high as 700 GB at another.

Targeting VMware ESXi servers

This week, security researcher MalwareHunterTeam found a Linux ELF encryptor for the Abyss Locker operation and shared it with BleepingComputer for analysis.

After looking at the strings in the executable, it is clear that the encryptor specifically targets VMware ESXi servers.

As you can see from the commands below, the encryptor utilizes the 'esxcli' command-line VMware ESXi management tool to first list all available virtual machines and then terminate them.

```
esxcli vm process list
esxcli vm process kill -t=soft -w=%d
esxcli vm process kill -t=hard -w=%d
esxcli vm process kill -t=force -w=%d
```

When shutting down the virtual machines, Abyss Locker will use the 'vm process kill' command and one of the soft, hard, or forced options.

The **soft** option performs a graceful shutdown, the **hard** option terminates a VM immediately, and **force** is used as a last resort.

The encryptor terminates all virtual machines to allow the associated virtual disks, snapshots, and metadata to be properly encrypted by encrypting all files with the following extensions: .vmdk (virtual disks), .vmsd (metadata), and .vmsn (snapshots).

In addition to targeting virtual machines, the ransomware will also encrypt all other files on the device and append the **.crypt** extension to their filenames, as shown below.

Encrypted files and ransom notes
Source: BleepingComputer

For each file, the encryptor will also create a file with a **.README_TO_RESTORE** extension, which acts as the ransom note.

This ransom note contains information on what happened to the files and a unique link to the threat actor's Tor negotiation site. This site is barebones, only having a chat panel that can be used to negotiate with the ransomware gang.

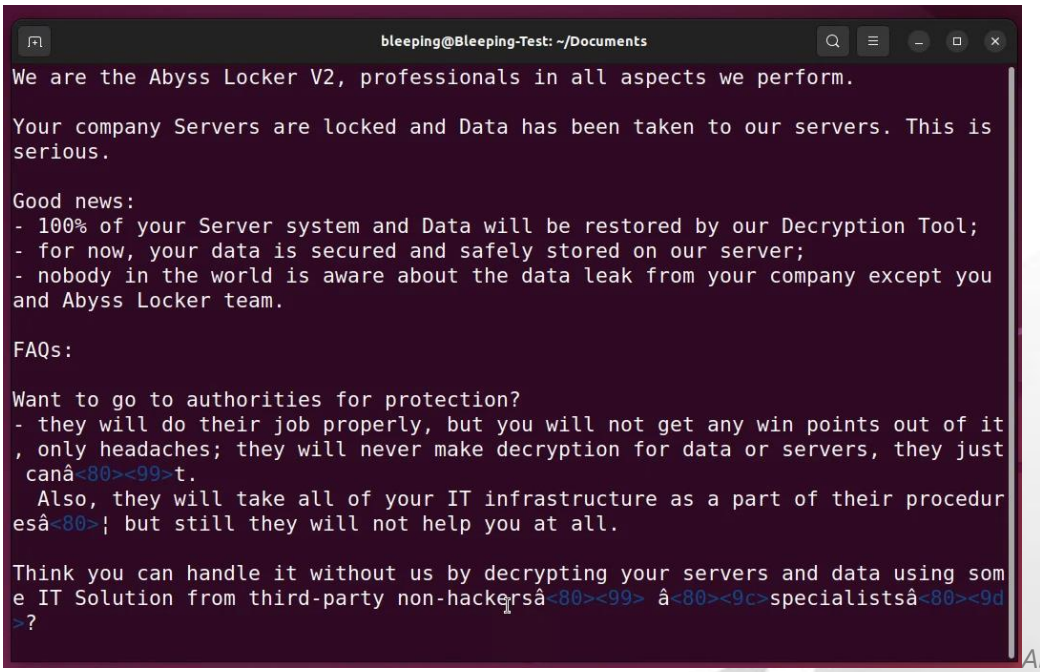

Abyss Locker ransom note
Source: BleepingComputer

Ransomware expert Michael Gillespie said that the Abyss Locker Linux encryptor is based on Hello Kitty, using ChaCha encryption instead.

However, it is not known if this is a rebrand of the HelloKitty operation or if another ransomware operation gained access to the encryptor's source code, as we saw with Vice Society.

Unfortunately, HelloKitty has historically been a secure ransomware, preventing the recovery of files for free.

Source: *https://www.bleepingcomputer.com/news/security/linux-version-of-abyss-locker-ransomware-targets-vmware-esxi-servers/*

If you want to learn more about ASOC and how we can improve your security posture, **contact us at tbs.sales@tbs.tech**.