Monthly CYBER THREAT INTELLIGENCE REPORT

May 2026

v1.0 May 2026

Prepared by
System Force IT Ltd

Classification:
TLP:AMBER

System Force I.T.
Secure IT Simplified

UNDERSTANDING CYBER THREAT INTELLIGENCE

Cyber threat intelligence (CTI) is the practice of gathering and analysing information about adversaries, their campaigns and the malware they use, then turning that analysis into guidance that helps an organisation reduce risk. A mature CTI capability gives everyone, from the boardroom to the service desk, a shared language for discussing threats and a clear set of actions for shrinking the attack surface.

Intelligence is usually grouped into three layers, each written for a different audience:

Strategic: a high-level picture for directors and decision makers, focused on business risk and the policy or investment choices needed to manage it.

Operational: insight into live or developing campaigns, including attacker goals and timing, used by incident responders and managers to prepare and react.

Tactical: the technical detail, such as indicators of compromise (IOCs) and attacker tactics, techniques and procedures (TTPs), used by engineers and analysts to hunt for and block specific threats.

This report blends all three layers, so it is useful whether you are setting security strategy or tuning a detection rule.

MONTHLY CRITICAL NEWS SUMMARY

The most significant publicly reported incidents of the month, with a severity rating and the themes each touches on.

Title	Description	Severity	Related
TeamPCP “Supply Chain Competition”	A criminal forum advertised a contest, run jointly with the threat group TeamPCP, that pays a Monero bounty (around USD 1,000) to whoever compromises the most widely used open-source packages. To drive participation, the organisers released a worm (“Shai Hulud”) on the forum’s own infrastructure and made its use mandatory to enter. Pairing a cash incentive with ready-made tooling effectively industrialises supply chain attacks against the open-source ecosystem.	Critical	Cybercrime; Supply chain
~3,800 GitHub internal repositories breached	On 19 May, TeamPCP claimed access to a large volume of private GitHub source code and offered roughly 4,000 repositories for sale from USD 50,000. GitHub later confirmed about 3,800 internal repositories were exposed. The cause was traced to an employee installing a trojanised Nx Console extension, poisoned via an npm supply chain compromise affecting TanStack; stolen CI/CD credentials then enabled spread to further projects. GitHub removed the extension and secured the device.	Critical	GitHub; Data breach; TeamPCP
CISA contractor leaks AWS GovCloud keys	A contractor linked to the US CISA left a public GitHub repository (“Private-CISA”) containing administrative AWS GovCloud access keys, plaintext passwords, tokens and internal build data. Poor practice contributed throughout: clear-text passwords, secret scanning turned off, and the repository used to sync data across environments. It stayed public for an extended period, and some credentials reportedly remained valid for up to 48 hours after disclosure.	Critical	Credential exposure; Government
Zara customer data breach	The extortion group ShinyHunters claimed a breach affecting roughly 197,000 Zara customers, traced to a database held by a former technology supplier. A 140GB archive of email addresses, locations, purchase histories and support tickets was published free on a criminal forum to promote a new platform. Parent company Inditex stated its own infrastructure was not breached and that passwords and payment details were not exposed.	High	Data breach; ShinyHunters

Title	Description	Severity	Related
“fast16” pre-Stuxnet sabotage malware	New research describes “fast16”, a Lua-based tool from 2005 that predates Stuxnet. It was built to quietly interfere with nuclear weapons testing by altering engineering software so that high-explosive simulation calculations produced corrupted results, using a large hook set to survive updates and spread across networked machines. It is a reminder that highly targeted industrial sabotage has been a real capability for over two decades.	High	OT tampering
Sandworm pivots into OT	The Russian state-linked group Sandworm is increasingly moving from compromised IT networks into operational technology. Rather than relying on zero-days, it exploits already-compromised environments, using older exploit chains and unresolved commodity infections as a foothold before targeting engineering workstations and industrial control systems. Affected systems generated high-confidence alerts for an average of 43 days before lateral movement, underlining why routine alerts should be treated as early warnings.	Critical	Russian state actor; OT
RubyGems pauses new sign-ups	The Ruby package registry temporarily halted new registrations after coordinated abuse. Bot accounts flooded it with more than 500 malicious junk packages, while a separate campaign (“GemStuffer”) abused the registry as a covert data drop, storing information scraped from UK council portals to avoid running its own command-and-control. RubyGems paused sign-ups to add rate limiting and web application firewall protection; existing users and packages were reported unaffected.	Critical	Supply chain; Ruby
Foxconn ransomware attack	Electronics manufacturer Foxconn confirmed an attack that briefly disrupted several North American sites. The Nitrogen group claimed responsibility, stating it stole 8TB across more than 11 million files, including confidential drawings and project material for major clients. Foxconn restored production. Nitrogen, active since 2023 and using leaked Conti source code, is pursuing double extortion, showing how compromising a manufacturing hub lets attackers pressure many downstream clients at once.	Critical	Ransomware; Nitrogen

HIGH-SEVERITY VULNERABILITY REVIEW

The month's highest-scoring vulnerabilities, with affected products and recommended remediation. Internet-facing edge devices and domain controllers should be prioritised.

CVE	Description	Affected	Remediation	CVSS
CVE-2026-0300	Palo Alto PAN-OS User-ID portal: buffer overflow allowing an unauthenticated attacker to run code with root privileges via crafted packets.	PAN-OS 10.2 and 11.x	Upgrade to a fixed PAN-OS release, or restrict access to the User-ID portal.	**9.8**
CVE-2026-20182	Cisco Catalyst SD-WAN Manager and Controller: authentication bypass granting admin rights and the ability to alter network configuration.	Versions before 20.18.2.2 / 26.1.1.1	Upgrade to a fixed SD-WAN release.	**10.0**
CVE-2026-41096	Microsoft Windows DNS: heap buffer overflow allowing remote code execution via crafted DNS responses.	Windows 11 / Server before May 2026 updates	Apply the May 2026 security updates.	**9.8**
CVE-2026-44277	Fortinet FortiAuthenticator: improper access control on API endpoints permitting unauthenticated command execution.	6.5.0–6.5.6; 6.6.0–6.6.8; 8.0.0–8.0.2	Upgrade to 6.5.7+, 6.6.9+ or 8.0.3+.	**9.8**
CVE-2026-23918	Apache HTTP Server (HTTP/2): double-free that may allow remote code execution via crafted requests.	Apache HTTP Server 2.4.66	Upgrade to 2.4.67 or later.	**8.8**
CVE-2026-0257	Palo Alto PAN-OS GlobalProtect: authentication bypass allowing an unauthorised VPN connection.	PAN-OS with GlobalProtect before fixed 10.2/11.1/11.2/12.1 releases	Upgrade to a fixed release (e.g. 10.2.18-h6+, 11.1.15+, 11.2.12+, 12.1.7+).	**7.8**
CVE-2026-41089	Windows Netlogon: stack buffer overflow allowing remote code execution on a domain controller via crafted requests.	Windows Server 2012–2025 acting as DCs with Netlogon	Apply the May 2026 Microsoft updates for affected Server versions.	**9.8**
CVE-2026-9082	Drupal database API: SQL injection on PostgreSQL-backed sites, potentially leading to data exposure, privilege escalation or RCE.	8.9.0 through 11.3.9 (multiple ranges)	Update to 11.1.10 / 11.2.12 / 11.3.10 or 10.4.10 / 10.5.10 / 10.6.9; apply vendor patches for 9.5 and 8.9.	**9.8**

RANSOMWARE ACTIVITY REVIEW

Top ten groups by number of claimed victims during May 2026.

Ransomware Group	Claimed Victims
Qilin	**110**
The Gentlemen	**77**
DragonForce	**55**
Akira	**31**
INC Ransom	**29**
Nova	**25**
Fulcrumsec	**23**
SafePay	**22**
Genesis	**21**
Cmd Organization	**16**
Total claimed victims (May 2026)	**704**
Most-targeted country	**United States**
Most-targeted sector	**Business Services**

VULNERABILITY SPOTLIGHT: COPY FAIL (CVE-2026-31431)

CVE-2026-31431, nicknamed Copy Fail, is a high-severity local privilege escalation flaw in the Linux kernel (CVSS 7.8). A proof of concept was disclosed on 29 April and the issue was quickly added to the CISA Known Exploited Vulnerabilities (KEV) catalogue.

It affects a wide range of kernels (4.14 through 6.19.12) across most mainstream distributions, including Ubuntu, Debian, Red Hat Enterprise Linux, SUSE, AlmaLinux, Fedora, Amazon Linux and Arch.

How it works

Linux keeps recently used file content in the page cache so it can be served from memory rather than re-read from disk. Copy Fail lets a low-privileged local user trigger a small, controlled write into that cached copy of a file, through an unintended interaction between several kernel components (the AF_ALG crypto socket interface, an in-place authenticated-encryption code path, and the splice mechanism that brings file-backed pages into the operation). Critically, the file on disk is never modified; only the in-memory copy is altered.

Although the corruption is only a few bytes, that can be enough to change how a privileged (“setuid-root”) program behaves in memory, for example by flipping a security check, allowing escalation to root. Because the on-disk file is untouched, tools that rely on comparing disk hashes may miss it. Compared with older issues such as Dirty COW, Copy Fail is more reliable because it does not depend on winning a timing race.

Mitigation

- Apply the kernel updates published by your distribution vendor as a priority. This includes container hosts and Kubernetes nodes, where the shared kernel can expose multiple workloads to the same flaw.

- Where immediate patching is not possible, the recommended temporary step is to disable the affected crypto module (algif_aead) until the update can be applied.

Note: *published exploit code has been omitted from this briefing by design. The priority for clients is patching, not reproduction.*

MALWARE DEEP DIVE: M3RX RANSOMWARE

Executive summary

M3RX is a newer ransomware-as-a-service (RaaS) operation seen targeting organisations across North America, Western Europe and parts of Asia, spanning manufacturing, healthcare, professional services, technology and education.

It uses double extortion (encrypting files and stealing data) and is written in Go, a language increasingly favoured by attackers for its portability and ease of deployment.

Static analysis

Examining the file and its strings reveals the following behaviours:

- Checks for administrative rights by inspecting its process token before attempting high-impact actions.
- Uses a mutex to ensure only one instance runs at a time.
- Exposes configurable options (target path, delay, thread count, encryption percentage, stealth and logging), indicating an operator-driven design; it can hide its console window and keep a progress log in the user profile.
- Maintains exclusion lists for critical system files and extensions (for example ntldr, bootmgr, ntuser.dat, and types such as .exe, .dll, .sys, .msi and .lnk) and for key directories (Windows, Program Files, ProgramData, Users, AppData), so it avoids breaking the operating system or its own ability to run.
- Generates randomised alphanumeric strings, likely for renaming files and creating mutexes, and shuts down in an orderly way so threads finish without corrupting output.
- For impact and anti-forensics, it deletes shadow copies, empties the recycle bin, overwrites data and then deletes itself.

Dynamic analysis

- When run without elevation it detects a medium integrity level and exits; run as administrator it proceeds and spawns conhost.exe.
- Its command-line help confirms options for delay, fast mode, hiding, logging, target path, encryption percentage (default 1%), thread count and timing.
- It encrypts local drives, reporting disk space and live progress (files processed, throughput, success rate) at the configured encryption depth.
- On interruption it waits for threads to finish cleanly.
- After encryption it renames files with randomised extensions and drops a ransom note (RECOVERY_NOTES.TXT) across user directories such as Desktop and Documents for multiple accounts.
- It uses PowerShell to overwrite and delete its own executable, an anti-forensics step.
- The note references both encryption and data theft and points victims to a Tor portal and leak site, confirming double extortion.

MITRE ATT&CK

Tactic	Techniques
Execution	Command and Scripting Interpreter; Native API
Defense Evasion	Indicator Removal
Discovery	File and Directory Discovery
Impact	Data Encrypted for Impact; Inhibit System Recovery

Detection and response

This activity presents several detection opportunities for a layered managed detection and response (MDR) service:

- The malicious binary is flagged by AV/AI engines when it is written to disk or attempts to run.
- Third-party threat intelligence corroborates the file as known-malicious.
- Behavioural monitoring detects unauthorised file operations, including tampering with decoy (“canary”) files and the appearance of suspicious extensions.

MALWARE DEEP DIVE: EVELYN STEALER

Executive summary

Evelyn is an information stealer first observed in December 2025.

It uses extensive anti-VM and anti-sandbox techniques to avoid analysis and targets a broad range of sensitive data, including browsers, cryptocurrency wallets, messaging apps, VPN clients and games.

Static analysis

Embedded strings and imports reveal the following:

- Verbose status strings describe its intended actions: clipboard capture, process listing, screenshots, installed-program enumeration, system info collection, and HTTP/FTP upload.

- Anti-analysis checks include using Sleep to outlast sandbox scan timeouts; querying display-adapter registry entries against known VM names; comparing the host name and username against lists of common sandbox values (53 and 65 entries respectively); checking whether the C: drive is under roughly 59GB; scanning for VM-related processes and registry keys (including Hyper-V); and calling IsDebuggerPresent.

- Targets span many browsers (Chrome, Edge, Brave, Opera, Vivaldi, Yandex and others), around 75 wallet extensions and applications, messaging clients (Discord, Telegram, WhatsApp), VPN clients (OpenVPN, NordVPN, ProtonVPN), FTP (FileZilla) and game clients (Epic Games, Roblox). It also searches for files containing sensitive keywords (such as password, seed, wallet and bank) across common document types.

- Hardcoded indicators include public-IP lookup services (api.ipify.org, ip-api.com), GitHub-hosted resources and a suspicious exfiltration domain.

Dynamic analysis

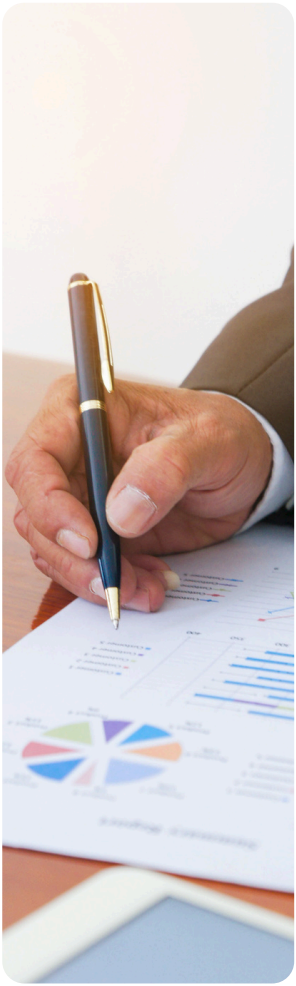

- On launch it runs its anti-analysis checks; if a VM, sandbox or debugger is detected it displays an error and exits.
- Otherwise it downloads a helper DLL (abe_decrypt.dll) from a GitHub-hosted URL into the user's Temp folder. The DLL is used to defeat Chromium app-bound encryption (ABE), which normally restricts decryption keys to the browser process.
- It launches Edge in a hidden, security-reduced configuration (headless, sandbox disabled, logging off, window moved off-screen and shrunk) and injects the DLL, then harvests Edge's Local State, cookies and login data. It repeats this for other Chromium browsers such as Chrome and Opera.
- Collected data is staged under C:\ProgramData\Evelyn\ in a structured layout (browsers, wallets, VPN, games, WhatsApp sessions, plus system info, processes, Wi-Fi passwords and a desktop screenshot).
- it retrieves a Python clipboard-hijacker script that watches the clipboard for cryptocurrency wallet addresses and swaps them for attacker-controlled addresses, and persists by copying itself into the Startup folder (WinCheck.pyw).
- It enumerates the host's public IP and country, packages the staged data into a ZIP named with host and environment details, and exfiltrates it over HTTPS to a remote server.

Key indicators

For blocking and hunting (defanged):

- Staging folder: C:\ProgramData\Evelyn\
- Temp artefacts: abe_decrypt.dll; service_task.pyw
- Startup persistence: WinCheck.pyw
- Exfiltration domain: ins0mnia[.]ru; User-Agent: "Evelyn/1.0"

MITRE ATT&CK

Tactic	Techniques
Execution	Command and Scripting Interpreter; Native API
Persistence	Boot or Logon Autostart Execution
Privilege Escalation	Process Injection
Defense Evasion	Debugger Evasion; Delay Execution; Process Injection; Virtualization/Sandbox Evasion
Credential Access	Credentials from Password Stores; Steal Web Session Cookie; Unsecured Credentials
Discovery	Browser Information; File and Directory; Process; Query Registry; System Information; System Owner/User; Virtual Machine; Virtualization/Sandbox
Collection	Archive Collected Data; Automated Collection; Clipboard Data; Data from Local System; Data Staged; Screen Capture
Exfiltration	Automated Exfiltration; Exfiltration Over Alternative Protocol

Detection and response

This activity presents several detection opportunities for a layered managed detection and response (MDR) service:

- The binary is detected on disk and at runtime.
- Unauthorised access to browser Local State, Web Data and Cookies is detected as credential and session theft.
- Web and DNS filtering blocks the outbound connection to the known-malicious exfiltration domain.

MALWARE DEEP DIVE: AUR0RA RANSOMWARE

Executive summary

Aur0ra is an emerging ransomware group hitting multiple sectors (business services, healthcare, consumer services, hospitality and tourism, and manufacturing), with mostly US victims and some in the Maldives and Australia.

The family is multi-platform (Windows, Linux and ESXi); the analysed sample was a Windows x64 build. It shows classic ransomware behaviour: file encryption, ransom notes, recovery inhibition and a Tor extortion portal, with the note claiming data theft (double extortion).

Static analysis

Embedded strings and imports reveal the following:

- Imports network-enumeration APIs (WNetOpenEnumW, WNetEnumResourceW, WNetCloseEnum), suggesting it can map accessible network resources.
- Imports privilege and access-control APIs (AdjustTokenPrivileges, OpenProcessToken, LookupPrivilegeValueW, SetEntriesInAclW, SetNamedSecurityInfoW).
- Imports Windows CryptoAPI functions (CryptAcquireContext, CryptGenRandom, CryptReleaseContext), consistent with encryption operations.
- Contains a built-in help menu with configurable options: target path, partial-encryption percentage, thread count, file-size limits and an ESXi-only mode.
- Contains commands for deleting volume shadow copies, resizing shadow storage and disabling the Windows System Restore scheduled task, all aimed at inhibiting recovery.

Dynamic analysis

The following recovery-inhibition commands were observed (these are well-known defensive indicators):

- cmd.exe running “vssadmin delete shadows /all /quiet” and “wmic shadowcopy delete” to remove volume shadow copies.
- “vssadmin resize shadowstorage /for=C: /on=C: /maxsize=401MB” to shrink shadow storage, which can purge existing restore data.
- “schtasks /Change /TN \Microsoft\Windows\SystemRestore\SR /Disable” to disable System Restore.
- Writes a ransom note named !!!README!!!DO_NOT_DELETE.txt across the machine; the note claims data was stolen and directs victims to a Tor portal with an access key.
- Encrypts file contents in place without renaming files or adding an extension, which can make the impact less obvious at a glance.
- Operates a Tor leak site listing victims to apply pressure.

MITRE ATT&CK

Tactic	Techniques
Execution	Native API; Command and Scripting Interpreter
Discovery	File and Directory Discovery; Process Discovery; Network Share Discovery
Impact	Data Encrypted for Impact; Inhibit System Recovery; Financial Theft

Detection and response

This activity presents several detection opportunities for a layered managed detection and response (MDR) service:

- The binary is detected on disk and at runtime, and corroborated by threat intelligence.
- Behavioural process monitoring flags the recovery-inhibition commands (shadow copy deletion and resize, and System Restore disablement).
- Decoy-file monitoring detects the in-place encryption of protected files by an unsigned process.

DARKNET WATCH: THE GENTLEMEN RAAS BREACH

This month produced an unusual case: the ransomware-as-a-service group “The Gentlemen” suffered a breach of its own infrastructure, exposing backend systems, internal chats and operational data. Attacks on organisations are routine; attacks that expose the attackers themselves are rare, and this fits a growing pattern of criminal-on-criminal compromise that highlights how fragile trust is within the underground economy.

Who they are

The Gentlemen emerged in mid-2025 under a RaaS model, with a core team building the malware and running infrastructure and negotiations while affiliates carry out intrusions for a share of payments. The group scaled quickly, claiming hundreds of victims by early 2026, helped by an aggressive 90/10 revenue split that attracted experienced operators.

The breach

In early May 2026 the group’s backend was compromised, reportedly via a hosting provider, exposing chat logs, affiliate rosters, negotiation records, tooling discussions and victim-management databases. The data was then offered for sale and reposted across multiple forums and file-sharing sites, making it impossible for the group to contain. The person or group responsible remains unidentified.

The leaked dataset effectively maps each intrusion from initial access to extortion: structured victim tracking, affiliate identifiers, negotiation transcripts and encryption task logs, while chat logs reveal live coordination around tooling, credential reuse and victim prioritisation.

Structure and tradecraft

Rather than a large network, the group is a tight team of around nine operators centred on a single administrator, yet it works with the discipline of much larger cartels.

Its approach relies on:

- Initial access through exposed edge devices (VPNs, firewalls) and stolen credentials.
- Rapid lateral movement, privilege escalation and disabling of security tooling.
- Use of legitimate administrative tools for stealth and persistence.
- Deployment across Windows, Linux and ESXi environments.

A notable trait is chain victimisation: data from one victim is reused to attack others, extending the value of a single breach across supply chains.

Assessment

The group projected resilience, but the leak shows a small, efficient operation that depends on coordination and proven techniques rather than novel capability. The same weaknesses it exploits in victims (poor segmentation, exposed credentials and mismanaged infrastructure) ultimately enabled its own compromise.

The episode is a valuable window into how a modern RaaS operation runs day to day, and it reinforces a wider trend of instability in the ransomware ecosystem.

The Gentlemen remains capable, but the incident shows these operations are not immune to the risks they impose on others.

SYSTEM FORCE IT RECOMMENDATIONS

Drawing the month together, we recommend the following priority actions:

- Patch the high-severity vulnerabilities above without delay, prioritising internet-facing edge devices (Palo Alto, Cisco SD-WAN, Fortinet) and domain controllers.
- Treat commodity malware and “low priority” alerts as potential precursors. The Sandworm and Gentlemen cases show attackers dwelling for weeks before acting.
- Harden the software supply chain: pin and verify dependencies, restrict who can install editor and IDE extensions, and protect CI/CD credentials.
- Reduce ransomware impact: enforce MFA on all remote access, segment networks, and keep offline, tested backups. Several families this month specifically target shadow copies and System Restore.
- Protect browser-stored credentials and crypto assets: deploy EDR that monitors access to browser Local State and cookies, and verify wallet addresses out of band before transacting.
- Keep Linux servers and container or Kubernetes hosts patched against Copy Fail (CVE-2026-31431).

If you would like System Force IT to review your exposure to any of the threats in this report, or to discuss managed detection and response, please contact your account manager.

APPENDIX

Risk ratings

Low	Medium	High	Critical

Traffic Light Protocol (TLP)

TLP is a standard set of labels used to indicate how sensitive information may be shared. The label on this report governs its distribution.

Label	When to use	How it may be shared
TLP:RED	Restricted to named recipients only, where wider sharing could harm privacy, reputation or operations.	Do not share beyond the specific people involved in the original exchange. Usually communicated in person or verbally.
TLP:AMBER	Limited sharing, where action requires support but wider exposure carries risk.	Share only within your own organisation, and with clients who need it to protect themselves. Sources may set tighter limits.
TLP:GREEN	Useful for awareness across a community or sector.	Share with peers and partners in your sector or community, but not via public channels.
TLP:CLEAR	Minimal or no foreseeable risk of misuse (formerly TLP:WHITE).	May be distributed without restriction, subject to standard copyright rules.

System Force I.T.
Secure IT Simplified

CONTACT

Quedgeley, Gloucester, United Kingdom

01452701355

www.systemforce.co.uk

sales@systemforce.co.uk

Security enquiries: [insert SOC / support contact]