

## CISO-Share Office Weekly Newsletter

Item 10 - Link 16

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

## Table of Contents

<b>CISO-Share Office Weekly Newsletter</b> .....	1
SC3 Daily Threat Summaries and Weekly Report .....	1
Chrome Extensions With 900,000 Downloads Caught Stealing AI Chats.....	2
Critical WordPress Modular DS Plugin Flaw Actively Exploited to Gain Admin Access .....	3
Cyber Insights 2026: External Attack Surface Management .....	5
New 'Reprompt' Attack Silently Siphons Microsoft Copilot Data .....	6
Cisco Patches Zero-Day RCE Exploited by China-Linked APT in Secure Email Gateways .....	7
Acumen Cyber Threat Intelligence Digest: Week 2 .....	8

---

## SC3 Daily Threat Summaries and Weekly Report



Please find SC3's daily threat summaries for this week for those who do not receive this information directly.

### **This week's reports:**

<https://www.cyberscotland.com/news/daily-threat-reports/>



SC3 - Monthly  
Ransomware Report



## Chrome Extensions With 900,000 Downloads Caught Stealing AI Chats

Two malicious Chrome extensions were observed exfiltrating browser data and users' conversations with ChatGPT and DeepSeek, OX Security reports.

### Main Article

Impersonating a legitimate extension from AITOPIA, the two extensions gathered over 900,000 downloads, potentially impacting as many users.

The applications, called 'Chat GPT for Chrome with GPT-5, Claude Sonnet & DeepSeek AI' and 'AI Sidebar with Deepseek, ChatGPT, Claude and more', are no longer available in the Chrome web store.

According to [OX Security](#), the extensions were abusing the AI-powered web development platform Lovable to host infrastructure components and anonymize their activity.

The legitimate AITOPIA extension they were impersonating allows users to chat with popular LLM models through a sidebar on top of visited websites.

The malicious applications copied the legitimate extension and added code that requested user consent to harvest "anonymous, non-identifiable analytics data" but instead stole the users' complete ChatGPT and DeepSeek conversations.

**Action Point:** "This data can be weaponized for corporate espionage, identity theft, targeted phishing campaigns, or sold on underground forums. Organizations whose employees installed these extensions may have unknowingly exposed intellectual property, customer data, and confidential business information," OX Security notes.

Users are advised to remove the malicious extensions from their Chrome browser as soon as possible.

---



## Critical WordPress Modular DS Plugin Flaw Actively Exploited to Gain Admin Access

A critical vulnerability in the **Modular DS WordPress plugin**—used by over **40,000 sites**—is now being **actively exploited**. The flaw, **CVE-2026-23550**, carries a **CVSS score of 10.0**, meaning it's as severe as it gets. The vulnerability allows **unauthenticated privilege escalation**, enabling an attacker to go from nothing to **full admin access** without logging in. It affects **all versions up to and including 2.5.1** and is patched in **2.5.2**.

### [Main Article](#)

The issue stems from several design weaknesses that interact badly when combined. The plugin provides an `/api/modular-connector/` route system that's meant to sit behind authentication. But when the plugin is in "direct request" mode, attackers can bypass the auth layer simply by adding two parameters to their request:

- `origin=mo`
- `type=<anything>`

This tricks the plugin into processing the request as a trusted "Modular" call. Worse, once the site is already connected to the Modular service (meaning tokens exist), **there's no cryptographic validation** to ensure requests actually come from Modular. That means any unauthenticated attacker can call sensitive routes like `/login/`, `/server-information/`, `/manager/`, and `/backup/`.

One of those routes—`/login/{modular_request}`—lets attackers auto-login as an admin. Once in, they can create new admin users, install backdoors, inject malware, deface the site, or redirect traffic to scams.

Patchstack observed **live exploitation** beginning **13 January 2026 at around 02:00 UTC**, where attackers used HTTP GET requests to the login endpoint and then attempted to create admin accounts. Two attacker IPs were highlighted:

- **45.11.89[.]19**
- **185.196.0[.]11**

Patchstack and the plugin maintainers note that the vulnerability wasn't just a simple bug—it resulted from overly permissive routing, poor authentication checks, and an admin login flow that automatically falls back to an administrator account when ambiguous.

Modular DS recommends immediate patching and a full compromise assessment. For sites already hit, attackers may have created hidden admin users, modified files, or planted malicious plugins.

**Action Points:**

1. Patch Immediately – Upgrade Modular DS plugin to v2.5.2 or later.
  2. Check for Compromise – Look for unexpected admin accounts.
    - Review logs for requests to /api/modular-connector/login/.
    - Search for suspicious plugins, altered files, or dropped payloads.
  3. Regenerate Secrets – Regenerate WordPress salts (invalidate all sessions).
    - Regenerate OAuth credentials.
  4. Harden WordPress – Disable unused plugins.
    - Enforce MFA for all admin accounts.
    - Limit access to admin endpoints.
  5. Monitor Traffic – Watch for scans or repeat hits from known attacker IPs.
-



## Cyber Insights 2026: External Attack Surface Management

Shadows are dark and dangerous places where bad guys attack anything or anyone they find. In 2026, AI will increase the number and size of shadows, together with the entire external attack surface.

### Main Article

External Attack Surface Management (EASM) is the process of finding and managing every asset an organization exposes to the internet. Those assets may be known (and therefore documented and may be secured) or unknown (and therefore invisible and almost certainly insecure). While EASM covers both categories, we are primarily concerned with the invisible assets.

“This includes domains, servers, APIs, and cloud assets that may not be tracked internally,” says Chris Boehm, field CTO at Zero Networks. “It matters because most companies do not have a complete inventory of what is visible from the outside, and attackers often find these gaps before defenders do.”


### Action Point:

“External attack surface management will remain a critical, but increasingly complex, issue in cyber security in the year ahead, largely because organizations have lost control of their environments,” warns Simon Phillips, CTO of Engineering at CybaVerse.

Control has been lost because business pressure and the need for agility to stay ahead of the competition results in new technology being adopted faster than security can apply governance. This includes the rapid adoption of SaaS solutions, the personal use of shadow IT, and the unsanctioned rise of shadow AI by individuals and developers downloading undisclosed copies of MCP.

AI is the double-edged sword in the picture. It will assist companies in finding their external attack surface, but it will also assist bad actors in locating and attacking the weak points.

The likelihood for 2026 is that the battle between attackers and defenders will increase in size, complexity and speed – with no sign of any decrease.



## New 'Reprompt' Attack Silently Siphons Microsoft Copilot Data

Dubbed [Reprompt](#), the attack bypassed the LLMs data leak protections and allowed for persistent session exfiltration even after the Copilot was closed, Varonis says.

[Main Article](#)

The attack leverages a Parameter 2 Prompt (P2P) injection, a double-request technique, and a chain-request technique to enable continuous, undetectable data exfiltration.

The Reprompt Copilot attack starts with the exploitation of the 'q' parameter, which is used on AI platforms to deliver a user's query or prompt via a URL. All it takes is for the user to click on the link.

"By including a specific question or instruction in the q parameter, developers and users can automatically populate the input field when the page loads, causing the AI system to execute the prompt immediately," Varonis explains.

A threat actor, the cybersecurity firm notes, could abuse the feature to make Copilot execute unwanted actions. The attack resulted in one-click compromise and, because it leveraged the active user session, it persisted after the chat was closed.

### Action Point:

Applying the latest security patches should keep you safe.

Microsoft has resolved the underlying issue. The attack does not now affect enterprise customers using Microsoft 365 Copilot.

"We appreciate Varonis Threat Labs for responsibly reporting this issue. We have rolled out protections that address the scenario described and are implementing additional measures to strengthen safeguards against similar techniques as part of our defense-in-depth approach," a Microsoft spokesperson told *SecurityWeek*.



## Cisco Patches Zero-Day RCE Exploited by China-Linked APT in Secure Email Gateways

Cisco has released patches for a **maximum-severity (CVSS 10.0)** remote command-execution vulnerability—**CVE-2025-20393**—affecting Cisco AsyncOS, the software powering Cisco Secure Email Gateway and Secure Email & Web Manager appliances. The worrying part? This wasn't just a theoretical risk. Cisco confirmed it was exploited **as a zero-day** by a **China-linked APT group known as UAT-9686**, with activity traced back to **late November 2025**.

### [Main Article](#)

The vulnerability exists in the **Spam Quarantine** feature, which fails to properly validate HTTP requests. If exploited, it gives an attacker **full root-level access** to the underlying OS—effectively total control of the appliance. Notably, exploitation requires three conditions: the appliance must be running a vulnerable release, must have Spam Quarantine enabled, and must expose that feature to the internet.

Once inside, UAT-9686 deployed a stack of post-exploitation tooling. This included tunnelers like **ReverseSSH (AquaTunnel)** and **Chisel**, a log-scrubber called **AquaPurge**, and a lightweight Python backdoor named **AquaShell**. AquaShell can receive encoded commands and execute them quietly, providing persistence and stealthy control.

Cisco has now issued fixed versions across AsyncOS branches 14.x through 16.x for both Secure Email Gateway and Email & Web Manager. They've also rolled out remediation to **remove persistence mechanisms** implanted by the attackers.

Beyond patching, Cisco is urging organisations to harden their appliances. Recommendations include firewalling the systems away from the open internet, monitoring logs for any unusual traffic, disabling unnecessary services, switching off HTTP access to the admin portal, enforcing strong authentication (SAML or LDAP), and changing any default admin credentials.

This incident is yet another reminder that email security appliances—often seen as protective layers—are high-value targets when they become initial access points. Their privileged network position, visibility into email flows, and potential lateral-movement opportunities make them prime candidates for targeted exploitation by state-aligned groups.

With a CVSS 10.0 score, confirmed zero-day exploitation by an advanced threat actor, and now broad disclosure, this vulnerability requires urgent action across all environments running impacted versions.

### Action Points:

**Patch immediately** to fixed AsyncOS versions: – Secure Email Gateway: **15.0.5-016, 15.5.4-012, 16.0.4-016**

– Email & Web Manager: **15.0.2-007, 15.5.4-007, 16.0.4-010**

**Isolate and harden appliances** – Place behind a firewall; block direct internet access to Spam Quarantine.

**Disable unnecessary services** – Disable HTTP admin access; enforce HTTPS only.

**Strengthen authentication** – Require SAML, LDAP, or comparable strong authentication.

**Monitor for compromise indicators** – Look for ReverseSSH/AquaTunnel, Chisel, AquaPurge, or unusual Python processes.

**Reset admin credentials**

– Replace any default or weak passwords.





## Acumen Cyber Threat Intelligence Digest: Week 2

The Week 2 digest highlights key cyber security developments across vulnerabilities, active threats, global news, and threat trends.

### Main article

**Significant vulnerabilities** were disclosed and patched: Angular addressed a high-severity XSS flaw (CVE-2026-22610) affecting core framework packages, while ServiceNow fixed a critical privilege escalation issue in its AI and Virtual Agent API (CVE-2025-12420). Appsmith patched a critical origin validation flaw (CVE-2026-22794) that could allow token theft; at the time of publishing none are known to be widely exploited.

Under **potential threats**, multiple active campaigns were observed. A **Magecart web-skimming operation** continues targeting e-commerce sites to steal payment and personal data via obfuscated JavaScript. Threat actors carried out **LinkedIn phishing** by abusing the platform's comment reply function to hijack credentials. Another phishing campaign impersonated employee performance reports to deliver **Guloader**, which installs the **Remcos RAT**, enabling remote control, keylogging and data capture on Windows systems.

In **general news**, the U.S. urged stronger international action against **North Korean IT worker scams and crypto thefts**, noting the operations' scale and use of stolen credentials to bypass sanctions. Widespread internet blackouts in **Iran** persisted amid protests, raising concerns about information access. New techniques such as **browser-in-the-browser phishing** were reported, showing increasingly sophisticated credential theft methods.

The report also includes threat actor trends and global threat intelligence signals, and concludes with **remediation recommendations** to update affected software and apply patches where available.

### **Action Point:** Remediation Actions

Following the information provided above, we recommend that the technologies mentioned be fully patched and updated. We also want to highlight and recommend applying the following patches where applicable:

- **CVE-2026-22610 (Angular XSS Vulnerability)** – This vulnerability can be addressed by upgrading Angular to patched versions 19.2.18, 20.3.16, 21.0.7 or 21.1.0-rc.0.
- **CVE-2025-12420 (ServiceNow AI Agents/Virtual Agent API)** – This vulnerability can be addressed by updating Now Assist AI Agents in versions 5.1.18, 5.2.19, alongside fixes in the Virtual Agent API from versions 3.15.2 and 4.0.4.
- **CVE-2026-22794 (Appsmith)** – This vulnerability can be addressed by upgrading Appsmith to version 1.93 or later, which introduces proper validation to prevent abuse of the Origin header.

---

**All the best from all of us at HEFESTIS and look out for a new ThreatScape update next week.**

---



## CISO-Share Office Weekly Newsletter

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are also a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

---

### SC3 Daily Threat Summaries and Weekly Report

<https://www.cyberscotland.com/news/daily-threat-reports/>



A daily breakdown of various cyber threats from the Scottish Cyber Coordination Centre (SC3)

#### Action Point:

All SC3 threat intelligence in one place.



### Web Check and Mail Check decommissioning update and feedback opportunity

As communicated in November 2025, the NCSC has made the decision to stop providing Mail Check and Web Check.

- Mail Check and Web Check will no longer be available on MyNCSC from 31st March 2026.
- Service users will need to source alternative commercial products by this date.
- The NCSC will continue to provide you with some external attack surface alerts via Early Warning and DNS Check, when we receive information relating to the security

of your assets.

- NCSC continue to work closely with commercial EASM providers, and learning from ACD2.0 research into attack surface management is refocusing efforts to develop new services and products where the NCSC can add unique value.

They are collecting feedback to understand how it may have affected users, both positively and negatively.

<https://bz5q106a.optimalworkshop.com/questions/ncsc-decomm-survey-2026>

Please note this is hosted by Optimal Workshop. The survey should take approximately 10-15 minutes to complete.

NCSC will provide updates and further guidance, but if you have any questions, please contact [servicetransition@digital.ncsc.gov.uk](mailto:servicetransition@digital.ncsc.gov.uk).

## Action Point:

Identifying a new External Attack Surface Management Product

Products that provide similar functionality to Mail Check or Web Check are often known as External Attack Surface Management (EASM) products. The NCSC has provided a buyer's guide to help you pick an External Attack Surface Management (EASM) product that best suits your organisation's specific needs.

There are many EASM products available commercially at a range of price points, including many that offer free tiers for small attack surfaces. Free offerings may be limited by the number of assets (domains or IP addresses) they can monitor, or the frequency of the checks e.g. weekly instead of daily.

You can also utilise the Check Your Cyber Security service. It checks your email domain for two important areas of cyber security relating to your emails; Email anti-spoofing and Email privacy. It can also check your IP address and website, as well as your web browser.



## Acumen Cyber Threat Intelligence Digest: Week 3 - 2026

<https://acumencyber.com/cyber-threat-intelligence-digest-january-2026-week-3>

This threat-intel digest highlights a wave of new vulnerabilities, active exploitation, and emerging threat activity across open-source platforms, Windows, Fortinet products, browser extensions, developer tooling, and phishing campaigns.

The first major issue is **CVE-2025-8110**, a **high-severity path-traversal flaw** in **Gogs**, an open-source Git server. Despite a pull request being opened in early January 2026, no formal patch has been released. A **public PoC** appeared in December and **CISA added it to KEV**, confirming exploitation in the wild. The vulnerability allows low-privileged users to write arbitrary files and potentially achieve **code execution**, making self-hosted developer environments particularly exposed.

Next is **CVE-2026-20805**, a Windows **Desktop Window Manager** information-disclosure issue. Though only medium-severity, it reached KEV status on 13 January with a PoC surfacing shortly after. Attackers need local code execution first, but exploitation could leak sensitive memory content. Microsoft has already patched the issue in January's Patch Tuesday.

Fortinet continues to feature in threat reports, this time with **CVE-2025-64155**, a **critical unauthenticated OS command-injection** in FortiSIEM. The vulnerability is fixed across several release branches, and while no exploitation has been reported yet, the nature of the flaw makes it extremely high-risk.

From a threat-actor perspective, multiple campaigns are active. One is the **CrashFix / NexShield campaign**, attributed to "KongTuke," which uses a fake Chrome extension to crash browsers and trick users into running a PowerShell command. This ultimately deploys **ModeloRAT** to domain-joined machines, using clever social engineering and multi-stage infection.

Another campaign targets software developers specifically — attackers are Trojanising **Visual Studio Code extensions** to drop **Evelyn Stealer**, a credential-harvesting malware. The infection leverages DLL sideloading, process hollowing, and browser injections to steal passwords, crypto wallets, and system data.

A third campaign is shipping-themed phishing delivering **fileless Remcos RAT** using the legacy **CVE-2017-11882** Equation Editor vulnerability. This underscores how attackers continue to rely on old bugs because they're still highly effective.

General news items include warnings about **Russian hacktivist DDoS attacks**, the EU's proposed **cybersecurity overhaul targeting high-risk suppliers**, and the identification of the leader of the **BlackBasta ransomware group** by German-Ukrainian investigators.

## Action Points

### Patch & Upgrade

- Apply patches for **FortiSIEM** (7.1.9+, 7.2.7+, 7.3.5+, 7.4.1+).
- Apply Microsoft's **January 2026 updates** for CVE-2026-20805.
- Track the pending **Gogs** patch and consider isolating or restricting its use.

### Threat Hunting

- Look for suspicious Chrome extensions (**NexShield**) and PowerShell execution.
- Monitor for malicious VS Code extensions and DLL sideloading events.
- Hunt for Remcos indicators, especially activity exploiting **CVE-2017-11882**.

## Hardening

- Enforce application-allowlisting for browser and developer extensions.
- Block outbound traffic to known attacker infrastructure and IPs.
- Ensure EDR rules for process hollowing, clipboard manipulation, and abnormal browser behaviour are active.

## Organisational Defence

- Prepare DDoS-resilience plans for public-facing services.
- Review supply-chain risk posture in light of the EU policy shift.



### Automated FortiGate Attacks Exploit FortiCloud SSO to Alter Firewall Configurations

<https://thehackernews.com/2026/01/automated-fortigate-attacks-exploit.html>

Arctic Wolf has raised an alert about a **new wave of automated attacks targeting Fortinet FortiGate appliances**, beginning on **15 January 2026**. These attacks focus on abusing **FortiCloud SSO** to perform **unauthorised firewall configuration changes**, and they strongly resemble a campaign from December 2025 that exploited **CVE-2025-59718** and **CVE-2025-59719** — both unauthenticated SAML bypass vulnerabilities.

These earlier flaws allow attackers to bypass SSO authentication entirely using **crafted SAML messages**, but the concerning development is that **customers are now reporting malicious SSO logins even on fully patched systems**. One Reddit user says Fortinet developers confirmed the issue “persists or is not fixed in version 7.4.10,” though Fortinet has not yet responded publicly.

The current campaign demonstrates a high level of automation. Attackers use a malicious account — **cloud-init@mail.io** — to perform SSO logins from four known IP addresses:

- **104.28.244.115**
- **104.28.212.114**
- **217.119.139.50**
- **37.1.209.19**

Within seconds of gaining access, the attackers:

1. **Exfiltrate firewall configuration files** via the GUI
2. **Create multiple persistent local accounts**, including:
  - secadmin
  - itadmin
  - support
  - backup

- remoteadmin
- audit

### 3. **Enable VPN access for these accounts**

The speed and consistency of these steps strongly indicate automated tooling rather than hands-on activity.

The vulnerabilities affect systems with FortiCloud SSO enabled across several Fortinet products:

- **FortiOS**
- **FortiWeb**
- **FortiProxy**
- **FortiSwitchManager**

This is particularly worrying because FortiCloud SSO is widely used, and successful exploitation gives attackers administrative access to critical network infrastructure. Exfiltrating firewall config files also gives attackers insight into an organisation's network topology, VPN setup, and segmentation — a goldmine for follow-on attacks.

Until Fortinet releases an updated advisory or patch, Arctic Wolf recommends **immediately disabling the “admin-forticloud-sso-login” feature**, which prevents remote SSO logins from being used in this attack chain.

With multiple reports of compromise coming from organisations already running patched firmware, the issue may reflect an incomplete fix or a new variant of the earlier exploit chain.

### **Action Points**

1. **Disable FortiCloud SSO immediately**
  - Turn off admin-forticloud-sso-login on all FortiGate appliances.
2. **Audit for compromise**
  - Look for logins from the known malicious IPs.
  - Check for the presence of suspicious accounts (secadmin, backup, audit, etc.).
  - Review firewall configuration export logs.
3. **Rotate credentials**
  - Reset local admin passwords and API keys.
  - Reissue VPN credentials.
4. **Inspect firewall configurations**
  - Check for unauthorised VPN rules, policy edits, or changed admin settings.
5. **Increase monitoring**
  - Alert on SSO logins, new account creation, and config exports.
  - Enable verbose logging where possible.
6. **Await official Fortinet guidance**
  - Monitor advisories for new patches or confirmation of unresolved vulnerabilities.



## Resurgence of a multi-stage AiTM phishing and BEC campaign abusing SharePoint

<https://www.microsoft.com/en-us/security/blog/2026/01/21/multistage-aitm-phishing-bec-campaign-abusing-sharepoint/>

Microsoft has uncovered a **resurgent, multi-stage adversary-in-the-middle (AiTM)** phishing and **business email compromise (BEC)** campaign targeting organisations in the **energy sector**. The campaign stands out because it abuses **SharePoint file-sharing links** — trusted, familiar infrastructure — to deliver phishing pages that appear legitimate and evade many traditional detection controls.

The operation starts with a **compromised trusted vendor account**, which attackers use to send SharePoint-style document-sharing emails. Because the link points to a real SharePoint domain requiring authentication, victims are far more likely to trust and click it. After clicking, users are redirected to an attacker-controlled credential prompt, enabling the theft of authentication cookies and session tokens — the core technique behind AiTM.

Once a victim account is compromised, attackers create **malicious inbox rules** that delete incoming emails and mark them as read, effectively blinding the user and allowing the attacker to operate silently. They then pivot to a **large-scale phishing campaign** using the compromised user's mailbox, sending more than **600 emails** across internal contacts, external partners, and distribution lists. These secondary victims who click the link are then targeted with another AiTM flow, multiplying the number of compromised accounts.

The attackers also monitor mailbox activity closely. They intercept and delete undelivered messages, out-of-office responses, and sceptical replies from recipients. They even respond to users questioning the authenticity of phishing emails — a classic BEC technique aimed at maintaining credibility and prolonging persistence.

Microsoft stresses an important point: **password resets are NOT enough** in AiTM attacks. Because attackers steal session cookies, they can maintain access even after credential changes. Attackers may also tamper with MFA settings, for example by adding their own one-time-password methods or registering their own phone numbers.

Defender XDR detected the attack through signals such as suspicious inbox rule creation, unusual sign-in IPs, malicious URL clicks, and cookie replay attempts. Microsoft emphasises the need for layered identity security including MFA, conditional access, continuous access evaluation, session revocation, and mailbox rule hygiene. Evidence of attacker infrastructure was also shared — notably IPs **178.130.46.8** and **193.36.221.10**.

This campaign demonstrates how AiTM operations can snowball across interconnected organisations: compromise one user, abuse their trusted identity, and rapidly spread across entire ecosystems.



## Action Points

1. **Revoke session cookies**  
Password resets alone won't help. Force sign-out across all devices.
2. **Remove malicious inbox rules**  
Look for rules that delete or auto-read incoming mail.
3. **Reset MFA configurations**  
Remove attacker-added OTP methods or phone numbers.
4. **Enable Conditional Access**  
Particularly risk-based policies, compliant device requirements, and trusted IP restrictions.
5. **Enable continuous access evaluation**  
Reduce reliance on long-lived session tokens.
6. **Use advanced anti-phishing protections**  
Defender for Office 365, Defender for Endpoint, and Microsoft Edge malicious site blocking.
7. **Hunt for indicators**
  - IPs: **178.130.46.8, 193.36.221.10**
  - Suspicious login locations, anonymizer IPs
  - Unexpected inbox rule creation
  - Cookie-replay sign-in attempts
8. **Educate staff**  
Highlight risks of SharePoint/OneDrive links and “trusted vendor” impersonation.



### Cisco Fixes Actively Exploited Zero-Day CVE-2026-20045 in Unified CM and Webex

<https://thehackernews.com/2026/01/cisco-fixes-actively-exploited-zero-day.html>

Cisco has released urgent patches for **CVE-2026-20045**, a **critical (CVSS 8.2)** vulnerability affecting multiple versions of **Cisco Unified Communications Manager (Unified CM)**, **Unity Connection**, and **Webex Calling Dedicated Instance**. This weakness has already been **actively exploited in the wild as a zero-day**, prompting Cisco and CISA to sound the alarm.



The issue stems from **improper validation of user-supplied input** in HTTP requests processed by the management interface. In practical terms, an attacker can send a sequence of crafted HTTP requests to a vulnerable device. If successful, they gain **user-level access to the operating system**, then escalate privileges all the way to **root**—effectively total system control.

Cisco notes that the flaw impacts a wide range of communication infrastructure:

- Unified CM
- Unified CM Session Management Edition
- Unified CM IM & Presence
- Unity Connection
- Webex Calling Dedicated Instance

Because these systems serve as core collaboration infrastructure in many organisations, exploitation offers attackers privileged access to internal communications, authentication integrations, voicemail systems, and other sensitive services.

Cisco has released fixed software versions across releases 12.5, 14, and 15. In some cases, customers must migrate to a later service update; in others, a downloadable patch (.cop) file is provided. Notably, **there are NO workarounds**. If you're running a vulnerable version, the only mitigation is to patch.

Cisco confirms it has observed **attempted exploitation**, and although it hasn't published technical details, it stresses that the vulnerability is being used in active attack campaigns. CISA added CVE-2026-20045 to the **Known Exploited Vulnerabilities (KEV)** catalogue and now requires all U.S. federal civilian agencies to patch by **11 February 2026**.

This comes just a week after Cisco patched another critical zero-day — **CVE-2025-20393** — impacting AsyncOS for Secure Email Gateway products. Together, the two incidents underline that Cisco appliances remain prime targets for APTs due to their privileged position inside enterprise networks.

For organisations using Cisco Unified CM or related telephony platforms, this is a high-priority patch cycle with material risk if left unaddressed.

## Action Points

### 1. Patch immediately

– Unified CM / SME / IM&P / Webex Dedicated Instance:

- Upgrade 12.5 to a fixed release
- 14 → apply **14SU5** or patch file
- 15 → **15SU4** (Mar 2026) or apply available patches

– Unity Connection:

- Upgrade 12.5
- Patch 14 / 15 with provided .cop files

### 2. Confirm exposure

– Identify systems with the web-based management interface accessible internally or

externally.

- Block external exposure immediately.

3. **Hunt for compromise**

- Look for suspicious HTTP requests to management interfaces.
- Monitor OS-level logs for privilege-escalation attempts.

4. **Strengthen access controls**

- Enforce MFA for administrative access.
- Restrict management interfaces to an isolated admin subnet.

5. **Review related Cisco advisories**

- Validate you have also patched **CVE-2025-20393** on AsyncOS appliances.

---

All the best from all of us at **HEFESTIS** and look out for our next **Threatscape** update next week.

*HEFESTIS Limited, Registered Office: Unit 27, Stirling Business Centre, Wellgreen, Stirling FK8 2DZ*

*Incorporated in Scotland SC603511*

## CISO-Share Office Weekly Newsletter

**A very warm welcome to all of you from all of us at the CISO-Share Office.**

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

## Table of Contents

<b>CISO-Share Office Weekly Newsletter .....</b>	<b>1</b>
SC3 Daily Threat Summaries and Weekly Report .....	1
New 'Reprompt' Attack Silently Siphons Microsoft Copilot Data .....	2
Living Off the Web: How Fake Captcha Turned Trust Into a Malware Delivery Channel .....	3
Cyber Insights 2026: Quantum Computing and the Potential Synergy With Advanced AI ....	4
Over 100 Organizations Targeted in ShinyHunters Phishing Campaign .....	5
Acumen Cyber Threat Intelligence Digest: Week 4 .....	6

## SC3 Daily Threat Summaries and Weekly Report



Please find SC3's daily threat summaries for this week for those who do not receive this information directly.

**This week's reports:**

<https://www.cyberscotland.com/news/daily-threat-reports/>

---



## New 'Reprompt' Attack Silently Siphons Microsoft Copilot Data

Dubbed [Reprompt](#), the attack bypassed the LLMs data leak protections and allowed for persistent session exfiltration even after the Copilot was closed, Varonis says.

### [Main Article](#)

The attack leverages a Parameter 2 Prompt (P2P) injection, a double-request technique, and a chain-request technique to enable continuous, undetectable data exfiltration.

The Reprompt Copilot attack starts with the exploitation of the 'q' parameter, which is used on AI platforms to deliver a user's query or prompt via a URL. All it takes is for the user to click on the link.

"By including a specific question or instruction in the q parameter, developers and users can automatically populate the input field when the page loads, causing the AI system to execute the prompt immediately," Varonis explains.

A threat actor, the cybersecurity firm notes, could abuse the feature to make Copilot execute unwanted actions. The attack resulted in one-click compromise and, because it leveraged the active user session, it persisted after the chat was closed.

**Action Point:** "Client-side monitoring tools won't catch these malicious prompts, because the real data leaks happen dynamically during back-and-forth communication — not from anything obvious in the prompt the user submits," Varonis says.

Microsoft has resolved the underlying issue. The attack does not affect enterprise customers using Microsoft 365 Copilot, Varonis notes.

"We appreciate Varonis Threat Labs for responsibly reporting this issue. We have rolled out protections that address the scenario described and are implementing additional measures to strengthen safeguards against similar techniques as part of our defense-in-depth approach," a Microsoft spokesperson told *SecurityWeek*.

---



## Living Off the Web: How Fake Captcha Turned Trust Into a Malware Delivery Channel

Fake Captcha pages look harmless — routine browser checks asking users to “verify” they’re human — but behind that familiar friction, attackers have quietly transformed trusted web workflows into a scalable malware delivery interface.

### Main Article

Recent [analysis](#) by Censys researchers shows that what once appeared to be a single, coordinated campaign is better understood as a fragmented ecosystem that lives off the web itself, inheriting trust rather than stealing it.

“Living Off the Web does not replace traditional lures or brand impersonation. It compounds them,” said Andrew Northern, principal security researcher at Censys in an email to eSecurityPlanet.

He added, “Threat actors still use trusted names and convincing pretexts, but increasingly they converge on abusing the mechanics of everyday internet use itself. CAPTCHAs, browser notifications, update prompts, and verification flows have become shared delivery surfaces because users are conditioned to interact with them to get things done.”

Andrew also explained, “This convergence spans diverse actors, toolchains, and objectives, signaling a broader shift that has been unfolding for years, from fake software updates to today’s web-native delivery interfaces.”

### **Action Point: Mitigating Fake Captcha and Trust-Based Attacks**

As Fake Captcha attacks continue to evolve, organizations need defenses that extend beyond traditional payload and malware detection.

These threats abuse trusted web workflows and user interactions, allowing malicious activity to unfold without obvious execution artifacts.

Effective mitigation requires shifting focus to browser behaviour, execution controls, and the handoff between user interaction and downstream activity.

- [Monitor](#) for security-themed verification pages appearing outside expected contexts and flag repeated “browser check” or “human verification” flows tied to unrelated infrastructure.
- Restrict browser notification permissions by default and closely monitor notification options that immediately follow verification or security prompts.
- Correlate browser interactions (e.g., clipboard access, permission grants, service worker registration) with downstream [endpoint](#) and network activity rather than relying solely on payload artifacts.
- Harden execution controls by limiting scripting engines, restricting MSI installation, and enforcing application allowlisting where possible.
- Reduce exposure by enforcing [least privilege](#), removing local install rights, and applying stronger controls to high-risk user groups.
- Strengthen detection and awareness by [training](#) users to recognize fake verification workflows and by logging and alerting on abnormal browser-initiated behavior.

Collectively, these steps help reduce exposure and detect Fake Captcha-driven attacks earlier in the kill chain.



## Cyber Insights 2026: Quantum Computing and the Potential Synergy With Advanced AI

Quantum computers are coming, with a potential computing power almost beyond comprehension. That's a given. The known threat is to current public key encryption methods, such as RSA and ECC, which will both be crackable through Shor's algorithm in short timeframes. It is believed that nation states and advanced criminal gangs are engaged in a widespread harvest now, decrypt later (HNDL) campaign – steal and store data and secrets today, even if they are encrypted, because they can be decrypted later with quantum computers.

### Main Article

But the timing is unquantified. Quantum computers exist today but are too 'small' to be a threat. Most projections do not expect a powerful quantum computer to be available within the next five years. But the fly in the ointment is the emergence of artificial intelligence, which could be used to speed the development of quantum (for example, by developing more efficient error correction modes), followed by automating the use of quantum power when it arrives.

In reality, nobody other than the major power intelligence agencies knows the current state of quantum development within other governments. We believe, and hope, that no adversarial nation is more advanced than ourselves.

The main thrust of this discussion will consider the potential of marrying advanced AI with powerful quantum computers. We'll probably see little evidence in 2026, but it may not be long beyond that. However, one thing is clear – we need to start considering the adversarial threats as well as the domestic potential that will come from the marriage of quantum and AI.

There are two areas to consider: the known threat to current encryption, where the use of AI might shorten the timescale to powerful quantum; and the unknown threat of powerful quantum automated by more advanced AI in the future.


**Action Point:** Quantum computers are not guaranteed but are likely. AGI is not guaranteed but is likely. The combination of the two, if and when they arrive, is not guaranteed – but is almost inevitable.

If or when that happens, the synergy of the two is pure conjecture. While the societal benefit could be enormous, the societal risk could be equally devastating. The power and speed of cyberattacks will be beyond human comprehension – and in this case, the first cut will literally be the deepest.

Will we be able to defend against future attacks? Probably not if cyber attackers or adversarial nations get there before we can use similar capabilities to defend.

Timing predictions for this new world order vary wildly, from 2026, 2027 and 2028 through 2030 to never. The effect is similarly debated, from quantum AGI versus quantum AGI largely cancelling each other (with an asymmetrical advantage to the adversary and a massive advantage to the first mover).

Meanwhile, while we wait, we should hope for the best and prepare for the worst.



## Over 100 Organizations Targeted in ShinyHunters Phishing Campaign

Over the past 30 days, Silent Push has [identified](#) domains suggesting that the threat actors have been preparing or conducting attacks against at least 100 organizations in sectors such as software and technology, financial, biotech and pharma, financial services, real estate, energy and utilities, healthcare, logistics and transportation, manufacturing, retail, and insurance.

### [Main Article](#)

Silent Push has named major companies such as Atlassian, Adyen, Canva, Epic Games, HubSpot, Moderna, ZoomInfo, GameStop, WeWork, Halliburton, Sonos, and Telstra.

The hackers have set up fake domains targeting these companies, but it's unclear whether any attacks were conducted or whether their attempts to gain access to systems were successful.

In the campaign, the cybercriminals used voice phishing (vishing) to target single sign-on (SSO) accounts associated with Okta and other identity platforms.

In attacks observed by Okta and others, threat actors used specialized phishing kits that enable them to intercept credentials and trick victims into helping them bypass multi-factor authentication.

**Action Point:** "The most critical of these features are client-side scripts that allow threat actors to control the authentication flow in the browser of a targeted user in real-time while they deliver verbal instructions or respond to verbal feedback from the targeted user," [Okta explained](#).

It added, "It's this real-time session orchestration that delivers the plausibility required to convince the threat actor's target to approve push notifications, submit one time passcodes (OTP) or take other actions the threat actor needs to bypass MFA controls."

"While this is not the result of a security vulnerability in vendors' products or infrastructure, we strongly recommend moving toward phishing-resistant MFA, such as FIDO2 security keys or passkeys where possible, as these protections are resistant to social engineering attacks in ways that push-based or SMS authentication are not. Administrators should also implement strict app authorization policies and monitor logs for anomalous API activity or unauthorized device enrolments," Carmakal added.





## Acumen Cyber Threat Intelligence Digest: Week 4

The **Cyber Threat Intelligence Digest: January 2026 – Week 4** from Acumen Cyber outlines significant vulnerabilities, active threat campaigns, and notable cyber-security developments observed during the week.

### Main Article

Several critical vulnerabilities are highlighted. A severe arbitrary file upload flaw in the RealHomes WordPress CRM plugin allows attackers to upload and execute malicious files and has been patched in version 1.0.1. An improper access control vulnerability affecting Oracle HTTP Server and the WebLogic Proxy Plug-in could allow unauthenticated attackers broad system access, with a proof-of-concept already publicly available. In addition, a WinRAR path traversal vulnerability has been actively exploited by both state-sponsored and financially motivated actors to deploy persistent malware; this issue is resolved in WinRAR version 7.13.

The digest also details multiple ongoing threat campaigns. One phishing operation abuses legitimate GoTo Resolve and LogMeIn remote management tools, beginning with Greenvelope-themed lures and escalating to full system compromise using trusted remote access software. The “MaliciousCorgi” campaign leverages malicious Visual Studio Code AI extensions to secretly exfiltrate source code and collect user information. Another campaign uses adversary-in-the-middle phishing techniques with spoofed SharePoint notifications to target energy sector organisations, enabling credential theft and business email compromise.

Broader cyber-security news includes a legal settlement by Google related to voice recordings, WhatsApp introducing a new anti-spyware lockdown feature, and UK government plans to centralise the policing of online crime and fraud.

### **Action Point:** Remediation Actions

Following the information provided above, we recommend that the technologies mentioned be fully patched and updated. We also want to highlight and recommend applying the following patches where applicable:

- **CVE-2025-67968 (Real Homes CRM Plugin)** – This vulnerability can be addressed by applying the 1.0.1 patch which introduced the necessary checks to prevent unauthorised file uploads.
- **CVE-2026-21962 (Oracle)** – This vulnerability can be addressed by updating Oracle Fusion Middleware to the most recent version.
- **CVE-2025-8088 (WinRAR)** – This vulnerability can be addressed by patching WinRAR to at least version 7.13.

**All the best from all of us at HEFESTIS and look out for a new ThreatScape update next week.**

CISO Share

## Threatscape

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

### SC3 Daily Threat Summaries and Weekly Report

Action Point: Please find SC3's daily threat summaries for this week for those who do not receive this information directly.

Steve McIntosh

2/6/2026

3 views

---

### NSA Releases Phase One and Phase Two of the Zero Trust Implementation Guidelines

The National Security Agency (NSA) has published Phase One and Phase Two of its Zero Trust Implementation Guidelines (ZIGs) , marking another step in its broader effort to help U.S. government and defence organisations reach the Department of War's ...

Steve McIntosh

2/6/2026

4 views

---

## Chrome Add-On Caught Stealing Amazon Commissions`

Security researchers at Socket have uncovered a Chrome extension called Amazon Ads Blocker that was quietly hijacking Amazon affiliate links to divert commissions to the extension developer — all while presenting itself as a simple ad-blocking tool....

Steve McIntosh

2/6/2026

5 views

---

## Russian Hackers Weaponize Microsoft Office Bug in Just 3 Days

In the latest illustration of how quickly attackers can exploit newly disclosed flaws, Russia's notorious APT28 cyber-espionage group has begun abusing a recently patched Microsoft vulnerability to steal emails and deploy malicious payloads against ...

David Robertson

2/4/2026

13 views

---

## Preventing Premium-Rate Telephone Fraud

Premium-rate and high-cost phone fraud is a growing threat to higher and further education institutions. Criminals exploit phone systems, softphones and staff behaviour to generate large, unexpected call charges, often routed through overseas premiu...

David Robertson

2/4/2026

8 views

---

## Attackers Harvest Dropbox Logins Via Fake PDF Lures

A new phishing scheme aims to trick organizations into giving up their Dropbox logins using a multistage obfuscation strategy. Main Article Data security vendor Forcepoint on Monday published research concerning an email-based social engineering cam...

David Robertson

2/4/2026

12 views

[Go To Site](#)

[Get the SharePoint Mobile App](#)

---

This email is generated through HEFESTIS Ltd's use of Microsoft 365 and may contain content that is controlled by HEFESTIS Ltd.

## CISO-Share Office Weekly Newsletter

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are also a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

---

### SC3 Daily Threat Summaries and Weekly Report

<https://www.cyberscotland.com/news/daily-threat-reports/>



A daily breakdown of various cyber threats from the Scottish Cyber Coordination Centre (SC3)

#### Action Point:

All SC3 threat intelligence in one place.

---



### Attackers Harvest Dropbox Logins Via Fake PDF Lures

<https://www.darkreading.com/cloud-security/attackers-harvest-dropbox-logins-fake-pdf-lures>

A new phishing scheme aims to trick organizations into giving up their Dropbox logins using a multistage obfuscation strategy.

Data security vendor Forcepoint on Monday published research concerning an email-based social engineering campaign observed in the wild. It follows a pattern [often seen](#): The threat actor sends an email to the target requesting the latter [open a linked PDF](#) to review a phony "request order."

The PDF includes a link to log in to a believable-yet-fake Dropbox phishing site; the target is asked to use his or her professional email address to log in and review the "order," with

reassurance that once the target does so, a response will automatically be sent to the email sender. The threat actor harvests the target's Dropbox credentials and location data, while the phishing site spits back an "incorrect username/password" message.

### Action Point:

One aspect that makes this campaign stand out is that neither the PDF nor the email nor the phishing site includes conventional malware of any kind. [Credential theft](#) is the end goal. While that might (reasonably) make one ask, "So what?" this and other aspects paint a portrait of an unexpectedly thoughtful scheme. And if they're able to bypass security checks and reach employee inboxes, it's a scheme worth being aware of.

Many [phishing](#) best practices remain useful here. Don't open a PDF attachment unless you can guarantee it came from a trusted source. Before opening any untrusted attachment or getting a suspicious email, get verbal or visual secondary confirmation from the person that sent it (such as via a phone call) or a relevant decisionmaker from within the organization. If you are given an urgent call to action to do something like log in to a website via your business credentials, take a moment to evaluate the request critically.



### Preventing Premium-Rate Telephone Fraud

Premium-rate and high-cost phone fraud is a growing threat to higher and further education institutions. Criminals exploit phone systems, softphones and staff behaviour to generate large, unexpected call charges, often routed through overseas premium numbers and disguised as legitimate service delivery. This can result in significant financial loss, operational disruption and reputational damage.

Institutions are vulnerable because they operate complex telephony estates, support many users, and increasingly rely on cloud collaboration platforms and mobile devices. Attacks commonly involve hacked phone systems making automated calls, phishing messages instructing staff to call fake "support" numbers, missed-call scams, and fraud linked to fake recruitment or services.

### Action Point:

The core defences are relatively straightforward:

- Clear policy on premium and international numbers.
- Strong technical controls on call destinations, spend, and administrative access.
- Routine monitoring of call and billing data.
- Simple incident response playbooks.
- Targeted staff and student awareness.

When applied consistently, these measures significantly reduce the likelihood and impact of premium-rate telephone abuse, regardless of the underlying telephony or UC platform.



## Russian Hackers Weaponize Microsoft Office Bug in Just 3 Days

<https://www.darkreading.com/cyberattacks-data-breaches/russian-hackers-weaponize-office-bug-within-days>

In the latest illustration of how quickly attackers can exploit newly disclosed flaws, Russia's notorious APT28 cyber-espionage group has begun abusing a recently patched Microsoft vulnerability to steal emails and deploy malicious payloads against organizations in Central and Eastern Europe.

CVE-2026-21509 is a security feature bypass vulnerability in Microsoft Office for which Microsoft [rushed an out-of-cycle patch](#) on Jan. 26 after confirming active zero-day exploitation. The US Cybersecurity and Infrastructure Security Agency (CISA) added the flaw to its database of known exploited vulnerabilities at the time.

### Speedy Exploit

According to Zscaler researchers, APT28 began exploiting the flaw just three days later, on Jan. 29, as part of a campaign they are tracking as Operation Neusplit. The attacks rely on specially crafted Microsoft Rich Text Format (RTF) documents to trigger the vulnerability and kick off a multistage infection chain that delivers different malicious payloads, Zscaler said in a [report this week](#).

**Action Point: Fast patching of security updates remains key to robust cybersecurity.**

[APT28](#), also known as Fancy Bear, Sofacy, and Sednit, is a Russia-linked advanced persistent threat (APT) group that the US government and others have linked to Russia's GRU military intelligence service. The [cyberespionage group](#) has been active since at least 2007 and is known for its ability to rapidly weaponize new vulnerabilities and constantly evolve its arsenal of malicious tools. It is associated with numerous attacks on government entities, military organizations, security firms, and critical infrastructure targets in North America, Europe, and elsewhere. Other high profile attacks include a breach of the Democratic National Committee and attacks on the World Anti-Doping Agency.



## Chrome Add-On Caught Stealing Amazon Commissions`

<https://www.techrepublic.com/article/news-amazon-ads-blocker-extension-affiliate-hijacking/>

Security researchers at Socket have uncovered a Chrome extension called **Amazon Ads Blocker** that was quietly hijacking Amazon affiliate links to divert commissions to the extension developer — all while presenting itself as a simple ad-blocking tool. At first glance, the extension appears legitimate: it hides sponsored ads on Amazon using basic CSS rules. But behind the scenes, it runs a second, undisclosed behaviour layer focused entirely on monetisation.



Whenever a user loads an Amazon product page, the extension scans for links matching common product patterns like /dp/ or /gp/product/. It then injects the developer's affiliate tag (**10xprofit-20**) into every link. If another creator's affiliate tag is already there, the extension overwrites it; if the link lacks a tag, it adds one. To ensure this behaviour can't be bypassed, it uses a **MutationObserver** to watch for dynamically loaded content (infinite scroll, page reloads, additional products) and reapply the hijack instantly.

Crucially, the extension never informs the user, never asks for consent, and never exposes this function in its interface. There are no pop-ups, settings, or permissions explaining that the extension modifies URLs or intercepts affiliate revenue. This is explicitly banned under Chrome Web Store policy, which prohibits automatic affiliate tag insertion or replacement.

Socket's investigation also found that this extension is just **one part of a much larger network** — at least **29 related extensions** targeting popular shopping platforms like Amazon, AliExpress, Shopify, Best Buy, and Shein. They all share infrastructure and behavioural patterns, suggesting a coordinated monetisation campaign rather than an isolated incident.

While affiliate hijacking doesn't compromise devices or steal sensitive data, it does represent **abuse of trust, violation of platform policies, and unauthorised manipulation of user browsing**, which can have broader implications for enterprise environments. Browser extensions remain a widely overlooked attack surface: they can read and modify traffic, intercept keystrokes, scrape pages, and interact with internal web apps — often with minimal oversight.

This event reinforces the importance of treating browser extensions not as “personal tools,” but as **software with access to sensitive enterprise data**, especially in organisations relying heavily on browser-delivered SaaS.

## Action Points

1. **Implement browser extension allowlisting** across managed devices — only approved, vetted extensions permitted.
  2. **Review installed extensions** for suspicious permissions (e.g., “Read and change all your data on all websites”).
  3. **Audit extension update histories** — sudden permission changes are red flags.
  4. **Establish an incident response flow** for browser-based threats: identify, collect evidence, remove, report to vendors, and communicate impact.
-



## NSA Releases Phase One and Phase Two of the Zero Trust Implementation Guidelines

<https://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-View/Article/4393480/nsa-releases-phase-one-and-phase-two-of-the-zero-trust-implementation-guidelines/>

The National Security Agency (NSA) has published **Phase One and Phase Two** of its **Zero Trust Implementation Guidelines (ZIGs)**, marking another step in its broader effort to help U.S. government and defence organisations reach the Department of War’s (DoW) *Target-level Zero Trust maturity*. These documents follow the release of the **Primer** and **Discovery Phase** earlier in January 2026 and are intended to guide system owners, cybersecurity teams, and enterprise stakeholders through what can otherwise be an overwhelming, highly technical transformation.

The core goal of the ZIGs is to provide a **practical, phased roadmap** for moving from early discovery and self-assessment to building, integrating, and maturing Zero Trust capabilities. Rather than mandating a rigid, one-size-fits-all model, NSA emphasises **modularity and customisation**, allowing organisations to pick and adapt activities depending on their environment, resourcing, and strategic priorities.

**Phase One** lays out **36 foundational activities** designed to establish a secure baseline environment capable of supporting Zero Trust. These activities map to **30 specific ZT capabilities**, ranging from identity hardening and asset visibility to network segmentation, data tagging, analytics, and early policy enforcement. Essentially, Phase One is about creating the “plumbing” needed before advanced Zero Trust functions can work effectively.

**Phase Two** moves from preparation to actual **solution integration**, detailing **41 activities** tied to **34 capabilities**. This stage focuses on embedding Zero Trust controls into operational systems — things like adaptive access control, continuous authentication, identity-driven policy engines, enhanced telemetry, and better authorisation workflows. Together, these support a transition toward automated, contextual decision-making rather than static perimeter-based security.

The NSA stresses that organisations *should first review the Primer and Discovery guidance* before diving into the Phases themselves. These earlier documents help teams understand their operational environments, existing security posture, and technology dependencies — preventing missteps later.

The overall message is clear: implementing Zero Trust isn’t about buying products or flipping a switch. It’s a structured, iterative programme that requires sequencing, alignment, and clarity on dependencies. NSA’s ZIG suite now provides a more complete blueprint for organisations at any stage of maturity, whether they’re just starting or already midway through transformation.

## Action Points

### Immediate Actions

1. **Review the Zero Trust Primer and Discovery Phase** to understand prerequisites before implementing Phases One and Two.
2. **Download and read the Phase One & Phase Two ZIGs** — they define required activities, dependencies, and capability outcomes.

### Strategic Programme Planning

3. **Conduct a Zero Trust capability gap assessment** against the 30 Phase One and 34 Phase Two capabilities.
4. **Prioritise foundational activities first**, especially around identity, device visibility, telemetry, and policy governance.

### Operational Measures

5. **Map NSA activities to your existing architecture**, including MFA, SIEM, data tagging, PAM, encryption, and access control.
6. **Establish cross-team ownership** (identity, networking, cloud, governance, SOC).

### Longer-Term

7. **Define your organisation's Target-level maturity goal** using the DoW CIO framework.
8. **Develop a phased roadmap** incorporating both foundational and advanced capabilities, aligned with NSA guidance.

---

All the best from all of us at **HEFESTIS** and look out for our next **Threatscape** update next week.

*HEFESTIS Limited, Registered Office: Unit 27, Stirling Business Centre, Wellgreen, Stirling FK8 2DZ*

*Incorporated in Scotland SC603511*

## SharePoint



CISO Share

### **ThreatScape Newsletter**

Dear all,

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).



### [SC3 Daily Threat Summaries and Weekly Report](#)

Please find SC3's daily threat summaries for this week for those who do not receive this information directly. Action Point: All threat reports, weekly and monthly, are available from this link -

<https://www.cyberscotland.com/news/sc3-threat-reports...>

David Robertson

2/12/2026

14 views



### [Microsoft to Enable 'Windows Baseline Security' With New Runtime Integrity Safeguards](#)

As part of the Secure Future Initiative announced in November 2023, the company is moving towards having runtime integrity safeguards enabled by default in Windows. Main Article The enhancement, called Windows Baseline Security Mode, will ensure tha...

David Robertson

2/13/2026



### [A method to assess 'forgivable' vs 'unforgivable' vulnerabilities](#)

Research from the NCSC designed to eradicate vulnerability classes and make the top-level mitigations easier to implement. Main article All systems contain vulnerabilities. In fact, the number of Common Vulnerabilities and Exposures (CVEs) in commod...

David Robertson

2/13/2026

6 views



### [47-Day SSL/TLS Mandates: A Step Towards Transitioning to Automation](#)

The article explains a major shift in SSL/TLS certificate policy: certificates are moving from 13-month validity to an extremely short 47-day lifespan , driven by new CA/Browser


Forum rules. This is one of the largest changes in digital trust in dec...

Steve McIntosh

2/13/2026

5 views

---



### [AI Drives Doubling of Phishing Attacks in a Year](#)

Security filters caught one phishing email every 19 seconds in 2025, more than double the rate a year previously, Cofense has revealed. AI technology is helping threat actors to increase the speed and scale of attacks, to the point where detected ph...

David Robertson

2/11/2026

11 views

---



### [Cyber Threat Intelligence Digest: February 2026 – Week 6 from Acumen Cyber \(11 Feb 2026\)](#)

The Cyber Threat Intelligence Digest: February 2026 – Week 6 from Acumen Cyber outlines significant vulnerabilities, active threat campaigns, and notable cyber-security developments observed during the week. Main Article The Week 6 Cyber Threat Inte...

David Robertson

2/13/2026

3 views

[Go To Site](#)



## CISO-Share Office Weekly Newsletter

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

## Table of Contents

<b>CISO-Share Office Weekly Newsletter</b> .....	1
SC3 Daily Threat Summaries and Weekly Report.....	1
Microsoft to Enable 'Windows Baseline Security' With New Runtime Integrity Safeguards...	2
A Method To Assess 'Forgivable' Vs 'Unforgivable' Vulnerabilities .....	3
47-Day SSL/TLS Mandates: A Step Towards Transitioning to Automation .....	4
AI Drives Doubling of Phishing Attacks in a Year.....	6
Acumen Cyber Threat Intelligence Digest: Week 6 .....	8

## SC3 Daily Threat Summaries and Weekly Report



Please find SC3's daily threat summaries for this week for those who do not receive this information directly.

### All Daily, Weekly and Monthly Reports:

<https://www.cyberscotland.com/news/sc3-threat-reports/>



SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu  
lletin-9-February-20lletin-10-February-2lletin-11-February-2lletin-12-February-2lletin-13-February-2l



SC3-Weekly-Vulnerability-Report-9-February-2026 SC3 - Monthly ransomware report



## Microsoft to Enable 'Windows Baseline Security' With New Runtime Integrity Safeguards

As part of the [Secure Future Initiative](#) announced in November 2023, the company is moving towards having runtime integrity safeguards enabled by default in Windows.

### [Main Article](#)

The enhancement, called Windows Baseline Security Mode, will ensure that only properly signed applications, drivers, and services can run, thus preventing tampering and unauthorized changes.

For those cases where exceptions are needed, users and administrators will have the option to override the safeguards.

"Developers can also check whether these protections are active and whether any exceptions have been granted — giving them insight and control over the conditions under which their apps run," Microsoft notes.

The tech giant announced the improvement simultaneously with revealing that Secure Boot certificates will begin to expire in June, and that [refreshed certificates](#) will be rolled out to supported Windows releases.

**Action Point:** The idea behind the newly detailed security and privacy improvements, Microsoft notes, is to provide users with better visibility and consent control over how applications access their files, camera, microphone, and other sensitive resources.

"We will begin by giving users and IT admins visibility into how apps and agents behave in the system. For developers, Windows will provide tools and APIs to streamline adoption. Their existing well-behaved apps will continue to work, giving developers the time and runway to adhere to the new, stronger security and privacy posture of Windows," the company says.

Microsoft will roll out these enhancements in phases, working together with developers and partners for guidance adjusted based on their feedback.

---



## A Method To Assess 'Forgivable' Vs 'Unforgivable' Vulnerabilities

Research from the NCSC designed to eradicate vulnerability classes and make the top-level mitigations easier to implement.

### Main article

All systems contain vulnerabilities. In fact, the number of Common Vulnerabilities and Exposures (CVEs) in commodity technology continues to rise. While there are a number of factors that are driving the increasing numbers, the NCSC expect this trend to continue unless interventions are made.

We know many vulnerabilities are complex and hard to avoid. But vulnerabilities that are trivial to find (and that occur time and time again) are ones the NCSC are aiming to drive down at scale.

These 'unforgivable vulnerabilities', a phrase coined by [Steve Christie in his 2007 MITRE paper](#), 'are beacons of a systematic disregard for secure development practices. They simply should not appear in software that has been designed, developed, and tested with security in mind'.

This paper extends the ideas in the MITRE paper and proposes a method to assess a vulnerability as 'forgivable' or 'unforgivable'. More importantly, this paper intends to generate discussion with vendors, and is a call on them to work to eradicate vulnerability classes and make the top-level mitigations discussed in this paper easier to implement.

**Action Point:** The method of assessing forgivable and unforgiveable vulnerabilities discussed in this paper could be matured to take account of other factors, including:

- the role of the product and the severity of the vulnerability (for example, if the product is an internet-facing service and the vulnerability can be exploited pre-authentication)
- how to determine the vulnerability management maturity of the vendor/developer (for example, is the vendor monitoring vulnerabilities throughout the expected lifecycle of the product)
- is being unaware of a vulnerability (that is later exploited in the wild) 'unforgivable', and should other factors also be taken into consideration (including know bug/vulnerability classes and their typical mitigations)



## 47-Day SSL/TLS Mandates: A Step Towards Transitioning to Automation

The article explains a major shift in SSL/TLS certificate policy: certificates are moving from 13-month validity to an extremely short **47-day lifespan**, driven by new CA/Browser Forum rules. This is one of the largest changes in digital trust in decades, and it fundamentally alters how organisations will need to manage certificates.

### Main Article

Historically, SSL/TLS certificates lived for years, then were shortened to 398 days, which was manageable with spreadsheets, reminders, and manual renewal rituals. But the upcoming shift — eventually bringing certificates down to **47 days** — makes manual handling practically impossible. Most organisations already struggle with certificate renewals today; with lifespans measured in weeks, manual processes become guaranteed outage generators.

The logic behind the mandate is straightforward:

- Shorter lifespans **limit exposure** when a certificate is compromised.
- Faster expirations **force adoption** of new security standards and cryptography.
- Frequent renewal cycles **push organisations into automation**.
- Reduced validity **shrinks attacker windows** for MITM, impersonation, and phishing abuse.

The rollout phases are gradual but meaningful. Starting **March 2026**, certificate validity begins shrinking in set stages: first to 398 days, then 200, then 100, and finally 47 days by **March 2029**. Although it appears slow on paper, the article stresses that organisations underestimate how quickly renewal frequency becomes overwhelming. What looks “far away” becomes operational pressure almost immediately.

Short-lived certificates change the entire nature of certificate management. Renewals become continuous, not occasional. Infrastructure must rotate certificates as part of normal operations. Load balancers, reverse proxies, Kubernetes sidecars, service meshes, IoT devices — all of them will need automated certificate flows. Without automation, organisations will experience cascading outages as a direct result of expired certificates.

The article also highlights that most organisations operate in reactive mode today: they only look at certificates when they break production. The new regime demands proactive, automated lifecycle management. The cultural shift is as important as the technical one.

Automation becomes non-negotiable. It prevents outages, reduces risk, stabilises uptime, and satisfies compliance requirements. The future of certificate management is ACME-driven issuance, renewal, deployment, revocation, and monitoring — all without human intervention.

SSL certificates are now expiring faster than avocados. Yes... avocados. You buy them green, blink twice, and suddenly they're brown and useless. That's exactly what's happening to SSL/TLS certificates. Not long ago, certificates lasted years.

## Action Points

### Immediate

1. **Create an inventory** of all public and internal certificates across apps, cloud workloads, devices, and networks.
2. **Identify manual renewal points** — spreadsheets, calendar reminders, “single-person knowledge” risks.

### Strategic

3. **Adopt a certificate automation platform** (ACME, PKIaaS, or centralised CLM system).
4. **Integrate automation into CI/CD** to ensure certificate rotation is seamless.

### Operational

5. **Update load balancer, reverse proxy, and cloud configs** to support automated renewal flows.
6. **Redesign private PKI policies** to eliminate long-lived internal certificates.

### Risk Management

7. **Evaluate outage risk** from legacy certificates and plan phased migration.
  8. **Educate teams** (DevOps, cloud, security, platform engineering) on the upcoming 47-day mandate and the inevitability of automation.
-



## AI Drives Doubling of Phishing Attacks in a Year

Security filters caught one phishing email every 19 seconds in 2025, more than double the rate a year previously, Cofense has revealed.

AI technology is helping threat actors to increase the speed and scale of attacks, to the point where detected phishing emails last year far outstripped 2024 figures of one every 42 seconds, the cybersecurity firm claimed.

### [Full Cofense Report \(PDF\)](#)

The security vendor's latest report, *The New Era of Phishing: Threats Built in the Age of AI*, is based on its own threat intelligence.

"Threat actors no longer experiment with AI in isolated ways. Instead, they use it as a core capability to generate, test, and deploy phishing campaigns at scale," the report warned.

Thanks to Infosecurity Magazine for providing this [summarised report](#).

"The result is phishing that is faster, more adaptive, and more convincing than ever before, giving rise to polymorphic, multi-channel campaigns that continuously change their appearance while preserving the same malicious intent."

AI is helping threat actors in several ways, the report claimed. Most obvious is the ability it gives them to compose emails in near-flawless local languages.

Cofense said that "conversational" phishing emails (ie those not including malicious attachments, QR codes or links) accounted for 18% of the total. That speaks to the growth of business email compromise (BEC) attacks.

Other trends include:

**Highly personalized campaigns:** Cofense observed a rise in campaigns where the same phishing website delivered different payloads depending on the type of machine/device it was accessed from. AI might also be helping campaigns to serve up different spoofed brands depending on the browser, or optimize credential harvesting pages specifically for mobile users, among other things.

**Polymorphism by default:** AI is helping threat actors to dynamically alter logos, signatures, wording, and URLs and files according to the specific victim. Three-quarters (76%) of initial infection URLs identified by Cofense were unique. AI is also scraping publicly available data from the web in order to personalise attacks

**Action Point:** Cofense also reported a 105% annual increase in detections of legitimate and malicious remote access tools (RATs) in 2025.

Software like ConnectWise ScreenConnect and LogMeln's GoTo Remote Desktop can be used to bypass traditional security. It is often combined with social engineering, whereby a user is tricked into downloading the tool to give an attacker access to 'fix' a non-existent issue.

"To continue to execute campaigns involving a large number of systems infected by legitimate RATs, threat actors increasingly rely on automation and AI in their workflows," the report noted.

Another trend is of attackers flocking to the .es TLD for credential phishing. Cofense observed use of .es domains in these attacks increase 19-fold from the fourth quarter of 2024 to the first quarter of 2025. That makes the domain the third-most commonly abused.

The report also recorded a 204% increase in phishing emails delivering malware last year, compared to 2024.

"Together, these patterns demonstrate why phishing must be analysed after delivery, where behavioural context and human validation expose threats that evade static, perimeter-based controls," [Cofense claimed](#).

---





## Acumen Cyber Threat Intelligence Digest: Week 6

The **Cyber Threat Intelligence Digest: February 2026 – Week 6** from Acumen Cyber outlines significant vulnerabilities, active threat campaigns, and notable cyber-security developments observed during the week.

### Main Article

The Week 6 Cyber Threat Intelligence Digest highlights significant vulnerabilities, emerging threats, geopolitical incidents, and broader security trends affecting organisations globally. Key **software vulnerabilities** patched this week include CVE-2026-1731 in BeyondTrust Remote Support and Privileged Remote Access, a critical OS command injection flaw now fixed in version 25.1.1; CVE-2026-0755 in Trend Micro's Gemini-Mcp-Tool, a critical remote code execution bug with no vendor patch yet (mitigation via access restrictions advised); and CVE-2026-21643 in Fortinet's FortiClientEMS, a critical SQL injection issue remediated in version 7.4.5.

Under **potential threats**, Microsoft Security details a new *CrashFix* variant of the ClickFix campaign: this attack weaponises browser crashes and fake recovery prompts to deploy ModeloRAT, a remote access trojan with persistence and C2 capabilities. Huntress reported an intrusion abusing compromised SonicWall SSLVPN credentials and a Bring Your Own Vulnerable Driver (BYOVD) technique to deploy a kernel-level EDR-killer payload disguised as firmware, indicating ransomware precursor activity. The European Commission disclosed an unauthorised breach of its central mobile device management infrastructure, potentially exposing staff contact details, though no managed device compromise has been confirmed.

**General news** captures broader regulatory and platform developments, including the EU threatening TikTok with heavy fines under the Digital Services Act over addictive features, and Discord mandating global age verification via video selfie or government ID to bolster child safety. Mandiant also reported North Korean threat group UNC1069 targeting a crypto executive with a deepfake Zoom scam deploying backdoors and credential-stealing tools.

The report concludes with threat actor trend graphs and mitigation actions, urging patching of the identified CVEs and proactive hunting by vulnerability management teams.

### **Action Point:** Remediation Actions

Following the information provided above, we recommend that the technologies mentioned be fully patched and updated. We also want to highlight and recommend applying the following patches where applicable:

- **CVE-2026-1731 (BeyondTrust)** – This vulnerability can be addressed by upgrading BeyondTrust RS/PRA to version 25.1.1 or later.
- **CVE-2026-0755 (Gemini-Mcp-Tool)** – This vulnerability can be addressed by restricting access to trusted environments, as there has been no patch specifically for Gemini-Mcp-Tool.
- **CVE-2026-21643 (FortiClientEMS)** – This vulnerability can be addressed by patching FortiClient EMS to version 7.4.5.

**All the best from all of us at HEFESTIS and look out for a new ThreatScape update next week.**

## SharePoint



### CISO Share

#### Threatscape

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).



#### [Acumen Cyber Threat Intelligence Digest: Week 7 2026](#)

This week's threat landscape is very much "patch or perish". Across all major platforms — Google, Apple and Mozilla — we've seen multiple actively exploited vulnerabilities getting emergency fixes. The common thread is simple: attackers are moving q...

Steve McIntosh

2/20/2026



#### [SC3 Daily Threat Summaries and Weekly Report](#)

Action Point: Please find SC3's daily threat summaries for this week for those who do not receive this information directly.

Steve McIntosh

2/20/2026




### [Prompt Control is the New Front Door of Application Security](#)

Traditional application security has always treated the “front door” as the network edge: APIs, authentication gateways, load balancers, firewalls. But according to new research discussed in the article, the real security pressure point in AI-driven...

Steve McIntosh

2/19/2026



### [Notepad++ Fixes Hijacked Update Mechanism Used to Deliver Targeted Malware](#)

Notepad++ has released a major security update (version 8.9.2 ) to fix weaknesses exploited in a targeted supply-chain attack attributed to a China-linked threat group known as Lotus Panda . This attack, active since June 2025 , involved hijacking t...

Steve McIntosh

2/19/2026



### [Password Managers Vulnerable to Vault Compromise Under Malicious Server](#)

A team of security researchers from ETH Zurich in Switzerland has analyzed popular password managers and identified ways in which threat actors could compromise users' vaults and access sensitive data. Main Article However, the researchers did not t...

David Robertson

2/18/2026



### [Over 300 Malicious Chrome Extensions Caught Leaking or Stealing User Data](#)

Security researchers have discovered more than 300 Chrome extensions that leak browser data, spy on their users, or outright steal users' data. Main Article Research focused on the analysis of network traffic generated by Chrome extensions has unco...

David Robertson

2/18/2026

[Go To Site](#)

## CISO-Share Office Weekly Newsletter

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are also a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

---

### SC3 Daily Threat Summaries and Weekly Report

<https://www.cyberscotland.com/news/daily-threat-reports/>



A daily breakdown of various cyber threats from the Scottish Cyber Coordination Centre (SC3)

#### Action Point:

All SC3 threat intelligence in one place.



### Acumen Cyber Threat Intelligence Digest: Week 7 2026

<https://acumencyber.com/cyber-threat-intelligence-digest-february-2026-week-7>

This week's threat landscape is very much "patch or perish". Across all major platforms — Google, Apple and Mozilla — we've seen multiple actively exploited vulnerabilities getting emergency fixes. The common thread is simple: attackers are moving quickly, and vendors are patching just as fast. Organisations need to keep pace.

#### Google Chrome: Actively Exploited CVE-2026-2441

Google pushed critical patches on **13 February 2026** to fix a **use-after-free vulnerability in the CSS component of Chrome**. Versions prior to **145.0.7632.75** (Windows/macOS)

and **144.0.7559.75** (Linux) are affected. Successful exploitation lets attackers run arbitrary code via a malicious webpage — exactly the sort of thing drive-by campaigns love. No actor attribution yet, and Google's keeping exploitation details tight.

## Apple: Actively Exploited Zero-Day CVE-2026-20700 (dyld)

Apple also had a busy week, patching an **actively exploited memory-corruption flaw** affecting the **Dynamic Link Editor (dyld)**. This hits a huge range of devices — modern iPhones, iPads, Macs and even visionOS. The bug allows arbitrary code execution if an attacker can manipulate memory, which makes it a prime target for privilege-escalation chains. Fixes shipped with **iOS/iPadOS/macOS/watchOS/tvOS/visionOS 26.3**.

## Mozilla: CVE-2026-2447 in libvpx

Mozilla patched a **heap buffer overflow** affecting **Firefox, Firefox ESR, and Thunderbird**. While it hasn't been exploited in the wild yet, it enables **remote code execution**. If you're running outdated versions, you'll want to get patched quickly — RCE bugs in browsers rarely stay unexploited for long once public.

## Action Points

- **Push Chrome updates immediately** — prioritise anyone running Linux where the patched version number differs.  
(Patch covers an active exploit.)
- **Force-update Apple devices** across your estate — especially student-owned or BYOD hardware.  
(The dyld issue is being exploited already.)
- **Update Firefox/Thunderbird** organisation-wide — even though it isn't actively exploited, this is a high-risk RCE vulnerability.
- **Review browser update compliance** via your device management platform (Intune/SCCM/JAMF).



## Prompt Control is the New Front Door of Application Security

<https://securityboulevard.com/2026/02/prompt-control-is-the-new-front-door-of-application-security/>

Traditional application security has always treated the “front door” as the network edge: APIs, authentication gateways, load balancers, firewalls. But according to new research discussed in the article, the real security pressure point in **AI-driven architectures** has moved **upstream** — **into the inference layers** where prompts, tokens, and model outputs are generated.

The article highlights a major shift: the **prompt layer** is now viewed as the most impactful point for both security (25% of respondents) and application delivery (29%). Token layers



follow closely (23% each). The output layer still matters, but ranks lower. Why? Because inference layers aren't infrastructure boundaries — they're **behavioural boundaries**. They determine *how* a model thinks, not just *what* it outputs.

Prompts are where intent enters the system. They define reasoning style, context retention, and the degree to which safeguards can be bypassed. Attackers increasingly target this layer with **prompt injection, context poisoning, and memory manipulation** — all before the model produces a single token. If organisations rely on output filtering alone, they're detecting problems after damage has already begun.

Token governance has also become a **security primitive**. Tokens directly map to **cost, capacity, and abuse surface**. Without token-level controls, attackers don't need to find vulnerabilities — they just force the model to generate excessive tokens until cost spikes, latency degrades, or availability collapses. It's a modern denial-of-service vector unique to AI workloads.

Output moderation still matters, but it's not enough. Filtering hallucinations or sensitive-data leaks after generation is useful, but purely reactive. Organisations relying heavily on output controls are implicitly accepting "assume breach" inside the model.

The research also shows authentication (55%) and observability (54%) remain foundational. But what needs authenticating and monitoring has changed: **prompts, tokens, sessions, and routing decisions**, not just API calls.

The big takeaway? In AI-powered systems, the "front door" has moved. The most resilient organisations treat **prompt handling, token controls, and inference routing** as first-class security domains, not afterthoughts. As AI agents begin chaining tasks and acting independently, the article warns that late adopters will struggle to regain control.

## Action Points

### Secure the Prompt Layer

1. Implement prompt-validation and sanitisation controls.
2. Enforce strict boundaries between user input and system instructions.

### Control Tokens (Cost, Abuse, and Safety)

3. Set per-request and per-session token limits.
4. Use token-shaping or streaming controls to detect abusive or anomalous usage.
5. Monitor for token-based DoS patterns.

### Harden Output Controls

6. Maintain toxicity, hallucination, and data-leak filters — but treat them as secondary safeguards.


### Strengthen Identity & Observability

7. Authenticate *who* is making inference requests (users, services, agents).
8. Log and analyse prompts, token flows, session activity, and routing decisions.

### Govern the Inference Layer

9. Treat prompt, token, and output layers as distinct security domains with dedicated policy sets.
10. Prepare for autonomous AI agents — implement guardrails now before scale makes control difficult.





## Notepad++ Fixes Hijacked Update Mechanism Used to Deliver Targeted Malware

<https://thehackernews.com/2026/02/notepad-fixes-hijacked-update-mechanism.html>

Notepad++ has released a major security update (version **8.9.2**) to fix weaknesses exploited in a **targeted supply-chain attack** attributed to a China-linked threat group known as **Lotus Panda**. This attack, active since **June 2025**, involved hijacking the Notepad++ update traffic at the hosting-provider level and redirecting certain users to malicious servers, where they were served a **backdoored update** containing a previously unknown malware strain called **Chrysalis**. The issue was only discovered in **December 2025**.

To prevent similar attacks in the future, Notepad++ maintainer Don Ho has introduced what he calls a “**double-lock**” **update validation model**. Previously, Notepad++ only verified the signed installer downloaded from GitHub (a feature introduced in 8.8.9+). Now, version 8.9.2 also validates the **signed XML metadata** provided by the update server (notepad-plus-plus[.jorg]). This means attackers would have to compromise both GitHub distribution and the official update metadata — raising the difficulty significantly.

Additional improvements were made to **WinGUP**, the auto-update component. These include:

- **Removing libcurl.dll** to eliminate DLL side-loading opportunities
- Dropping insecure SSL options **CURLSSLOPT\_ALLOW\_BEAST** and **CURLSSLOPT\_NO\_REVOKE**
- **Restricting plugin installations** so only programs signed with the same certificate as WinGUP can manage plugins

The update also fixes **CVE-2026-25926**, a high-severity vulnerability caused by Notepad++ calling **Windows Explorer** without specifying an absolute path. This could allow attackers to plant a malicious explorer.exe in the working directory and achieve **arbitrary code execution**.

The broader supply-chain incident is tracked as **CVE-2025-15556**, rated 7.7. Both Rapid7 and Kaspersky attributed the tampered updates to Lotus Panda, which selectively targeted victims rather than mass-distributing malware. The attack underlines how threat actors are evolving: instead of poisoning the software itself, they compromised **the hosting infrastructure** and shaped update traffic based on victim profiles.

Notepad++ urges all users to **upgrade to v8.9.2** and only download installers from the official site. This incident illustrates that even small but popular tools like Notepad++ can become high-value supply-chain targets, especially for espionage-motivated adversaries.

### Action Points

#### Immediate

1. **Upgrade to Notepad++ 8.9.2** across all endpoints.

2. **Verify installer origin** (only download from official domain).
3. **Scan systems** that installed updates between June–Dec 2025 for indicators of the Chrysalis backdoor.

## **Threat Hunting**


4. Look for unusual outbound connections from Notepad++.
5. Check for rogue explorer.exe files or suspicious working-directory behaviour.
6. Monitor plugin folders for unsigned or unexpected binaries.

## **Supply-Chain Hygiene**

7. Harden update pathways for all developer/admin tools.
8. Require signed-metadata validation where available.
9. Add update-mechanism review to your vendor risk process

## **Long-term**

10. Educate staff on risks associated with “trusted” utilities — even lightweight tools can be weaponised.



## **Password Managers Vulnerable to Vault Compromise Under Malicious Server**

<https://www.securityweek.com/password-managers-vulnerable-to-vault-compromise-under-malicious-server/>

**A team of security researchers from ETH Zurich in Switzerland has analyzed popular password managers and identified ways in which threat actors could compromise users’ vaults and access sensitive data.**

However, the researchers did not test the password managers against external or client-side attacks. Instead they targeted zero-knowledge encryption, a security model where the service provider is unable to access the user’s encrypted data and the data should be protected even if the provider’s servers are compromised.

As such, the ETH Zurich researchers conducted an [analysis](#) of popular cloud-based password managers under the assumption that the servers storing user vaults are “fully malicious”. The researchers targeted password managers from Bitwarden, Dashlane, LastPass, and 1Password, each having millions of users and overall accounting for a significant share of the market. Although 1Password was included in the research, the analysis focused on the other password managers.

Several types of attacks were conducted against each of the tested password managers to degrade security guarantees, undermine expected protections, and fully compromise user accounts.

The experts targeted features used for account recovery and SSO login, as well as features designed for backward compatibility. They conducted attacks leveraging improper vault integrity and attacks enabled by sharing features, which allow families or businesses to use the same credentials.

## Action Point:

1Password has also been analyzed and the researchers managed to achieve full compromise of vault confidentiality and integrity, allowing an attacker to obtain passwords and other sensitive data stored in the vault, as well as to add items to the vault.

However, Jacob DePriest, CISO and CIO of 1Password, told *SecurityWeek* that the attack vectors identified by the researchers had already been documented in the company's publicly available [Security Design White Paper](#).

"We are committed to continually strengthening our security architecture and evaluating it against advanced threat models, including malicious-server scenarios like those described in the research, and evolving it over time to maintain the protections our users rely on," DePriest said.

He added, "For example, 1Password uses Secure Remote Password (SRP) to authenticate users without transmitting encryption keys to our servers, helping mitigate entire classes of server-side attacks. More recently, we introduced a [new capability](#) for enterprise-managed credentials, which from the start are created and secured to withstand sophisticated threats."



## Over 300 Malicious Chrome Extensions Caught Leaking or Stealing User Data

<https://www.securityweek.com/over-300-malicious-chrome-extensions-caught-leaking-or-stealing-user-data/>

**Security researchers have discovered more than 300 Chrome extensions that leak browser data, spy on their users, or outright steal users' data.**

Research focused on the analysis of network traffic generated by Chrome extensions has uncovered 287 applications transmitting the user's browsing history or search engine results pages (SERP).

Some of them, security researcher Q Continuum [explains](#), would essentially expose the data to unsecured networks, while others would send it to collection servers, either due to intended functionality, for monetization purposes, or with malicious intent.

The extensions have over 37.4 million users, the researcher says. Of these, roughly 27.2 million users installed 153 extensions that were confirmed to leak browser history upon installation.

Q Continuum, who also flagged over 200 additional extensions as suspicious due to shared author details with the data-leaking ones, observed four scrapers connecting to the honeypot set up for the research.

Based on the observations, the researcher believes that a data broker rather than extension developers might be directly involved in the monetization of these applications

## Action Point:

Posing as AI assistance tools, all extensions had “the same internal structure, JavaScript logic, permissions, and backend infrastructure”, suggesting they are part of a single, coordinated operation.

One extension would render a full screen iframe pointing to a remote domain and allowing the attacker to load remote content to manipulate the UI directly.

The extension can also extract data from the active tab, supports message-triggered voice recognition, and includes explicit tracking pixel scripts.

According to LayerX, 15 extensions were seen specifically targeting Gmail, extracting email content and transmitting it to third-party infrastructure.

---

All the best from all of us at **HEFESTIS** and look out for our next **Threatscape** update next week.

*HEFESTIS Limited, Registered Office: Unit 27, Stirling Business Centre, Wellgreen, Stirling FK8 2DZ*

*Incorporated in Scotland SC603511*

## CISO-Share Office Weekly Newsletter

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

## Table of Contents

<b>CISO-Share Office Weekly Newsletter</b> .....	1
SC3 Daily Threat Summaries and Weekly Report .....	1
Hundreds of FortiGate Firewalls Hacked in AI-Powered Attacks: AWS .....	2
A Method To Assess 'Forgivable' Vs 'Unforgivable' Vulnerabilities .....	3
New 'Sandworm_Mode' Supply Chain Attack Hits NPM .....	4
Claude's New AI Vulnerability Scanner Sends Cybersecurity Shares Plunging .....	5
Acumen Cyber Threat Intelligence Digest: Week 8 .....	6

## SC3 Daily Threat Summaries and Weekly Report



Please find SC3's daily threat summaries for this week for those who do not receive this information directly.

## All Daily, Weekly and Monthly Reports:

<https://www.cyberscotland.com/news/sc3-threat-reports/>



SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu  
lletin-23-February-2lletin-24-February-2lletin-25-February-2lletin-26-February-2lletin-27-February-2l



SC3-Weekly-Vulner  
ability-Report-23-Fe

---



## Hundreds of FortiGate Firewalls Hacked in AI-Powered Attacks: AWS

The attacks, observed between January 11 and February 18, did not target known vulnerabilities. Instead, they focused on the exploitation of exposed device configurations across globally dispersed appliances.

### Main Article

According to AWS, the campaign was carried out by an unsophisticated threat actor that relied on multiple commercial gen-AI services to implement known attack techniques.

The hackers were seen scanning for management interfaces accessible via ports 443, 8443, 10443, and 4443, and using common credentials for initial access.

“The campaign’s targeting appears opportunistic rather than sector-specific, consistent with automated mass scanning for vulnerable appliances,” [AWS notes](#).


In some cases, multiple FortiGate devices belonging to the same organization were compromised. AWS says that some IP clusters point either to managed service provider deployments or to large organizational networks.

**Action Point:** “These plans reference academic research on offensive AI agents, suggesting the actor follows emerging literature on AI-assisted penetration testing. The AI produces technically accurate command sequences, but the actor struggles to adapt when conditions differ from the plan,” AWS notes.

On the threat actor’s infrastructure, AWS identified multiple scripts likely generated using AI, used to parse configurations, extract credentials, automate VPN connections, perform mass scanning, and aggregate results.

“The volume and variety of custom tooling would typically indicate a well-resourced development team. Instead, a single actor or very small group generated this entire toolkit through AI-assisted development,” AWS says.

The attacks, it notes, were likely mounted by a financially motivated, Russian-speaking threat actor with low-to-medium technical capability, based on the extensive reliance on AI across all operational phases.



## Einstein-Style AI Agents Could Reshape HE/FE — For Better and Worse

Tools like *Einstein* — full-autonomy AI agents that can log into VLEs, watch lectures, read materials, write assignments, and even post in discussions — represent a fundamental shift in how students might engage with education. And let's be honest: in HE/FE, this is both exciting and deeply disruptive.

### Main Article

On the positive side, autonomous AI agents point to a future where learners can get real, meaningful support rather than generic chatbot answers. The idea of an AI that can genuinely *understand* course content, pull out key ideas, and help a student study more effectively is powerful. Used ethically, it could level the playing field for students with additional learning needs, language barriers, or chaotic personal circumstances. Imagine pairing this with structured academic integrity guidelines — it could become a legitimate learning companion rather than a shortcut.

But the risks in HE/FE are enormous. A tool that logs into Canvas, completes assignments, and auto-submits them isn't "assistive"; it's full-scale academic outsourcing. This is plagiarism reimagined as a service. If students can pass modules without engaging with any learning materials themselves, the integrity of qualifications collapses.

For staff, detecting AI-generated work becomes significantly harder because submissions come from the student's own account, with context-aware, original responses.

Ultimately, the sector must shift from assessment that tests output to assessment that tests understanding — more vivas, more in-class work, more practical demonstrations. AI like this isn't going away. HE/FE is quickly facing a choice: adapt quickly, or risk making assessment meaningless.

### **Action Point:**

AUPs specifically prohibit sharing credentials with other humans, do they also cover sharing them with non human services, like AI agents?

This is something that needs to be considered and explored further to identify and manage. How would we spot someone using an AI like this?





## New 'Sandworm\_Mode' Supply Chain Attack Hits NPM

Dubbed [Sandworm\\_Mode](#), the attack was deployed through 19 packages published under two aliases, which relied on typosquatting to trick developers into executing the malicious code.

[Main Article](#)

According to cybersecurity firm Socket, the attack bears the hallmarks of the Shai-Hulud campaign that hit roughly 800 NPM packages in [September](#) and [November](#) 2025.

Sandworm\_Mode abuses stolen NPM and GitHub credentials for propagation and relies on a weaponized GitHub Action to harvest and exfiltrate CI secrets and to inject dependencies and workflows into repositories.

The malicious packages, all of which have been removed from the registry, rely on typosquatting to pose as popular developer utilities, crypto tools, and AI coding utilities, such as Claude Code and OpenClaw.

**Action Point:** Sandworm\_Mode executes a multi-stage attack, where the initial credential and crypto key exfiltration is followed by deep harvesting of secrets from password managers, MCP server injection, persistence via Git hooks, worm propagation, and multi-channel exfiltration.

"This two-phase design is deliberate: the most financially damaging operation, crypto key theft, runs instantly and unconditionally, while the noisier operations are deferred to evade short-lived sandbox analysis," Socket explains.

The code also contains a configurable but inactive dead switch capability to trigger home-directory wiping when losing access to GitHub and NPM.

The same as Shai-Hulud, Sandworm\_Mode propagates by infecting existing packages but can also use carrier packages for propagation, adding a dependency reference to trigger a pull request workflow in GitHub Action and harvest and exfiltrate all repository secrets, [EndorLabs](#) explains.

Developers are advised to remove any of the malicious packages they might have installed, to check their packages for recent changes to JSON files, rotate all GitHub and NPM credentials, tokens, and CI secrets, and check for unexpected workflows.



## Claude's New AI Vulnerability Scanner Sends Cybersecurity Shares Plunging

Anthropic announced on Friday that its AI-powered coding assistant Claude Code is being enhanced with a new capability designed for finding vulnerabilities.

### Main Article

The new capability is named [Claude Code Security](#) and it's currently available in limited preview to Enterprise and Team customers. It's designed to scan code for vulnerabilities and suggest patches. Developers can review the patch suggestions and decide whether they want to apply them.

Similar tools have been available for some time. [GitHub](#) has been offering AI-powered vulnerability remediation capabilities for years, and Google has also been making [significant progress](#) in this area.

While the new Claude capability is limited to finding vulnerabilities in code, the markets reacted to the announcement and the shares of major cybersecurity companies fell over fears that AI could replace their solutions.

Broader software stocks have faced pressure in recent weeks amid AI disruption concerns, and cybersecurity firms are now experiencing similar volatility.

**Action Point:** Joe Silva, CEO of vulnerability management firm Spektion, believes this moment represents a fundamental shift in application security that goes beyond tooling, and challenges the core assumptions of how defenders and attackers operate.

"Think of this as the ultimate red-team tool and one that can reason about code like a seasoned analyst, not just match patterns. That's powerful and it's exactly why this announcement is sending ripples through the cybersecurity market," Silva said.

"However, don't mistake this for a plateau," Silva added. "In adversarial environments, capabilities are symmetrical but speed to operationalize is asymmetrical in favor of attackers. The very AI skills defenders laud today will be weaponized by attackers tomorrow to find unpredictable vectors, to pivot at machine speed, to uncover dangers static tools never even dreamed of."

---



## Acumen Cyber Threat Intelligence Digest: Week 8

The **Cyber Threat Intelligence Digest: February 2026 – Week 8** from Acumen Cyber outlines significant vulnerabilities, active threat campaigns, and notable cyber-security developments observed during the week.

### Main Article

This week's threat landscape is very much "patch or perish". Across all major platforms — Google, Apple and Mozilla — we've seen multiple actively exploited vulnerabilities getting emergency fixes. The common thread is simple: attackers are moving quickly, and vendors are patching just as fast. Organisations need to keep pace.

#### **Google Chrome: Actively Exploited CVE-2026-2441**

Google pushed critical patches on **13 February 2026** to fix a **use-after-free vulnerability in the CSS component of Chrome**. Versions prior to **145.0.7632.75** (Windows/macOS) and **144.0.7559.75** (Linux) are affected. Successful exploitation lets attackers run arbitrary code via a malicious webpage — exactly the sort of thing drive-by campaigns love. No actor attribution yet, and Google's keeping exploitation details tight.

#### **Apple: Actively Exploited Zero-Day CVE-2026-20700 (dyld)**

Apple also had a busy week, patching an **actively exploited memory-corruption flaw** affecting the **Dynamic Link Editor (dyld)**. This hits a huge range of devices — modern iPhones, iPads, Macs and even visionOS. The bug allows arbitrary code execution if an attacker can manipulate memory, which makes it a prime target for privilege-escalation chains. Fixes shipped with **iOS/iPadOS/macOS/watchOS/tvOS/visionOS 26.3**.

#### **Mozilla: CVE-2026-2447 in libvpx**

Mozilla patched a **heap buffer overflow** affecting **Firefox, Firefox ESR, and Thunderbird**. While it hasn't been exploited in the wild yet, it enables **remote code execution**. If you're running outdated versions, you'll want to get patched quickly — RCE bugs in browsers rarely stay unexploited for long once public.

#### **Action Point: Remediation Actions**

- Push Chrome updates immediately — prioritise anyone running Linux where the patched version number differs.  
(Patch covers an active exploit.)
- Force-update Apple devices across your estate — especially student-owned or BYOD hardware.  
(The dyld issue is being exploited already.)
- Update Firefox/Thunderbird organisation-wide — even though it isn't actively exploited, this is a high-risk RCE vulnerability.
- Review browser update compliance via your device management platform (Intune/SCCM/JAMF).

**All the best from all of us at HEFESTIS and look out for a new ThreatScape update next week.**

## CISO-Share Office Weekly Newsletter

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

## Table of Contents

<b>CISO-Share Office Weekly Newsletter</b> .....	1
SC3 Daily Threat Summaries and Weekly Report.....	1
Latest from NCSC: Advises UK organisations to take action following conflict in the Middle East .....	2
ClickFix Attack Uses Windows Terminal to Evade Detection .....	4
AI Agent Overload: How to Solve the Workload Identity Crisis.....	5
Middle East Conflict Highlights Cloud Resilience Gaps.....	6
Acumen Cyber Threat Intelligence Digest: Week 10 .....	7

## SC3 Daily Threat Summaries and Weekly Report



Please find SC3's daily threat summaries for this week for those who do not receive this information directly.

## All Daily, Weekly and Monthly Reports:

<https://www.cyberscotland.com/news/sc3-threat-reports/>



SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu  
lletin-9-March-2026lletin-10-March-202lletin-11-March-202lletin-12-March-202lletin-13-March-202l



SC3 - Monthly SC3-Weekly-Vulner  
Ransomware Reportability-Report-9-Mai



National Cyber  
Security Centre  
a part of GCHQ

## **Latest from NCSC: Advises UK organisations to take action following conflict in the Middle East**

**In response to the evolving events in the Middle East, the NCSC is advising that UK organisations review their cyber security posture.**

Future updates will also available via the link below:

### Main Article

How has the cyber threat changed?

- As a result of the ongoing conflict in the Middle East, there is likely no current significant change in the direct cyber threat from Iran to the UK, however due to the fast-evolving nature of the conflict, this assessment may be subject to change.
- There is almost certainly a heightened risk of indirect cyber threat for those organisations and entities who have a presence, or supply chains, in the Middle East
- Iranian state and Iran-linked cyber actors almost certainly currently maintain at least some capability to conduct cyber activity.

How should organisations respond?

- Organisations should prepare to respond to the risk of collateral impacts in the UK from Iran-linked hacktivists by reading previously issued advisories on [DDoS attacks](#), [phishing activity](#) and [ICS Targeting](#).
- For organisations exposed to higher risk, for example with those with offices or supply chains in the region, you should adjust your cyber security posture accordingly. You should take the steps outlined in our [actions to take when threat is heightened guidance](#), and consider proportionate action to [increase monitoring](#) and [review your external attack surface](#).
- The NCSC continues to encourage UK organisations to sign up to its [Early Warning service](#), to receive timely notifications of security issues affecting their networks.
- In addition, given this is an evolving situation, CNI organisations may wish to pre-emptively review the NCSC's recently published guidance on [actions to take now to prepare CNI organisations for severe cyber threat](#).
- For physical and personnel security risks, please refer to [guidance](#) issued by the National Protective Security Authority (NPSA). In particular following the [sabotage guidance](#) will help you protect your site from physical threats.

**Action Point:**

Organisations are advised to review their risk posture, take proportionate action and report any concerning activity to the NCSC's Incident Management team using Report a cyber incident.

[Report a cyber incident](#)

---



## ClickFix Attack Uses Windows Terminal to Evade Detection

Like traditional ClickFix attacks, the campaign relies on fake CAPTCHA pages, troubleshooting prompts, and verification lures to trick victims into executing malicious PowerShell commands.

### Main Article

What sets the new campaign apart, however, is the fact that victims are instructed to open Windows Terminal directly, instead of relying on the Windows Run dialog.

“Rather than the traditional Win + R → paste → execute technique, this campaign instructs targets to use the Windows + X → I shortcut to launch Windows Terminal (wt.exe) directly, guiding users into a privileged command execution environment that blends into legitimate administrative workflows and appears more trustworthy to users,” Microsoft [says](#).

The new approach, observed in the wild in February, allows attackers to bypass protections designed to prevent Run dialog abuse, the tech giant notes.

The execution of the malicious command in Windows Terminal spawns a PowerShell process that decodes embedded hex commands, triggering a multi-stage attack chain that leads to a Lumma Stealer infection.

**Action Point:** The code achieves persistence using scheduled tasks, contains anti-malware evasion routines, and targets browser data and other sensitive information for exfiltration.

In another variant of the attack, the malicious commands executed in Windows Terminal lead to a batch script executed via command prompt and through MSBuild.exe.

“The script connects to Crypto Blockchain RPC endpoints, indicating etherhiding technique. It also performs QueueUserAPC()-based code injection into chrome.exe and msedge.exe processes to harvest Web Data and Login Data,” Microsoft says.

Another recently observed ClickFix attack variant, dubbed [InstallFix](#), relies on cloned AI tool websites to trick victims into executing malicious commands, also leading to information-stealer infections.





## AI Agent Overload: How to Solve the Workload Identity Crisis

Authenticating workloads is becoming more and more complex, particularly given things like AI agents and the wide range of identity permissions they need. Organizations need to be thinking ahead on securing workloads in complicated modern environments,

but it's not an easy task.

### Main Article

Researchers at Zscaler hope to explore this evolution in an upcoming RSAC 2026 Conference session entitled, "[What Are You, Really? Authenticating Workloads in a Zero Trust World.](#)"

In computing terms, workloads cover the tasks applications and services conduct in order to do their job, and the IT resources those tasks consume. Workloads can refer to a wide range of things, from processing front-end user requests on a Web server (like managing a shopping cart) to cloud-native microservices, complex data analysis, AI training, and more.

**Action Point:** Curry tells Dark Reading that, from a defender standpoint, there are many options to solve these problems and remediate the weaknesses. At a basic level, he says organizations should be looking for secrets, taking inventory of AI agents (as well as other NHI processes and services), adopting standards, and working toward zero-trust. They should also be talking to their platform providers about also adopting workload authentication standards.

---



## Middle East Conflict Highlights Cloud Resilience Gaps

Businesses that counted on the cloud's distributed nature to guarantee their data's availability have had a cold dose of reality during the past two weeks.

### [Main Article](#)

On Feb. 28, following military strikes by the US and Israel, Iran's Internet traffic [fell to less than 1% across all major networks in the country](#), according to Cloudflare Radar, which tracks Internet traffic internationally. Within 24 hours, Iran responded, targeting infrastructure in the United Arab Emirates, Bahrain, and other Gulf States, hitting two Amazon Web Services' facilities in the UAE with drone strikes, while a third facility in Bahrain suffered "physical impacts to [its] infrastructure," Amazon Web Services stated March 2 on its AWS Health Dashboard.

"These strikes have caused structural damage, disrupted power delivery to our infrastructure, and in some cases required fire suppression activities that resulted in additional water damage," [AWS stated](#). "We are working closely with local authorities and prioritizing the safety of our personnel throughout our recovery efforts."

**Action Point:** One of the first casualties of the attacks on cloud infrastructure may be the push by many countries to keep data and digital services in their borders.

"Forcing a country's data to remain trapped within its physical borders turns it into a massive strategic liability that can be erased in a single bombing campaign," Raines says. "As a result, I expect we'll see governments rapidly shift toward 'Allied Data Sovereignty,' and rewrite laws to ensure critical national data can be legally backed up and hosted in allied nations during a crisis to ensure it survives."

Companies that rely on their cloud platforms — and most do — need to rethink their disaster recovery and data governance, she says.



## Acumen Cyber Threat Intelligence Digest: Week 10

**The Cyber Threat Intelligence Digest: March 11 2026 – Week 10 from Acumen Cyber outlines significant vulnerabilities, active threat campaigns, and notable cyber-security developments observed during the week.**

### Main Article

A major focus of the digest is a sophisticated adversary-in-the-middle (AitM) phishing campaign targeting Amazon Web Services (AWS) Management Console credentials. Researchers reported that attackers sent emails disguised as AWS security alerts, warning recipients about suspicious cross-account activity involving Identity and Access Management roles. Victims who followed the embedded links were redirected through multiple domains before arriving at a counterfeit AWS login page designed to closely replicate the legitimate interface. The phishing kit intercepted credentials in real time by relaying authentication requests to the genuine AWS service and returning responses to the victim, enabling attackers to bypass standard authentication workflows. In at least one case, compromised credentials were used within roughly twenty minutes to access an AWS account via a VPN exit node, demonstrating the attackers' rapid operational tempo.

The digest also describes a malvertising campaign distributing the Amatera Stealer malware through a technique known as "InstallFix." Attackers created cloned installation pages for a developer tool and promoted them through paid search advertisements. The fraudulent sites presented installation commands that appeared legitimate but actually retrieved malicious payloads from attacker-controlled infrastructure. Execution methods varied by operating system: on Windows the command invoked mshta.exe to run a remote script, while on macOS it decoded and executed a compressed payload via the Z shell before downloading additional malware.

Additionally, the report notes malware capabilities observed in broader threat activity, including environment validation checks, anti-analysis techniques, and attempts to disable security features. Some samples altered Windows Defender registry settings to reduce protection, downloaded further payloads from command-and-control servers, and used techniques such as process hollowing or vulnerable drivers to achieve elevated privileges and terminate security tools.

Overall, the digest emphasises the continued evolution of phishing infrastructure, social-engineering-driven malware distribution, and increasingly sophisticated post-exploitation techniques used to evade detection and maintain persistence in compromised environments.

### **Action Point:** Remediation Actions

- CVE-2026-1492 (WordPress) – This vulnerability can be addressed by updating the User Registration and Membership plugin to version 5.1.3 or later.
- CVE-2026-29000 (pac4j-jwt) – This vulnerability can be remediated by updating to versions 4.5.9, 5.7.9, and 6.3.3.
- CVE-2025-36105 (IBM Containers) – This vulnerability can be addressed by updating IBM Planning Analytics Advanced Certified Containers to version 3.1.5.

**All the best from all of us at HEFESTIS and look out for a new ThreatScape update next week.**



### CISO Share

#### Threatscape

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).



#### [Iran intelligence backdoored US bank, airport, software outfit networks](#)

Security researchers have uncovered an active Iranian cyber-espionage campaign that has been sitting inside multiple U.S. and Canadian organisations—as well as an Israeli site—since early February 2026. The activity is linked to MuddyWater (aka Seed...

Steve McIntosh

3/6/2026



#### [Defending Against Iranian Cyber Threats in the Wake of Operation Epic Fury](#)

The article outlines the cyber fallout following Operation Epic Fury, a coordinated U.S. and Israeli military and cyber operation launched on 28 February 2026 against

Iranian military and government assets.  
Cyber Command led with digital attacks be...

Steve McIntosh

3/6/2026

4 views

---



### [SC3 Daily Threat Summaries and Weekly Report](#)

Action Point: Please find SC3's daily threat summaries for this week for those who do not receive this information directly.

Steve McIntosh

3/6/2026

2 views

---



### [Cisco Confirms Active Exploitation of Two Catalyst SD-WAN Manager Vulnerabilities](#)


Cisco has confirmed that attackers are actively exploiting two vulnerabilities in Catalyst SD-WAN Manager—formerly vManage—and organisations running older or unpatched versions should assume they're exposed. The two flaws under active exploitation ...

Steve McIntosh

3/6/2026

2 views

---



## [Microsoft Reveals ClickFix Campaign Using Windows Terminal to Deploy Lumma Stealer](#)

Microsoft has dropped details on a new social-engineering campaign called ClickFix , and it's one worth paying attention to because it sidesteps a lot of the usual detection logic we rely on. Instead of telling victims to open the classic Run dialog...

Steve McIntosh

3/6/2026

5 views



## [Backup strategies are working, and ransomware gangs are responding with data theft](#)

Cyber insurance data from Coalition paints a pretty familiar picture: BEC, fraud, and ransomware are still the heavy hitters – but the tactics and losses are shifting. Business Email Compromise (BEC) was the number-one claim in 2025, making up 31% o...

Steve McIntosh

3/6/2026

6 views

[Go To Site](#)

## CISO-Share Office Weekly Newsletter

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are also a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

---

### SC3 Daily Threat Summaries and Weekly Report

<https://www.cyberscotland.com/news/daily-threat-reports/>



A daily breakdown of various cyber threats from the Scottish Cyber Coordination Centre (SC3)

#### Action Point:

All SC3 threat intelligence in one place.



### 54 EDR Killers Use BYOVD to Exploit 35 Signed Vulnerable Drivers and Disable Security

<https://thehackernews.com/2026/03/54-edr-killers-use-byovd-to-exploit-34.html>

A new analysis from ESET shows just how widespread **EDR killers** have become in modern ransomware operations. The researchers identified **54 different EDR-disabling tools**, and more than half rely on the increasingly common **Bring Your Own Vulnerable Driver (BYOVD)** technique. In total, attackers are abusing **35 different legitimate but vulnerable drivers**, many of them signed by reputable vendors—meaning Windows trusts them, even though attackers weaponise them.

Why does this matter? Because ransomware actors know that encryptors themselves are noisy and easy to detect. Rather than trying to make the ransomware stealthy, groups instead run a



separate “EDR killer” component **before** deploying the payload. These tools disable endpoint protection, kill processes, tamper with kernel callbacks, and neuter defensive agents, giving threat actors moments—or minutes—of complete freedom.

BYOVD is the most popular method because it’s **reliable and simple**. Attackers load an old, signed driver with a known vulnerability. Since the driver is legitimate, Windows allows it to run, giving the attacker *Ring 0* (kernel-level) access. With that, they can turn off almost any security control, regardless of vendor.

ESET highlights four major categories of EDR killers:

1. **BYOVD-based binaries** — Used heavily by closed ransomware groups like DeadLock and Warlock, as well as cybercriminals selling “EDR-killer-as-a-service.”
2. **Script-based tools** — Abuse built-in commands like taskkill, net stop, and sc delete. Some even force reboots into **Safe Mode**, where security tools don’t load—but this is noisy and risky, so less common.
3. **Anti-rootkit utilities** — Legitimate tools (e.g., GMER, PC Hunter) weaponised to terminate protected processes.
4. **Driverless EDR killers** — An emerging class that simply cuts off outbound traffic from EDR agents, putting them into a “coma” without touching drivers.

The article notes attackers **aren’t trying to hide encryptors anymore**—they’re putting all their innovation into the pre-encryption EDR-killer stage. These tools are cheap, fast to adapt, and easily reused across campaigns.

Defending against this requires **layered security**, because blocking a single EDR killer won’t help—attackers can just switch to another. The goal is detecting and disrupting earlier phases of the intrusion before the EDR killer even launches.

## Action Points

### 1. Block known vulnerable drivers

Use Windows Defender Application Control (WDAC), kernel blocklists, and Microsoft’s vulnerable driver list.

### 2. Monitor for BYOVD behaviour

Alerts for unexpected driver loads, especially from non-standard paths or old vendor signatures.

### 3. Detect pre-ransomware behaviours

EDR killers appear *late* in the kill chain—focus on earlier indicators like credential theft, lateral movement, or LSASS access.

### 4. Protect Safe Mode

Disable unauthorised Safe Mode reboots and monitor for services being tampered with.

## 5. Monitor EDR health signals

Alert if the EDR agent stops checking in or suddenly drops telemetry.

## 6. Adopt defence-in-depth

EDR alone cannot stop BYOVD attacks—identity protection, network segmentation, and logging are critical.



### Shadow AI Risk: How SaaS Apps Are Quietly Enabling Massive Breaches

<https://www.securityweek.com/the-shadow-ai-problem-how-saas-apps-are-quietly-enabling-massive-breaches/>

**A new report from Grip Security analyzes 23,000 SaaS application environments. The statistics it found include:**

**100% of analyzed companies operate SaaS environments with embedded AI; there has been a year-over-year 490% spike in public SaaS attacks; and 80% of documented incidents involve PII and/or customer data.**

“But what really surprised me,” says Chad Holmes, product marketing consultant at Grip Security, “is that organizations have an average of 140 AI-enabled SaaS environments.” If an AI-enabled app is breached, any integral agentic AI can be used first to access data from connected systems, secondly to cascade from the one breach to a breach of every other AI-enabled environment within the organization – and potentially to expand further into AI-enabled environments in other organizations. This is chaos.

The poster boy example of this cascading chaos is the Salesloft Drift incident (the ‘Great SaaS Breach of 2025’). Ultimately more than 700 organizations were affected, including [security firms](#) Cloudflare, Palo Alto Networks, Zscaler and CyberArk. [UNC6395](#) attackers compromised Salesloft’s internal systems, starting with their [GitHub](#) repositories and moving from there into the Drift AWS environment. Here they stole the active OAuth and refresh tokens used by customers to connect the Drift Chatbot to local installations of Salesforce and other apps such as Slack.

Armed with the legitimate pre-approved OAuth token, the attackers were able to impersonate Drift and log directly into Salesforce installations into companies also using the Drift chatbot. One breach of a SaaS app (Drift) cascaded into hundreds of compromises in different companies across the globe.


Now it is fair to say that this incident prompted a drive to improve the security of SaaS apps and their AI implementations, but Grip is not convinced it will be enough. “One thing we’re seeing,” comments Holmes, “as we’ve moved outside the traditional perimeter, outside firewalls and network level protections, [identity is the new perimeter](#). The focus is on identity, and if we have that identity, we can log into any environment anywhere.” For attacks against SaaS AI, the key ‘identity’ is a valid OAuth token.

It's worth briefly examining why this is such a threat, if control hasn't been achieved. Much of it is down to the current need for speed in business. SaaS developers are tempted to rapidly include agentic AI within their own products to improve efficiency ahead of their competition, and they don't always make the implications apparent to their customers. The customer may have installed shadow AI without knowledge. 'Shadow' is the epithet used to describe the use of AI or autonomous agentic AI without formal oversight from the IT and security departments – and if the customer is unaware of the AI within the SaaS app, it is automatically 'shadow'.

## Action Point:

But the chaos can be brought under control. "The way out is not more policy or slower innovation. It is a shift in how AI is governed in practice," adds the report. The key is increased visibility into, and understanding of, SaaS shadow AI, and more dynamic governance.

"Leaders who succeed replace static approvals with continuous oversight, discovery, and risk-based controls. AI becomes a managed third-party risk, monitored continuously, aligned to business outcomes, and governed with the same rigor as any critical supplier."



## Iranian Hackers Likely Used Malware-Stolen Credentials in Stryker Breach

<https://www.securityweek.com/iranian-hackers-likely-used-malware-stolen-credentials-in-stryker-breach/>

**Newly uncovered evidence suggests that the recent cyberattack targeting US medical technology giant Stryker involved compromised credentials obtained via infostealer malware.**

The [attack on Stryker](#), a major manufacturer of surgical equipment and orthopedic implants for hospitals worldwide, came to light on March 11, with the Iran-linked hacker group Handala immediately taking credit.

Handala, which is believed to be an anti-Israel hacktivist persona under the control of Iran's Ministry of Intelligence and Security (MOIS), claimed to have wiped more than 200,000 devices, forcing Stryker to shut down offices in dozens of countries. The hackers also claimed to have stolen a significant amount of data.

While some early reports indicated that the hackers used wiper malware in the attack — Handala has been known to use such malware — Stryker said it found no evidence of malware being deployed on its systems.

According to some reports, the attackers wiped systems by abusing Stryker's Microsoft Intune instance, which is used to remotely manage desktop and mobile endpoints and applications within the organization.

Bleeping Computer [reported](#) earlier this week that the attackers compromised an Intune administrator account and created a new global admin account, which they used to wipe managed devices.

## Action Point:

While pro-Iran hackers have [ramped up attacks](#) against Israel, the US, and other allies after the war began, this appears to be the most significant attack against the United States. Handala has been highly active since the start of the conflict, particularly against Israel, claiming to have hacked a wide range of organizations. However, its claims are often difficult to fully verify.

Forbes reported on Tuesday that two leaders of Iranian cyber operations have been [killed](#) in the recent airstrikes. One of them is Mohammad Mehdi Farhadi Ramin, [charged](#) by the US in 2020 for his role in state-sponsored hacking, and Yahya Hosseiny Panjaki, who oversaw the MOIS unit that controlled hacker groups such as Handala.



## Microsoft New tools and guidance: Announcing Zero Trust for AI

<https://www.microsoft.com/en-us/security/blog/2026/03/19/new-tools-and-guidance-announcing-zero-trust-for-ai/>

Microsoft has announced a major evolution of its security guidance: **Zero Trust for AI (ZT4AI)**—an extension of traditional Zero Trust principles designed specifically for AI systems, agents, and the data that powers them. The message from Microsoft is simple: AI introduces new trust boundaries and risks, and trying to secure it with the same old models won't cut it.

Security leaders are consistently asking Microsoft: *"We're adopting AI fast—how do we secure it just as fast?"* This release is their attempt to answer that question with updated tools, architecture, and prescriptive guidance.

At the centre of the announcement is a refreshed **Zero Trust Workshop**, now with a dedicated **AI pillar** covering **700 controls**, 116 logical groups, and 33 "swim lanes." It tackles the full lifecycle: how organisations handle AI identities, control access to prompts/models, protect sensitive data flowing through AI, and govern behaviour of agentic systems.

Alongside that comes an expanded **Zero Trust Assessment tool**, adding **Data** and **Network** as formal pillars (Identity and Devices were already there). This matters because AI agents rely heavily on vast data and network reach—two areas where missteps can lead to prompt injection, data leakage, or agents acting outside intended

boundaries. A dedicated AI pillar for the Assessment tool is due **summer 2026**.

Microsoft also released a **Zero Trust for AI reference architecture**, showing how least-privilege, continuous verification, and layered monitoring operate across AI workloads—from training to inference to agent actions. It's designed to give security, IT, and engineering teams a shared blueprint, particularly as organisations roll out semi-autonomous or autonomous agents.

A useful part of the drop is the set of **practical AI security patterns**. These include threat modelling techniques built for AI, guidance for securing agentic systems, AI observability requirements, safety engineering principles, and defences for indirect prompt injection—something standard mitigations often fail to stop.

All of this is framed as a structured journey: *strategy* → *assessment* → *implementation*. Microsoft's goal is to give security teams a clear path instead of abstract high-level advice.

The launch aligns with several RSA Conference 2026 sessions Microsoft is hosting focused on AI security, risk, and Zero Trust maturity.

## Action Points

### 1. Use Microsoft's Zero Trust Workshop

Run the updated version with the AI pillar to baseline strengths and weaknesses.

### 2. Run the expanded Zero Trust Assessment

Evaluate identity, devices, data, and network controls—especially before deploying internal AI agents.

### 3. Apply Zero Trust to AI workloads


- Verify AI agent identities
- Apply least privilege to prompts, plugins, and data
- Assume breach to defend against prompt injection & poisoning

### 4. Use Microsoft's AI security patterns

Adopt AI-specific threat modelling and observability standards.

### 5. Plan for autonomous agent governance

Establish guardrails, monitoring, and policies before agents act across systems.



## Apple Warns Older iPhones Vulnerable to Coruna, DarkSword Exploit Kit Attacks

<https://thehackernews.com/2026/03/apple-warns-older-iphones-vulnerable-to.html>

Apple has issued a warning to anyone still running **older, unpatched versions of iOS**, after researchers confirmed that attackers are using powerful exploit kits—named **Coruna** and **DarkSword**—to compromise outdated iPhones via malicious websites. These aren't theoretical risks: Apple says simply **clicking a malicious link or visiting a compromised webpage** could let attackers steal sensitive data from devices that aren't up to date.

These exploit kits are being used in **web-based “drive-by” attacks**, making them easy for threat actors to weaponise at scale. Historically, this level of sophistication was mostly seen in state-sponsored mobile spyware, targeting diplomats, journalists and high-value individuals. The worrying development, according to security firm iVerify, is that these high-end iOS exploits are now in the **secondary market**, meaning less-skilled cybercriminals can use them too.

The attacks rely on vulnerabilities affecting older iOS versions—issues Apple has already patched. Devices running **iOS 15 through iOS 26** are safe if fully updated, as those versions contain fixes for the exploited flaws. But millions of older devices still run outdated versions due to age, storage constraints, user habits, or lack of awareness. Apple's message is blunt: *“Keeping your software up to date is the single most important thing you can do to maintain the security of your Apple products.”*

For users stuck on older devices that **can't** upgrade to the newest iOS versions, Apple has released special updates such as **iOS 15.8.7, iPadOS 15.8.7, iOS 16.7.15, and iPadOS 16.7.15**. More “critical security updates” are expected to roll out in the coming days for devices still on iOS 13 or 14.

Apple also advises enabling **Lockdown Mode** where available. It's designed for high-risk users, but in this case it can provide meaningful protection by blocking risky web content and reducing exposure to malicious code execution paths.

Researchers warn that these exploit kits have been adopted quickly by **multiple threat actors in different countries**, and the simplicity of deploying them makes widespread mobile compromise a realistic concern. The takeaway: if the device isn't patched, it's at real risk.

## Action Points

### 1. Update immediately

- Newer devices: just install the latest iOS update.
- Older devices: install iOS/iPadOS **15.8.7** or **16.7.15**.

### 2. For legacy devices stuck on iOS 13/14:

- Update to **iOS 15**, and expect a critical patch shortly.

### 3. Enable Lockdown Mode

Especially for high-risk users, execs, researchers, and anyone targeted in past phishing/smishing campaigns.

### 4. Reinforce staff awareness

Drive-by browser exploits mean **malicious links are enough**—no installs required.





## DoJ Disrupts 3 Million-Device IoT Botnets Behind Record 31.4 Tbps Global DDoS Attacks

<https://thehackernews.com/2026/03/doj-disrupts-3-million-device-iot.html>

The U.S. Department of Justice has taken down the command-and-control infrastructure behind **four major IoT botnets**—AISURU, Kimwolf, JackSkid, and Mossad—responsible for some of the **largest DDoS attacks ever recorded**, including the staggering **31.4 Tbps attack** Cloudflare reported in November 2025. This was a coordinated, court-authorised, multi-country effort involving the U.S., Canada, Germany, and a long list of private-sector partners like AWS, Cloudflare, Akamai, Google, Oracle, Lumen, and more.

These botnets weren't small operations. Together they had infected **over 3 million devices**, mostly off-brand Android TVs, webcams, routers, and DVRs—classic IoT kit with weak security. AISURU and Kimwolf alone were responsible for hundreds of thousands of attacks, and Kimwolf in particular represented a significant step forward in botnet evolution: instead of scanning the internet for vulnerable devices, it infiltrated **residential proxy networks**, effectively using home routers and media boxes to hide malicious traffic behind legitimate household IP addresses. That made them harder to detect and far more resilient.

The operators rented out their botnets through “**cybercrime-as-a-service**,” letting other criminals buy access to launch DDoS attacks. Some attacks hit 30 Tbps and 14 billion packets per second—levels capable of overwhelming even major ISPs and cloud providers.

Researchers also identified individuals they believe were involved: one potential operator in Canada (though he claims impersonation) and another in Germany. No arrests have been announced yet.

Lumen's Black Lotus Labs said it null-routed **almost 1,000 C2 servers**, and noted that copycat botnets have sprung up rapidly because so many IoT devices remain vulnerable. JackSkid and Mossad were still hitting **100,000–250,000 devices per day** in early March 2026.

Akamai warned that these hyper-volumetric attacks can cripple core internet infrastructure and even overwhelm top-tier DDoS protection systems. The risk isn't just large organisations—anyone depending on internet services can feel the impact if a backbone provider gets hammered.

### Action Points

#### Audit your IoT estate

Identify cameras, TVs, printers, smart boards, sensors, etc. Replace or isolate anything unpatchable.

#### Enforce network segmentation

IoT should *never* sit flat on production, student, or corporate networks.



## **Block outbound traffic anomalies**

Watch for spikes in packets per second or unusual destinations—classic botnet indicators.

## **Patch and harden edge devices**

Routers, firewalls, and Wi-Fi systems are favourite targets for these botnets.

## **Use DDoS-resilient architecture**

If you're running public-facing services, review DDoS controls with providers.

---

All the best from all of us at **HEFESTIS** and look out for our next **Threatscape** update next week.

*HEFESTIS Limited, Registered Office: Unit 27, Stirling Business Centre, Wellgreen, Stirling FK8 2DZ*

*Incorporated in Scotland SC603511*



## CISO Share

### ThreatScape Newsletter

Dear all,

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

Kind Regards

David



### [SC3 Daily Threat Summaries and Weekly Report](#)

Please find SC3's daily threat summaries for this week for those who do not receive this information directly. Scottish Cyber Activity Report 2026 - gov.scot CyberScotland updates News from CyberScotland and our Partner network FutureScot Cyber Secu...

David Robertson

3/27/2026

9 views



### [Changes to Cyber Essentials for April 2026](#)

The Cyber Essentials scheme is updated annually to stay aligned with evolving threats. While the scheme's five core controls remain unchanged, the April 2026 updates aim to enhance clarity, consistency, and effectiveness. This blog outlines the ann...

David Robertson

3/24/2026

16 views




### [March 2026 Information and Cyber Security Update](#)

Threat Update March 2026 The cyber risk environment facing UK universities, colleges, and public-sector organisations is becoming increasingly difficult as financial pressures, resource scarcity and rapid digital innovation requirements collide. Att...

David Robertson

3/26/2026

12 views



### [SANS: Top 5 Most Dangerous New Attack Techniques to Watch](#)

RSAC 2026 CONFERENCE – San Francisco – Each year SANS researchers head to the RSAC Conference to reveal the five top attack techniques. But 2026 marks a distinct

shift: all are powered by artificial intelligence.  
"We would be lying to you if we poin...

David Robertson

3/26/2026

8 views

---



### [Chinese Hackers Caught Deep Within Telecom Backbone Infrastructure](#)

A China-linked state-sponsored threat actor has deployed kernel implants and passive backdoors deep within telecommunication backbone infrastructure worldwide for long-term persistence, Rapid7 reports. Main Article The stealth digital sleeper cells ...

David Robertson

3/27/2026

8 views

---



### [Stryker Says Malicious File Found During Probe Into Iran-Linked Attack](#)

Medical technology giant Stryker has shared an update regarding its investigation into the recent Iran-linked cyberattack, revealing that a malicious file used by the attackers has been identified. Main Article The Stryker incident came to light on ...

David Robertson

3/26/2026

12 views

---



## [Cyber Threat Intelligence Digest: Week 12](#)

Highlights of Cyber Threat Intelligence Digest  
Main Article The digest outlines several newly disclosed vulnerabilities, active threat campaigns, and broader security developments across software, AI tooling, and online services. Citrix reported two...

David Robertson

3/26/2026

7 views

[Go To Site](#)

## CISO-Share Office Weekly Newsletter

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

## Table of Contents

<b>CISO-Share Office Weekly Newsletter</b> .....	1
SC3 Daily Threat Summaries and Weekly Report .....	1
SANS: Top 5 Most Dangerous New Attack Techniques to Watch .....	2
March 2026 Information and Cyber Security Update .....	4
Changes to Cyber Essentials for April 2026 .....	7
Chinese Hackers Caught Deep Within Telecom Backbone Infrastructure (Technical).....	10
Stryker Says Malicious File Found During Probe Into Iran-Linked Attack.....	13
Acumen Cyber Threat Intelligence Digest: Week 12 .....	14

## SC3 Daily Threat Summaries and Weekly Report



Please find SC3's daily threat summaries for this week for those who do not receive this information directly.

## All Daily, Weekly and Monthly Reports:

<https://www.cyberscotland.com/news/sc3-threat-reports/>



SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu SC3-Daily-threat-bu  
lletin-23-March-2026lletin-24-March-2026lletin-25-March-2026lletin-26-March-2026lletin-27-March-2026



SC3-Weekly-Vulnerability-Report-23-March-2026



## **SANS: Top 5 Most Dangerous New Attack Techniques to Watch**

**RSAC 2026 CONFERENCE – San Francisco – Each year SANS researchers head to the RSAC Conference to reveal the five top attack techniques. But 2026 marks a distinct shift: all are powered by artificial intelligence.**

"We would be lying to you if we pointed out a trend in attacks that did not involve AI," SANS president and presentation moderator Ed Skoudis explained to the audience during a keynote session covering the Top 5. "That is just where we are in the industry."

### Main Article

Surprisingly the article contains the Top 5 most dangerous attacks....

First, authorization sprawl in cloud and SaaS environments creates excessive and poorly tracked user permissions. Attackers exploit these hidden access paths to move laterally without detection, highlighting weaknesses in identity and access management.

Second, ransomware targeting industrial control systems (ICS) poses growing risks to critical infrastructure. As organizations automate operations, they often remove manual backups, creating single points of failure that attackers can exploit to disrupt essential services.

Third, destructive ICS attacks go beyond financial motives, aiming to cause physical damage or safety failures. These attacks, often linked to nation-state actors, manipulate industrial processes and safety mechanisms, raising concerns about real-world consequences.

Fourth, attackers are increasingly erasing or avoiding forensic artifacts, making incident response and investigation significantly harder. Without sufficient logging and visibility, defenders struggle to determine how breaches occurred or how far they spread.

Finally, AI-related regulatory challenges introduce a new type of risk. While defenders rely on AI for detection, emerging privacy laws may restrict its use, potentially giving adversaries—who freely exploit AI—an advantage.

Overall, the report stresses that cybersecurity is now a strategic business issue, requiring improved visibility, stronger governance, and cross-functional coordination. Organizations must adapt quickly to these evolving threats to maintain resilience in increasingly complex digital environments.



## Action Point: So how can defenders respond?

"They have their artificial intelligence," Lee said. "Now we build ours."

He pointed to [Protocol SIFT](#), an open source initiative from SANS Institute designed to help defenders catch up with AI-wielding attackers. It uses AI to organize workflows, surface insights, and coordinate tools. Meanwhile, humans are responsible for validating results and making decisions.

"The goal is to accelerate analysts, not replace them, and early results suggest that the model can significantly compress response times," Lee said.

In one response exercise involving a sophisticated, two-week attack scenario, an analyst used Protocol SIFT to wrap up the entire investigation in a little less than 15 minutes, including identifying the malware, mapping the attacker's movements, and aligning the tactics, techniques, and procedures (TTP) activity to known frameworks, and determining next steps. It's the ability for defenders to move react quickly and coordinate across the global security community that will give defenders a true edge over attackers, Lee added.

---



## March 2026 Information and Cyber Security Update

### Threat Update March 2026

The cyber risk environment facing UK universities, colleges, and public-sector organisations is becoming increasingly difficult as financial pressures, resource scarcity and rapid digital innovation requirements collide. Attackers are taking advantage of these conditions, exploiting stretched budgets and system vulnerabilities across campuses. The education sector continues to be a prime target because it holds large volumes of personal data and valuable research.

Ransomware and data-extortion attacks remain the most damaging form of cybercrime in 2026. Modern groups now do more than lock systems—they steal data first and then threaten to publish or sell it if a ransom isn't paid. Some also contact regulators or the press to increase pressure. Supply chains have become a major weakness, with single provider breaches (for example, student management or learning platforms) having the potential to affect dozens of institutions at once. Attacks on shared IT services and outsourced support arrangements are proving particularly disruptive.

Third-party risks have intensified. Criminals and state actors are using convincing social-engineering techniques and bypassing identity verification controls to gain trusted access to systems shared across education and public services. Once inside a supplier platform, attackers can quietly move through networks for weeks or months before discovery. This connected ecosystem—where many institutions depend on common services—creates a “multiplier effect” when one compromise occurs.

The rapid adoption of artificial intelligence tools has introduced new threats. Attackers now use AI to automate research about staff, craft believable phishing emails, and identify weak points faster than before. Fake study assistants and generative-AI chatbots, some posing as legitimate platforms such as “Einstein,” are being abused to send credible messages that mimic lecturers or students. At the same time, educators and learners are unintentionally leaking sensitive data by feeding it into external AI tools not controlled by the organisation.

National cyber bodies, including the NCSC and the Scottish Cyber Coordination Centre, report a continuing rise in attacks against UK and Scottish education and public sector networks. Typical breaches involve stolen or weak passwords, compromise of cloud-based collaboration systems, or techniques that defeat multi-factor authentication. Scottish institutions are considered particularly exposed because of some remaining legacy infrastructure, shared services, and perception of limited in-house cybersecurity capability. Moreover, some threat actors are focusing on research linked to health, energy, and emerging technology—areas critical to the UK's economic and national security agendas. In these cases, the goal is often stealthy, long-term access rather than obvious disruption.

Within this context, many of our members are actively improving their cyber defences. Efforts have focused on around-the-clock monitoring through upgraded Security Operations and SIEM services, tightening control of privileged accounts and supplier access, and strengthening network segmentation to stop malware spreading internally. There is increased emphasis on reliable backup testing and wider organisational awareness. Staff and students are now receiving more consistent, mandatory training centred on daily risks and behaviours, along with clearer policies for non-university devices (BYOD), remote access, and third-party connections. Information and cyber risk management is being embedded at operational and leadership levels rather than sitting solely within IT teams.

**Recent geopolitical tensions have further raised the urgency. Conflicts in the Middle East** are driving global volatility across both cyberspace and supply chains. Intelligence shows a pattern of state-linked and activist groups expanding their operations well beyond the immediate conflict zone. Cyber campaigns have moved from opportunistic disruption to targeted, coordinated attacks on critical infrastructure, research, and cloud environments. Even large cloud providers have experienced regional disruption—highlighting the link between physical conflict and digital resilience.

For any UK institutions with operations or collaborations in the Middle East, this means exposure to potential internet outages, degraded cloud services, and increased surveillance or espionage. Stricter local data-sovereignty laws may also make incident response or cross-border recovery more challenging. Therefore, institutions are advised to validate local backups, review data flows that cross regions, and keep monitoring at a higher level for targeted or spill-over attacks from regional tensions.

Global supply chain disruption has become another major concern. The same conflict is affecting energy markets, semiconductor availability, and logistics chains, which in turn impact the cost and delivery times of essential security hardware like firewalls or authentication tokens. Reduced redundancy in cloud and SaaS ecosystems means fewer failover options if something goes wrong. Cybersecurity managers should plan for supplier degradation or failure, require stronger contractual assurances for incident notifications, and diversify sourcing strategies where possible. This ties directly to the UK NCSC priority on supply-chain assurance.

Closer to home, the Scottish public-sector cyber environment is experiencing sustained pressure. Incident reports show increasing compromise attempts against universities, colleges, public sector and other shared-service participants. Attackers are focusing on identity systems, cloud storage, and collaboration platforms—the very technologies underpinning hybrid teaching and administration. AI-crafted phishing attempts are growing more accurate and harder to detect, while supply-chain attacks targeting education vendors are becoming more frequent. The latest intelligence supports ongoing improvement priorities for Identity and Access Management, Data Security, and continuous Security Monitoring. These are now urgent, not theoretical, risks for the sector.

A further emerging threat is the misuse or careless use of AI tools within institutions. As staff and students interact with public or semi-trusted AI platforms, there is an increasing risk of accidental data leakage, exposure of login information, and creation of “shadow AI” systems operated outside compliance boundaries. Attackers take advantage of this by generating synthetic identities or using AI to automate reconnaissance. Policies must now clarify which AI tools are allowed, how sensitive data is classified before use, and how monitoring can detect unusual outbound flows. Awareness campaigns should explain in plain language what responsible AI use looks like and what kinds of data should never be pasted into public models.

Finally, institutions with overseas campuses or operations—especially in politically volatile regions—should consider cyber resilience a duty-of-care issue as much as a technical one. Regional communications outages, hacktivist activity, and increased surveillance can directly affect staff safety and business continuity. Oversight of these operations must explicitly include cyber-physical coordination and secure communications planning.

**Action Point:** Taken together, all these developments make clear that the cybersecurity environment for the UK HE and public sectors in 2026 is both dynamic and highly interlinked with global events. The immediate recommendation is not to change overall strategy, but to heighten awareness at Board and leadership levels and to ensure governance keeps pace with the shifting risk backdrop. Improvement programmes underway—covering supply chain, identity, cloud dependency, and incident response—remain appropriate. However, their implementation now carries enhanced urgency, with additional scrutiny advised on supplier resilience, data-sovereignty constraints, and regional exposure for collaborative or international operations.



## Changes to Cyber Essentials for April 2026

The Cyber Essentials scheme is updated annually to stay aligned with evolving threats. While the scheme's five core controls remain unchanged, the April 2026 updates aim to enhance clarity, consistency, and effectiveness. This blog outlines the annual updates to the [Requirements for IT Infrastructure](#) document, which serves as the standard for achieving Cyber Essentials certification. It also contains essential new information about changes to the assessment framework.

[Main Article](#)

[Cyber Essentials Supply Chain Playbook](#)

### A summary for HE/FE is included below

While you may have already seen a preview of the updated Requirements, the [National Cyber Security Centre](#) (NCSC) has now added further adjustments to the certification process, marking scheme, and Cyber Essentials Plus assessment methodology. It is important to understand and implement these changes to ensure compliance with the updated requirements.

### Cyber Essentials 2026 Update (v3.3 "Danzell")

Summary & Implications for Higher and Further Education

#### Effective Date

- Applies to certifications started from 27 April 2026
- No change to the 5 core control areas, but significantly tighter enforcement and clarity

#### Key Changes

##### 1) Mandatory MFA Enforcement

- Multi-factor authentication (MFA) required wherever available
- Failure to implement MFA = automatic certification failure
- Applies especially to:
  - o Cloud services (M365, Google Workspace)
  - o Admin accounts
  - o Remote access

##### 2) Cloud Services Fully In Scope

- Formal definition of cloud services introduced
- Cannot be excluded from scope
- Covers:
  - o Email platforms
  - o Identity providers
  - o SaaS applications (e.g. VLEs, MIS systems)

## 3) Stricter Patching Requirements

- Critical/high vulnerabilities must be remediated within 14 days
- Introduction of auto-fail conditions for non-compliance

## 4) Stronger Scoping Rules

- Clearer definitions of in-scope systems:
  - Internet-facing systems
  - Devices accessing organisational data/services
  - Reduced flexibility to exclude systems without justification

## 5) Enhanced Cyber Essentials Plus (CE+)

- Increased technical validation (less reliance on self-declaration)
- Broader sampling and stricter testing methodology
- Reduced opportunity for pre-audit remediation

## 6) Identity & Access Management Focus

- Greater emphasis on:
  - MFA coverage
  - Access control policies
  - Emerging passwordless approaches

## 7) Backup & Recovery Expectations

- Backups must be:
  - Defined
  - Tested
  - Proven recoverable

## 8) Governance & Accountability

- Senior leadership must attest to ongoing compliance
- Not just point-in-time certification

## Implications for HE/FE Organisations

### Students & BYOD

- Student-owned devices → Out of scope
- Students themselves → Not directly assessed

**However:**

## **Identity & Access = In Scope**

- Student accounts accessing organisational systems are in scope
- Institutions must secure:
  - Authentication (MFA) – MFA for student accounts is in.
  - Access to cloud services
  - Identity lifecycle management

## **Cloud-First Reality**

- Heavy reliance on cloud platforms in HE/FE now fully recognised
- Core systems (email, VLE, collaboration tools) must be secured and included

## **Increased Risk of Certification Failure**

Common HE/FE weak points now leading to failure:

- Lack of universal MFA for students or staff
- Incomplete cloud service scoping
- Patch management gaps in distributed environments
- Weak control over legacy or lab environments

## **Cyber Essentials Plus Becomes More Demanding**

- Expect:
  - More rigorous endpoint validation
  - Wider sampling across departments
  - Greater scrutiny of real-world controls

## **Strategic Impact**

Shift in Cyber Essentials Philosophy:

- From compliance checklist → operational security baseline
- From device-centric → identity and cloud-centric
- From point-in-time → continuous assurance

## **Recommended Actions for HE/FE**

- **Enforce MFA across all users (including students where feasible)**
- **Inventory and include all cloud services in scope**
- **Align patching processes to ≤14 days for critical updates**
- **Review and tighten scope definitions**
- **Validate backup and recovery processes**
- **Prepare for more rigorous CE+ audits**

## **Bottom Line**

Cyber Essentials 2026 does not expand scope to include student/BYOD devices directly—but it significantly increases expectations around securing the systems and identities they use, making identity and cloud security central to compliance in HE/FE.

## **Action Point:**

If you're preparing for certification or recertification, it's vital to review these updates carefully.





## Chinese Hackers Caught Deep Within Telecom Backbone Infrastructure (Technical)

A China-linked state-sponsored threat actor has deployed kernel implants and passive backdoors deep within telecommunication backbone infrastructure worldwide for long-term persistence, Rapid7 reports.

### [Main Article](#)

The stealth digital sleeper cells have not been attributed to any known APT but are meant for high-level espionage, including against government networks, the cybersecurity firm says.

The persistent tools were deployed as part of apparent discreet breaches that are characterized by recurring elements, suggesting an ongoing operation aimed at “embedding stealthy access mechanisms deep inside telecom and critical environments” for extended access.

As part of its investigation, Rapid7 uncovered passive backdoors and kernel-level implants that have been used in combination with credential harvesters and cross-platform command frameworks.

“Together, these components form a persistent access layer designed not simply to breach networks, but to inhabit them,” the cybersecurity firm [warns](#).

One of the central pieces of the campaign is [BPFdoor](#), a stealthy Linux backdoor that was publicly detailed in 2021, and which uses Berkeley Packet Filter (BPF) functionality for packet inspection within the kernel, reacting to specific packets only.

The cybersecurity firm underlines that BPFdoor’s capabilities make it more threatening than a typical, stealthy backdoor, turning it into an access layer to telecom backbone infrastructure.

“Rather than targeting individual servers, the operators appear to focus on the underlying platforms that power modern telecommunication networks: bare-metal systems running telecom workloads, cloud-native Kubernetes environments hosting Containerized Network Functions, and the signaling protocols that coordinate subscriber identity, mobility, and communication flows,” Rapid7 notes.

This is not the first time Chinese hackers have been caught deep inside critical infrastructure. In early 2024, [CISA confirmed](#) that Volt Typhoon had been “pre-positioning” across US organizations, only months after [Mandiant warned](#) of the hacking group being “clearly dug in”.

In 2024, the networks of [nine US telecom firms](#) were hacked by Salt Typhoon, a Chinese state-sponsored group that continued [targeting telecoms providers](#) in 2025.

**Action Point: This one is definitely for our technical readers.**

Cybersecurity teams need a layered response that accepts “assumed compromise” of backbone infrastructure and focuses on hardening, deep visibility, and rigorous threat-hunting rather than just blocking a single implant family like BPFdoor.

## 1. Strategic posture and scoping

- Treat backbone and core telecom platforms (bare metal, Kubernetes CNFs, signaling stacks) as high-value, likely-compromised assets, not just transport.
- Prioritize cross-agency and vendor coordination (CISA, NSA, NCSC advisories) to align on Chinese state-sponsored TTPs, pre-positioning patterns, and shared indicators for Volt Typhoon / Salt Typhoon-style campaigns.
- Define clear containment thresholds: when detection of BPFdoor-class implants on a backbone node triggers traffic re-routing, out-of-band management, and emergency change control.

## 2. Hardening initial access paths

- Aggressively patch and configuration-harden public-facing appliances and apps abused in these campaigns: Ivanti, Cisco, Fortinet, VMware, Palo Alto, Apache Struts, and other web-facing platforms.
- Enforce strong MFA and conditional access for all admin and “valid account” access paths, plus strict IP allow-listing and just-in-time elevation for device and cloud orchestrator logins.
- Remove or isolate end-of-life devices and software; CISA repeatedly highlights EoL risk and stresses basic hygiene (patching, MFA, logging, lifecycle) as core mitigations for Chinese pre-positioning campaigns.

## 3. Detecting kernel-level implants and passive backdoors

- Deploy and tune EDR/XDR or kernel telemetry on telecom Linux hosts and Kubernetes worker nodes where possible, with attention to: raw socket usage, anomalous packet filtering behavior, and unexpected service masquerading.
- Use Rapid7’s released scanner and similar tooling to detect BPFdoor variants and related implants, looking for behavioral indicators tied to stealth activation mechanisms and kernel packet-filter hooks.
- Hunt for processes posing as hardware monitoring, container runtime components, or other core infrastructure services, since BPFdoor samples can mimic bare-metal platforms and containerization components to blend into noise.

## 4. Network-level detection despite encrypted triggers

- Implement deep packet inspection and network telemetry focused on odd BPF usage patterns and non-standard packet flows at backbone and SSL-termination points (load balancers, proxies), since BPFdoor now hides triggers in HTTPS and abuses termination chokepoints.
- Monitor for crafted packets that carry magic markers at consistent offsets (for example, the 26th-byte marker Rapid7 describes) or use unusual “do not forward” style fields in internal communications, which may indicate backdoor activation chains.
- Instrument ICMP, SCTP, and other “lower-priority” protocols used for signaling or covert control, because newer variants combine HTTPS triggers with ICMP-based signals and target telecom signaling (e.g., SCTP for 4G/5G subscriber tracking).

## 5. Identity, credential, and lateral-movement controls

- Monitor and constrain use of Linux beacon frameworks like CrossC2 / Cobalt Strike-derived beacons, command frameworks, and SSH brute-forcers that intruders are pairing with BPFdoor for staging and lateral movement.
- Deploy strong privileged access management for telecom NOC, Kubernetes, and NFV admins; record and analyze sessions; and restrict direct SSH into backbone nodes, preferring controlled bastions.
- Continuously hunt for anomalous credential use: pre-populated brute-force utilities tuned to telecom environments, unusual keylogger output paths, and suspicious SSH key deployments and sudo patterns.

## 6. Logging, telemetry, and hunting at telecom scale

- Elevate logging for backbone routers, core servers, CNF hosts, and signaling infrastructure (including SCTP, Diameter, 5G core functions), and centralize into SIEM/“big data” platforms capable of multi-year retention and advanced analytics.
- Develop specific hunt playbooks for “dormant access layer” behavior: long-lived kernel hooks with almost no beacons, implants that never open listening ports, and rare activation events tied to magic packets.
- Perform retroactive hunts using newly released indicators (from Rapid7 and joint advisories) to find historical pre-positioning and quietly active sleeper cells in both IT and OT segments.

## 7. Segmentation, resilience, and response

- Segment backbone management planes, signaling networks, and customer data planes as much as operationally feasible; strictly limit cross-segment admin paths and monitoring channels.
  - Design traffic-engineering runbooks to quickly re-route around compromised backbone segments or nodes, maintaining service continuity while forensic teams analyze systems suspected of harboring deep implants.
  - When BPFdoor-class activity is confirmed, perform full incident response: golden-image rebuild of affected hosts, rotation of all credentials and keys, verification of firmware and boot chain where possible, and validation that implants are not present in supply-chain images or orchestration templates.
-



## Stryker Says Malicious File Found During Probe Into Iran-Linked Attack

Medical technology giant Stryker has shared an update regarding its investigation into the recent Iran-linked cyberattack, revealing that a malicious file used by the attackers has been identified.

### Main Article

The Stryker [incident](#) came to light on March 11, with the hacker group Handala taking credit for the attack.

Handala, widely believed to be a hacktivist persona controlled by Iran's Ministry of Intelligence and Security (MOIS), claimed to have wiped more than 200,000 devices, forcing Stryker to shut down offices in dozens of countries.

Some early reports indicated that Handala used wiper malware in the attack — the group has been known to use such malware — but Stryker said it found no evidence of malware or ransomware being deployed on its systems.

The most likely scenario based on the evidence available to date is that the hackers wiped systems by abusing Stryker's Microsoft Intune instance, which is used to remotely manage desktop and mobile endpoints and applications within the organization. There is some indication that Handala leveraged credentials [obtained via infostealer malware](#) to gain access.

Stryker said the incident [disrupted](#) order processing, manufacturing, and shipping. In a [statement](#) issued on Monday, the company said it has made "meaningful progress" in restoring impacted systems.

**Action Point:** In addition to sharing updates on its investigation and restoration efforts, Stryker has made public an assessment from Palo Alto Networks to show that the incident has been contained.

The US government has officially [linked Handala to Iran's MOIS](#) and has taken down several websites used by the threat actor.

In addition, the FBI has issued an [alert](#) sharing information about attacks allegedly carried out by MOIS threat actors such as Handala, including the malware they use.

*"FBI obtained malware samples through investigations. The samples were categorized as masquerading malware (stage 1), persistent implant (stage 2), and related stage 2 malware that contained additional or unique functions. Stage 1 usually masqueraded as commonly used applications like Pictory, KeePass, and Telegram and contained the binaries for the next stage of malware."*

*The persistent implant malware spawned following the masquerading malware's execution and possible user interaction with the malicious application. At this stage, the Iran MOIS cyber actors configured a command and control (C2) using a Telegram bot, allowing bidirectional communication between the compromised device and `api.telegram[.]org`. FBI considered the masquerading malware and persistent implant to be core functionality for the malware campaign."*



## Acumen Cyber Threat Intelligence Digest: Week 12

### Main Article

The digest outlines several newly disclosed vulnerabilities, active threat campaigns, and broader security developments across software, AI tooling, and online services. Citrix reported two flaws in NetScaler ADC and Gateway appliances: CVE-2026-3055, an internally discovered input validation bug causing memory overflow in SAML IdP configurations, and CVE-2026-4368, a race condition leading to user session mix-ups in Gateway and AAA virtual server setups.

Administrators are urged to upgrade to fixed versions. NVIDIA disclosed CVE-2025-33244, a critical deserialization issue in Apex for PyTorch on Linux prior to version 2.6, which can enable arbitrary code execution, data exposure, and denial-of-service, with updates recommended and no exploitation yet reported. Mesop, a Python UI framework, faces CVE-2026-33054, a path traversal bug in versions 1.2.2 and earlier that allows arbitrary file operations and potential application DoS, addressed in Mesop 1.2.3.

Threat activity includes the “Claude Fraud” campaign abusing Claude.ai branding and a malicious VS Code extension to deploy information stealers against developers on macOS and Windows, and a LiteLLM PyPI supply-chain compromise in which malicious package versions exfiltrated credentials and attempted Kubernetes persistence, linked to TeamPCP.

Attackers also backdoored Trivy GitHub Action tags to run an infostealer in CI/CD pipelines. Beyond technical threats, the UK plans a pilot on youth social media restrictions to inform possible policy changes; Crunchyroll is investigating a third-party support breach affecting millions of users; and the UK NCSC warns that AI-driven “vibe coding” could reshape SaaS while increasing insecure software risks if not managed carefully.

### **Action Point:** Remediation Actions

- CVE-2026-3055, CVE-2026-4368 (Citrix) – These vulnerabilities can be addressed by updating NetScaler ADC and NetScaler Gateway deployments to 14.1-66.59 or later.
- CVE-2025-33244 (NVIDIA) – This vulnerability can be remediated by updating PyTorch to version 2.6 or later.
- CVE-2026-33054 (Mesop) – This vulnerability can be addressed by updating Mesop to the newly released 1.2.3.

---

**All the best from all of us at [HEFESTIS](#) and look out for a new [ThreatScape](#) update next week.**

## CISO-Share Office Weekly Newsletter

A very warm welcome to all of you from all of us at the CISO-Share Office.

Here are also a few of the articles which caught the attention of the CISO-Share team this week.

If you have any questions or comments, please feel free to contact us via [CISO-Office@hefestis.ac.uk](mailto:CISO-Office@hefestis.ac.uk).

---

### SC3 Daily Threat Summaries and Weekly Report

<https://www.cyberscotland.com/news/daily-threat-reports/>



A daily breakdown of various cyber threats from the Scottish Cyber Coordination Centre (SC3)

#### Action Point:

All SC3 threat intelligence in one place.

---



### Acrobat Reader zero-day exploited in the wild for many months

<https://www.helpnetsecurity.com/2026/04/09/acrobat-reader-zero-day-exploited/>

Security researchers have confirmed that a **previously unknown (zero-day) vulnerability in Adobe Acrobat Reader** has been **actively exploited in the wild since at least November 2025**—possibly earlier. This isn't a lab demo or proof-of-concept; attackers have been using it quietly for months.

The discovery was made by **Haifei Li**, one of the creators of **EXPMON**, a sandbox system designed to detect advanced file-based exploits. A suspicious PDF submitted to EXPMON triggered multiple deep-inspection alerts. Further analysis showed that once the PDF is opened in Acrobat Reader, it executes **heavily obfuscated JavaScript embedded inside the**



## document.

That JavaScript fingerprints the victim system—collecting details such as **OS version, language settings, Acrobat Reader version, and local file paths**—and sends the data to an attacker-controlled server. Based on the results, the server can then selectively deliver **additional exploits**, including **remote code execution or sandbox escape payloads**. This behaviour strongly suggests a **targeted, adaptive attack**, not mass spam.

When Li analysed the sample, the attacker's server didn't deliver a follow-on exploit, likely because the environment didn't meet the attacker's criteria. However, testing confirmed that the infrastructure *is capable* of delivering live exploits when conditions are right.

A second researcher, **Giuseppe Massaro**, found related samples on VirusTotal. These PDFs used **Russian-language decoy documents** (rendered as images) themed around **gas supply disruption and emergency response**. That strongly suggests the intended targets were **Russian-speaking organisations**, likely in **government, energy, or critical infrastructure sectors**.

Crucially, the exploit **works against the latest available version of Acrobat Reader**, and **no patch has been released yet**. Adobe has been notified, but at the time of writing, there is no fix.

## Action Points

### Treat PDFs as active content

Reinforce the message that PDFs can carry exploits—especially unsolicited or unexpected ones.

### Restrict PDF handling

- Block PDFs from untrusted sources at email gateways where possible.
- Consider opening PDFs in browser-based viewers or sandboxed environments.

### Network-level mitigations

- Block traffic to known attacker IPs: **169.40.2.68** and **188.214.34.20**.
- Monitor or block outbound traffic with the “**Adobe Synchronizer**” User-Agent string.

### Endpoint monitoring

Watch for suspicious behaviour such as:

- AdobeCollabSync.exe making outbound connections
- PDF JavaScript calling APIs like `RSS.addFeed()` or `util.readFileIntoStream()`

### Prepare for patch response

Be ready to **deploy Adobe updates quickly** once a fix is released.

---





## Axios Supply Chain Attack Pushes Cross-Platform RAT via Compromised npm Account

<https://thehackernews.com/2026/03/axios-supply-chain-attack-pushes-cross.html?m=1>

This article covers a **software supply-chain attack** that abused a **compromised npm account** to distribute malware via packages related to **Axios**, one of the most widely used JavaScript HTTP libraries in the world. Axios itself wasn't "hacked" in the traditional sense, but the trust around the ecosystem was abused — which is exactly why supply-chain attacks are so effective.

Attackers gained access to a legitimate npm maintainer account and published **malicious packages that appeared related to Axios**. These packages were designed to look harmless and useful, increasing the chance developers would install them without much scrutiny. Once installed, the malicious code delivered a **cross-platform Remote Access Trojan (RAT)** that worked on **Windows, macOS, and Linux**.

The malware was stealthy and practical rather than flashy. After execution, it collected basic system information, established persistence, and connected back to attacker-controlled infrastructure for command-and-control. From there, attackers could remotely execute commands, steal data, and potentially pivot further into developer environments or production systems.

What makes this particularly concerning is **where the malware landed**. npm packages often run during build processes, CI/CD pipelines, and developer workflows. That means a single poisoned dependency can end up:

- Inside developer laptops
- Inside build servers
- Inside cloud environments
- Potentially baked into production applications

This isn't about end-users clicking dodgy links — it's about **developers doing their jobs** and pulling in code they trust.

The attack also highlights a wider trend: threat actors are increasingly targeting **package registries (npm, PyPI, RubyGems)** because the return on investment is huge. One compromised maintainer or dependency can give access to thousands of downstream organisations.

The malicious packages were eventually identified and removed, but by that point the damage window was already open. Anyone who installed the packages during that period needs to assume exposure.

## Action Points

### Treat supply-chain risk as a primary threat

This is no longer a niche or “developer-only” issue — it’s an enterprise risk.

### Lock down dependency management

- Use allow-lists for approved packages
- Avoid pulling in unvetted or look-alike packages
- Pin dependency versions

### Monitor npm and CI/CD environments

Look for unexpected outbound connections, new persistence mechanisms, or strange build behaviour.

### Audit recent installs

Identify whether affected packages were installed on developer machines, build servers, or production pipelines.

### Enforce least privilege in build systems

Build agents should not have broad access to credentials, secrets, or internal networks.

### Educate developers

Make sure teams understand that “popular ecosystem” does not equal “safe by default”.



### LinkedIn secretly scans for 6,000+ Chrome extensions, collects data

<https://www.bleepingcomputer.com/news/security/linkedin-secretly-scans-for-6-000-plus-chrome-extensions-collects-data/>

A new report has raised eyebrows by revealing that **LinkedIn runs hidden JavaScript on its website to scan visitors’ browsers for installed Chrome extensions**—and it’s not a small list. Testing by *BleepingComputer* confirmed that LinkedIn checks for **over 6,000 extensions** when users visit the site, up from roughly 2,000 last year. This is done using a known browser fingerprinting technique: the site tries to load files linked to specific extension IDs to see which ones are present.

Why does this matter? Because LinkedIn accounts are tied to **real identities, job roles, and employers**. According to the report’s authors, this means LinkedIn could theoretically map which companies use which third-party tools—especially sales, scraping, and enrichment extensions that compete with LinkedIn’s own products. The report claims LinkedIn has already used this data to identify and threaten users of certain third-party tools.

Beyond extensions, the script also gathers **detailed device and browser data**: CPU cores, available memory, screen resolution, timezone, language settings, battery status, audio information, and storage capabilities. Put together, that’s enough to build a **very strong**

**browser fingerprint**, which can be used to track users across sessions—and potentially across sites.

LinkedIn doesn't deny detecting extensions, but it **strongly disputes the intent**. The company says the scanning is used to protect the platform from scraping, abuse, and performance issues, and to enforce its Terms of Service. LinkedIn also says it does **not** use the data to infer sensitive information about users. It claims the report comes from a developer whose extension was restricted for violating LinkedIn's rules, and notes that a German court rejected that developer's legal challenge.

Regardless of motive, one thing is undisputed: **LinkedIn is fingerprinting browsers at scale**, and this isn't an isolated practice. The article notes similar techniques have previously been used by major organisations (including banks and e-commerce platforms) to detect fraud or unwanted software.

For organisations, this sits at the uncomfortable intersection of **privacy, acceptable use, and third-party tooling**—especially where staff install browser extensions on work devices without much oversight.

## Action Points

### Review browser extension policy

Limit which extensions are allowed on managed devices, especially data-scraping or automation tools.

### Audit existing extensions

Identify what's installed today and remove anything unnecessary or risky.

### Brief staff on browser privacy

Make it clear that websites can detect extensions and fingerprint devices—even without malware.

### Separate work and personal browsing

Encourage staff to use managed browsers or profiles for work accounts like LinkedIn.

### Reassess third-party LinkedIn tools

Check whether marketing, or research teams are using tools that could breach LinkedIn's terms.

### Factor this into privacy impact assessments

Especially where LinkedIn is used heavily for recruitment, outreach, or research.

---



## Cyber Threat Intelligence Digest: Week 14

<https://acumencyber.com/cyber-threat-intelligence-digest-april-2026-week-14>

This week's digest is a strong reminder that **internet-facing management systems and identity workflows remain prime targets**, and attackers are moving fast once access is gained.

The most urgent issue is **active exploitation of a critical FortiClient EMS vulnerability (CVE-2026-35616)**. This flaw allows attackers to bypass authentication controls in the EMS API and execute arbitrary code on the server. In simple terms: if your FortiClient EMS is exposed and unpatched, an attacker could fully take it over — and potentially gain access to every endpoint it manages. Over **2,000 exposed instances** were identified online, mostly in the US and Germany, and exploitation was observed *before* public disclosure.

Another high-risk flaw affects **Apache ActiveMQ (CVE-2026-34197)**. It enables remote command execution via the Jolokia management API. While it technically requires credentials, many deployments still use **default admin credentials**, and some older versions are effectively **unauthenticated** due to a separate flaw. This significantly lowers the barrier to exploitation.

**Docker Engine** also patched a serious issue (**CVE-2026-34040**) that lets attackers bypass authorisation controls and spin up **privileged containers**, potentially breaking out to the host system. This is particularly relevant for research, DevOps, and teaching environments that rely heavily on container platforms.

On the threat side, several campaigns stand out:

- A **malicious Chrome extension ("ChatGPT Ad Blocker")** harvested users' ChatGPT conversations and exfiltrated them via a Discord webhook. It abused user interest in blocking ads and used GitHub for dynamic command-and-control.
- A new **phishing-as-a-service platform (VENOM)** targeted executives using fake SharePoint notifications and QR codes. It supports advanced **AiTM** and **OAuth token theft**, allowing long-term access even after password resets.
- **SMS QR-code phishing** impersonating state agencies tricked victims into paying small "fines" and handing over personal and card details.
- The UK's **NCSC** warned that Russian military intelligence is actively compromising **home and small-office routers** to hijack traffic and conduct espionage, often exploiting weak SNMP configurations.

The broader trend is clear: attackers are blending **trusted platforms, legitimate tools, and identity abuse** to stay under the radar.

## Action Points

**Patch immediately**

- FortiClient EMS → apply the hotfix / upgrade to 7.4.7 when released
- Apache ActiveMQ → update to 6.2.3 or 5.19.4
- Docker Engine → update to 29.3.1

## Lock down management interfaces

- Remove EMS, ActiveMQ, Docker APIs from direct internet exposure.

## Review identity protections

- Monitor for AiTM indicators, rogue MFA devices, and OAuth abuse.
- Revoke sessions and refresh tokens when accounts are compromised.

## Control browser extensions

- Restrict extension installs and audit existing ones, especially AI-related tools.

## Secure routers and edge devices

- Disable or restrict SNMP, change defaults, and apply firmware updates.



## Chaos malware expands from routers to Linux cloud servers

<https://www.helpnetsecurity.com/2026/04/08/chaos-malware-cloud-misconfigured-servers/>

Chaos is a **Go-based malware botnet** that's been on the radar for a while, mainly for infecting **routers and edge devices**. What's changed now—and why this article matters—is that researchers have seen **Chaos pivot into Linux cloud servers**, significantly raising the risk profile.

Darktrace spotted the new behaviour in March 2026 using its **CloudyPots honeypot network**. One of the honeypots was deliberately running a **misconfigured Apache Hadoop service**, exposed to the internet and allowing **remote code execution**. Attackers found it quickly.

The intrusion itself was simple and efficient. The attacker sent an HTTP request directly to Hadoop's resource manager endpoint, submitted a fake "application," and embedded shell commands inside it. Those commands downloaded the Chaos malware binary, ran it, and then **deleted it from disk**—a deliberate step to make forensic analysis harder.

Technically, this is a **new generation of Chaos**. Earlier versions focused on router-friendly architectures (ARM, MIPS, PowerPC). This variant is a **64-bit x86 Linux binary**, clearly aimed at cloud and server environments. While some older features (like SSH brute-forcing) were removed, the core capabilities remain:

- **Persistence via systemd**

- **Multiple DDoS attack modes** (HTTP, TLS, TCP, UDP, WebSockets)
- **Command-and-control communication** via infrastructure linked to Hong Kong

The most worrying addition is a **SOCKS5 proxy feature**. Once instructed by its C2 server, Chaos can turn a compromised server into a live proxy. That lets attackers:

- Route attacks through your cloud IPs
- Hide their true origin
- Bypass rate-limiting and IP-based defences
- Pivot into internal networks that trust the compromised host

This reflects a broader criminal trend: **botnets are no longer just for DDoS**. Proxy access can be rented out, monetised, and reused for fraud, credential stuffing, or follow-on attacks.

Crucially, none of this required a zero-day exploit. The entire attack succeeded because a **management interface was exposed and poorly secured**. Hadoop is just the example—this pattern applies to many admin consoles and cloud services if left open.

## Action Points

### Hunt for exposed management interfaces

Audit cloud services for unauthenticated or internet-reachable admin endpoints.

### Lock down Hadoop and similar platforms

Never expose resource managers or admin APIs to the public internet.

### Monitor for proxy behaviour

Look for unexpected SOCKS listeners or outbound proxy-like traffic from servers.

### Watch for short-lived binaries

Malware that downloads, runs, and deletes itself is a red flag.

### Assume cloud servers are high-value targets

Apply the same monitoring and hardening you would to on-prem infrastructure.

### Revisit DDoS assumptions

Botnets now enable pivoting and anonymity, not just traffic floods.



### BlueHammer: Windows zero-day exploit leaked

<https://www.helpnetsecurity.com/2026/04/08/bluehammer-windows-zero-day-exploit-leaked/>

A new Windows zero-day exploit called **BlueHammer** has leaked publicly, and that's the bit that really matters. Once a working exploit is out in the open, it stops being a niche tool for



advanced actors and quickly becomes something **commodity attackers, ransomware crews, and opportunists** can pick up and reuse.

BlueHammer targets **Windows internals** and allows attackers to gain **high-level privileges** on affected systems. In plain English: if an attacker can already get a foothold on a machine (for example via phishing, a malicious file, or stolen credentials), BlueHammer can help them **break out of restrictions**, disable security controls, and take full control of the system.

Security researchers believe the exploit has already been used quietly in targeted attacks, but the leak changes the risk profile overnight. Historically, this is the point where we see:

- Rapid weaponisation
- Proof-of-concepts turned into reliable tools
- Integration into malware loaders and ransomware chains

There is **no patch available yet**, which means defenders can't simply "update and move on". This puts pressure on **defence-in-depth controls** — EDR, privilege management, and behaviour-based detection — rather than signature-based protection.

The exploit also reinforces a recurring theme in modern Windows attacks: attackers don't need to be stealthy for long. They focus on **privilege escalation** early, then quickly:

- Disable or blind security tooling
- Dump credentials
- Move laterally
- Deploy ransomware or data-exfiltration tools

The concern for organisations like universities and colleges is scale. Even if BlueHammer requires local access, **phishing still works**, unmanaged endpoints still exist, and once an exploit is widely available, **skill level stops being a barrier**.

In short: this isn't about one bug. It's about how quickly leaked zero-days **compress attacker timelines** and reduce the warning defenders usually get.

## Action Points

### Assume exploitation will increase

Treat BlueHammer as "will be used", not "might be used".

### Focus on privilege escalation detection

Monitor for unusual elevation behaviour, token manipulation, or SYSTEM-level access.

### Harden local admin usage

- Remove standing local admin rights
- Enforce just-in-time elevation
- Audit service accounts and scheduled tasks



## **Protect security tooling**

Ensure tamper protection is enabled on EDR and Defender components.

## **Prepare for rapid patching**

Be ready to deploy Microsoft fixes quickly once released.

---

All the best from all of us at **HEFESTIS** and look out for our next **Threatscape** update next week.

*HEFESTIS Limited, Registered Office: Unit 27, Stirling Business Centre, Wellgreen, Stirling FK8 2DZ*

*Incorporated in Scotland SC603511*