

CISO-Share Office Weekly Newsletter

Item 09 - Link 09

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.

Table of Contents

CISO-Share Office Weekly Newsletter	1
SC3 Daily Threat Summaries and Weekly Report.....	1
100 Chrome Extensions Steal User Data, Create Backdoor.....	2
McGraw-Hill confirms data breach following extortion threat.....	3
School IT system targeted in cyber attack ahead of exam season	4
NCSC Advisory. TLP AMBER: Mythos AI Model.....	5

SC3 Daily Threat Summaries and Weekly Report



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-13-April-2026.lletin-14-April-2026.lletin-15-April-2026.lletin-16-April-2026.lletin-17-April-2026.



SC3-Weekly-Vulnerability-Report-13-April-2026



100 Chrome Extensions Steal User Data, Create Backdoor

Over 20,000 users installed malicious Chrome extensions designed to provide a backdoor, steal information, or inject ads, cybersecurity firm Socket reports.

[Main Article](#)

The nefarious extensions have been published using five different accounts, namely GameGen, InterAlt, SideGames, Rodeo Games, and Yana Project, but appear to be part of a single, coordinated campaign, based on shared command-and-control (C&C) infrastructure.

Socket identified [108 extensions](#) performing various types of malicious activities. Half of them were designed to steal Google accounts via OAuth2, and 45 were injected with a universal backdoor that opens arbitrary URLs when the browser starts.

The remaining extensions were designed to exfiltrate Telegram sessions, inject ads into YouTube and TikTok pages, inject content scripts into all visited pages, or to proxy translation requests through an attacker-controlled server.

“The 108 extensions are published across several product categories: Telegram sidebar clients, slot machine and Keno games, YouTube and TikTok enhancers, a text translation tool, and page utility extensions. Each targets a different type of user, but all share the same backend,” Socket says.

The extensions provide the expected functionality to avoid raising suspicion, but malicious code running in the background connects to the threat actor’s C&C to perform the nefarious activities.

Action Point: The background script of 45 extensions contains an identical function that, upon browser start, opens a URL received from the C&C in a new tab.

“There is no restriction on what URL the server can return. This channel survives browser restarts and operates independently of whether the user ever opens the extension,” Socket notes.

The cybersecurity firm says it reported all the malicious extensions, but they were not immediately removed from the Chrome Web Store.



McGraw-Hill confirms data breach following extortion threat

Education company McGraw-Hill has confirmed in a statement to BleepingComputer that hackers exploited a Salesforce misconfiguration and accessed its internal data.

Main Article

The company assured that the breach did not affect its Salesforce accounts, customer databases, or internal systems, and that the amount of exposed data is limited and non-sensitive.

"McGraw-Hill recently identified unauthorized access to a limited set of data from a webpage hosted by Salesforce on its platform. This activity appears to be part of a broader issue involving a misconfiguration within Salesforce's environment that has impacted multiple organizations that work with Salesforce," a McGraw-Hill spokesperson told BleepingComputer.

A global education company focused on learning content and platforms, McGraw-Hill offers textbooks, digital learning platforms, and K-12 school and university systems. The company is a major player in education publishing, with an annual revenue of \$2.2 billion.

The statement about the cyberattack comes in response to the extortion group ShinyHunters announcing McGraw-Hill as a victim on its dark-web portal and threatening to leak stolen data by April 14 unless a ransom is paid.

Action Point:

The ShinyHunters data extortion group has carried out several confirmed high-profile security breaches since the start of the year, including those against [Rockstar Games](#), [Hims & Hers](#), the [European Commission](#), [Telus Digital](#), [Wynn Resorts](#), [Canada Goose](#), [Match Group](#), [Panera Bread](#), and [CarGurus](#).

In March, the threat group also breached the American firm [Infinite Campus](#), which also operates a K-12 student information system.



School IT system targeted in cyber attack ahead of exam season

An IT system used by schools across Northern Ireland has been targeted in a cyber attack, the Education Authority (EA) has said.

Main Article

On Thursday, schools received a message that as part of "work to manage an IT security issue" the EA would be carrying out a password reset for all users.

An EA spokesperson said immediate steps were taken to contain the issue and a full investigation is underway, but it could not confirm whether any personal data has been affected.

It has resulted in all schools and pupils being logged out of their accounts, meaning pupils cannot log in to get work or resources provided by their teachers in the run up to exam season.

All online and IT systems in schools in Northern Ireland are provided through the C2K network, managed by the EA.

"We are taking this very seriously and are engaging with the Information Commissioner's Office and relevant authorities as part of our response," the EA spokesperson said.


Action Point:

Kian Hawes, a pupil at Lumen Christi College in Londonderry, said the Easter break usually gives students a chance to catch up on revision and that his teachers had been encouraging them to do so.

But the 14-year-old, who is also the education officer for the Secondary Students Union of Northern Ireland, said the issue with C2K would hinder students ability to revise.

"There's no doubt that it's definitely going to pile pressure on to students," he said.

"And they're going to be frustrated at the fact that now with a week and a half off they're not able to start revision, or get in-depth revision done."



NCSC Advisory. TLP AMBER: Mythos AI Model

Anthropic's latest AI model, Mythos, strengthens cyber-attack capabilities and demonstrates why organisations must prioritise basic cyber security measures in response.

Please find the full notice, containing further information and recommended resilience actions.

TLP AMBER

Mythos AI Model

HANDLING: The information below can be shared by your organisation at TLP:AMBER. TLP AMBER information is subject to limited disclosure, recipients can share it on a need-to-know basis within their organisation and its clients.

Anthropic's model Mythos, is the latest release in a series of recent improvements in frontier AI cyber capability. Whilst the trajectory of continually increasing capabilities is in line with NCSC's expectations, we do believe that this model represents a marked improvement for some elements of cyber capabilities.

In particular, the model has been shown to be effective at identifying vulnerabilities and associated exploits in technology at scale and pace. It has also been able to complete some cyber range challenges, suggesting it has potentially effective network exploitation capability.

AI is changing the cyber landscape in meaningful ways, and while there will be long term cyber benefits, the path to reaching this point brings significant risk and requires urgent action.

All organisations need to get the basics right - ensuring they adopt the five basic controls in Cyber Essentials. Patching quickly, for example, will continue to be important. But as AI continues to change the nature of cyber threat, it reinforces the need for organisations to follow NCSC's best practice advice and guidance.

As outlined in the latest version of the CAF published in August, we urge all organisations to:

- [Reduce what an attacker can access or exploit](#). Utilise scanning tools to understand your exposure. Apply access controls. Where possible, utilise a defence-in-depth approach.
- Apply [patches and security updates](#) as soon as they are available. Optimise processes to allow updates to be deployed at pace and ideally no more than 10 days.
- [Understand your assets and supply chain](#). Ensure you can prioritise your systems and subsystems based on exposure and criticality.
- Perform comprehensive logging and monitoring of devices and systems using modern detection and response tooling. Optimise response pathways to increase response pace and limit adversary impact.
- Strictly apply access and privileged access controls.
- Prepare mitigation measures, such as disconnecting or otherwise isolating network regions, to limit the impact of compromise should it occur.
- Ensure cyber incident and business continuity plans are up to date and exercised in preparation for any incident.
- Make cyber risk a board level priority using the [Cyber Governance Code of Practice](#).

As we said in a recent [blog](#), there is now an imperative for cyber defenders to prepare for the impact of Frontier AI. The majority of the effort needs to be on improving cyber hygiene – architecting, configuring, testing and patching to ensure our digital systems are secure and resilient. Successful cyber defence hasn't changed with the arrival of AI but getting it right has become even more important.

We will continue to update our advice and guidance as capabilities develop.

TLP AMBER

Action Point:

Please reach out to your NCSC POC for broader sharing of the alert on a case-by-case basis.

Report any concerning activity to the NCSC's Incident Management team using [report a cyber incident](#).



Acumen Cyber Threat Intelligence Digest: Week 15

The Cyber Threat Intelligence Digest highlights a surge in actively exploited vulnerabilities, sophisticated attack techniques, and evolving threat actor activity.

Main Article

Several critical vulnerabilities were exploited shortly after disclosure. Adobe patched CVE-2026-34621, a prototype pollution flaw in Acrobat and Reader enabling arbitrary code execution when users open malicious files. A Marimo WebSocket vulnerability (CVE-2026-39987) was exploited within hours of disclosure, granting unauthenticated remote shell access. Additionally, attackers leveraged CVE-2025-0520 in ShowDoc to upload webshells, demonstrating continued exploitation of file upload weaknesses.

Threat actors are increasingly leveraging social engineering and trusted platforms. A phishing campaign targeting Trust Wallet users used QR codes and deep links to trick victims into approving malicious smart contracts, enabling wallet draining via unlimited token permissions. Another campaign compromised CPUID's download infrastructure, distributing trojanised installers that deployed stealthy, memory-resident malware capable of credential theft.

A notable campaign, REF6598, targeted cryptocurrency and financial professionals using LinkedIn and Telegram personas. Victims were lured into using a malicious Obsidian vault with trojanised plugins, leading to deployment of the PHANTOMPULSE remote access trojan, highlighting abuse of legitimate collaboration tools for stealthy compromise.

Broader industry issues include privacy violations by major tech firms allegedly ignoring opt-out signals, and supply chain risks, exemplified by a breach involving Rockstar Games via a third-party analytics provider. Meanwhile, North Korean-linked actors executed a highly sophisticated \$280M cryptocurrency theft through long-term social engineering and supply chain compromise.

Threat actor tracking shows moderate increases in activity, with ransomware groups and emerging actors gaining capability. Overall, the report underscores rapid exploitation timelines, increased social engineering sophistication, and persistent risks in third-party and trusted ecosystems.

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-2026-34621 (Adobe) – This vulnerability can be remediated by updating Acrobat DC and Acrobat Reader DC to version 26.001.21411, alongside version 24.001.30362 on Windows and 24.001.30360 on macOS for Acrobat 2024.
- CVE-2026-39987 (Marimo WebSocket) – This vulnerability can be remediated by updating to version 0.23.0.
- CVE-2025-0520 (ShowDoc) – This vulnerability can be addressed by updating ShowDoc to version 2.8.7.

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.

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.



When Mythos Finds Thousands of Zero-Days, EU Regulators Won't Wait for Your SOC to Catch Up

<https://securityboulevard.com/2026/04/when-mythos-finds-thousands-of-zero-days-eu-regulators-wont-wait-for-your-soc-to-catch-up/>

The core message of this article is uncomfortable but simple: **AI-driven vulnerability discovery has broken traditional SOC and compliance timelines**, especially in Europe.

Anthropic's **Mythos** model is expected to uncover **hundreds or even thousands of previously unknown (zero-day) vulnerabilities at once**. When that happens, the clock doesn't start when you patch — it starts the moment the findings exist. And under EU regulations, those clocks are brutally short.

If you're covered by **NIS2**, **CRA**, or **DORA**, a single Mythos finding can trigger multiple legal obligations simultaneously:

- **NIS2**: 24-hour early warning, 72-hour full report, fines up to €10m or 2% of global turnover, plus personal liability for senior leaders.
-
- **CRA (Cyber Resilience Act)**: 24-hour notification to ENISA, accelerated remediation, potential product recalls, fines up to €15m.
-
- **DORA**: initial reporting in **4 hours**, ongoing updates, and penalties that can accrue daily — up to **€10m per day** for large financial institutions.
-

The maths is what really hurts. A 10-person SOC can realistically process **around 320 findings in 24 hours**. Mythos is expected to drop **500+ findings in a single disclosure event** — and that's a conservative estimate. At roughly 30 minutes per finding, you're immediately underwater. Miss the DORA window and penalties can snowball into tens of millions before the backlog clears.

The article makes a key point: **every Mythos finding is a regulatory event**, not “just another vuln”. Traditional vulnerability management is built for CVEs arriving gradually, with time to read advisories and prioritise. Mythos doesn't work like that. It produces **actionable, exploit-ready findings at machine speed**, often without the guardrails regulators assumed would exist.

To make matters worse, the same issue can trigger **three separate reporting regimes**, each with different classification rules, evidence requirements, and audiences. CVSS scores don't map cleanly to regulatory definitions of “significant incident”, so **manual judgement is unavoidable** — and at scale, impossible within the deadlines.

The conclusion is blunt: **manual triage is financially and operationally insolvent** in an AI-driven disclosure world. This problem won't be limited to Mythos either — tools like OpenAI's Codex Security are already surfacing tens of thousands of findings.

Action Points

Stop treating vulnerability floods as an IT problem

This is a **regulatory and board-level risk**, not a SOC efficiency issue.

Map regulatory overlap in advance

Know which assets and services fall under NIS2, CRA, and DORA — before the clock starts.

Stress-test your SOC capacity

Run tabletop exercises with **500+ simultaneous findings** and real reporting deadlines.

Automate classification, not just detection

Speed matters more than perfect accuracy in the first hours.

Prepare evidence trails

Regulators care about *when* and *how* decisions were made, not just outcomes.

Brief senior leadership now

Personal liability changes behaviour — make sure boards understand the exposure.



Acumen Cyber Threat Intelligence Digest: Week 16

<https://acumencyber.com/cyber-threat-intelligence-digest-april-2026-week-16>

This week's threat picture reinforces a consistent theme: **attackers are blending social engineering, supply-chain abuse, and unpatched systems to move fast and quietly.**

On the vulnerability side, a **proof-of-concept exploit for CVE-2026-3805** was published for **curl**, a tool that's everywhere — from servers to build pipelines. The bug is a **use-after-free flaw** in curl's SMB handling that can leak heap memory or crash systems. While it's not a straight "remote code execution" bug, curl's ubiquity makes it important. The fix is already available in **curl 8.19.0**, but older versions remain exposed.

SAP also patched a **critical SQL injection vulnerability (CVE-2026-27681, CVSS 9.9)** affecting Business Planning and Consolidation and Business Warehouse. Exploitation hasn't been seen yet, but the flaw could allow **arbitrary code execution** if abused. This is a reminder that "not actively exploited" doesn't mean "safe to delay".

Apache Tomcat admins also get homework. **CVE-2026-34486** affects Tomcat clustering (Tribes) and could allow **unauthenticated remote code execution** due to a cryptographic validation failure that "fails open". Patching and restricting access to the Tribes receiver port are both essential.

On the threat activity front, attackers are leaning heavily into **trust abuse**:

- A **fake Claude AI installer** campaign used ClickFix social-engineering to trick users into running commands that ultimately deliver an **information stealer**, bypassing AMSI and hiding execution in memory.
- A **malicious npm package and a fake video game** were used to distribute **XWorm**, a full-featured remote access trojan, hitting both developers and general users.
- **Ransomware affiliates** linked to the Gentlemen RaaS operation combined SystemBC proxies and lateral-movement tooling to achieve enterprise-wide compromise once they already had domain admin access.

There's also a worrying patch gap: **over 1,300 internet-facing SharePoint servers** remain unpatched against **CVE-2026-32201**, a spoofing flaw already exploited as a zero-day. Low complexity, no user interaction, and internet exposure is a bad mix.

Finally, the **Vercel breach** reappears as a cautionary tale: third-party AI tools, OAuth tokens, and environment variables now sit firmly in the supply-chain risk category.

Action Points

Patch with urgency

- curl → update to **8.19.0**
- SAP BPC/BW → apply April patches
- Apache Tomcat → update and lock down TCP/4000
- SharePoint → patch exposed servers immediately

Treat AI installers and dev tools as untrusted

Educate users that “AI” branding doesn't equal legitimacy.

Lock down supply-chain risk


Audit npm usage, preinstall hooks, and build pipelines.

Reduce internet exposure

Continuously inventory and protect internet-facing assets.

Assume attackers abuse trust, not just exploits

Design controls around social engineering, OAuth abuse, and lateral movement.



AI-powered defense for an AI-accelerated threat landscape

<https://www.microsoft.com/en-us/security/blog/2026/04/22/ai-powered-defense-for-an-ai-accelerated-threat-landscape/>

Microsoft's message is blunt: **AI has changed the pace of cyber risk**. Modern AI models can now discover vulnerabilities, chain together low-severity flaws, and generate working exploits far faster than humans. That means the time between “bug found” and “system compromised” is shrinking dramatically. Waiting weeks or months to patch simply isn't viable anymore.

The upside is that defenders can use the *same* AI capabilities to swing the balance back. Microsoft argues that security teams need to move from reactive patching to **continuous exposure reduction**, using AI to find weaknesses earlier, prioritise what actually matters, and automate fixes at scale.

A big part of this sits inside Microsoft's **Secure Future Initiative (SFI)**. Over the last two years, Microsoft has been embedding AI into its own security engineering, including vulnerability discovery and remediation. They're also working with external model providers rather than betting on a single AI. Through **Project Glasswing**, Microsoft is testing advanced models like **Claude Mythos Preview** using its open-source **CTI-REALM** benchmark to see whether models can genuinely help with real-world detection engineering.

Microsoft groups its defensive approach into **three practical areas**.

First is **AI-led vulnerability discovery**. Advanced models are being built directly into Microsoft's Security Development Lifecycle to identify bugs earlier and push fixes faster. For SaaS and PaaS customers, this happens automatically. For anyone running Microsoft software on-prem or self-hosted, staying fully patched is no longer "best practice" — it's a baseline requirement.

Second is **reducing exposure**, not just patching. Microsoft highlights five areas where AI-driven attacks gain leverage: unpatched systems, open-source code, customer source code, internet-facing assets, and poor baseline hygiene. To tackle this, Microsoft Security Exposure Management brings together tools like **Defender External Attack Surface Management**, **GitHub Advanced Security**, **Copilot Autofix**, and **Baseline Security Mode** across M365 and Entra. A new "**Secure Now**" experience ties guidance directly to action.

Third is **AI-powered defence at scale**. Microsoft is building multi-model AI scanning and prioritisation tools to help organisations turn huge volumes of findings into *actionable fixes*, not noise. A preview of this is expected in **June 2026**.

The takeaway: **security teams must act faster than attackers**, and AI is now essential to doing that.

Action Points

Patch relentlessly

Assume attackers can weaponise new bugs fast — delays increase risk sharply.

Focus on exposure, not just alerts

Prioritise internet-facing assets, open-source dependencies, and identity hygiene.

Use posture management tools properly

Continuously discover assets and reduce what's exposed.

Automate where possible

Manual security processes won't keep up with AI-driven threats.

Treat AI as a defensive necessity

Not hype — a requirement to operate at modern speed.

Pass the key, passwords have passed their sell-by date

https://www.theregister.com/2026/04/23/ncsc_passkey_tech_now_reliable/



The UK's **National Cyber Security Centre (NCSC)** has made a big call: **passkeys should now be the default way people log in**, and passwords should be avoided wherever passkeys are available. This is the first time the NCSC has explicitly told organisations and consumers to move away from passwords altogether, rather than just "use stronger ones".

This isn't a sudden decision. The NCSC has been circling passkeys for a while but previously held back due to practical issues: confusing naming across platforms, patchy device support, and poor compatibility with credential managers. According to the NCSC, those gaps have now closed enough for passkeys to be viable at scale.

So why the confidence? The NCSC's technical report concludes that passkeys are **at least as secure as, and generally more secure than**, the traditional combination of a password plus two-step verification (2SV). Passkeys use **cryptographic key pairs** tied to a user's device. There's nothing to remember, nothing to reuse, and crucially **nothing to phish**. Even if a user is tricked into visiting a fake site, there's no secret they can hand over.

The user experience is also better. Logging in with a biometric prompt or device PIN is faster and avoids the endless cycle of forgotten passwords, resets, and MFA fatigue. The NCSC notes that passkeys can be **up to eight times quicker** to use than passwords.

Industry adoption helped tip the balance. The NCSC name-checks **Google, eBay, and PayPal** for improving passkey roll-outs, and says around **50% of UK Google users** have already registered at least one passkey. Microsoft quietly made passkeys the default almost a year ago. In other words, this isn't experimental anymore — it's mainstream.

That said, passwords aren't dead everywhere. Where passkeys aren't yet supported, the NCSC still recommends **password + 2SV**, backed by a **password manager** to ensure credentials are unique. That reduces the impact of infostealer dumps and credential reuse attacks — still one of the biggest causes of compromise.

The wider message is about resilience. With the NCSC warning that the UK continues to face **multiple nationally significant cyber incidents each week**, basic hygiene that actually works at scale matters more than ever.

Action Points

Make passkeys the default where possible

Enable and promote passkeys for staff- and student-facing services.

Update authentication policy

Explicitly state: *no passwords where passkeys are supported*.

Keep MFA where passkeys aren't available

Use password managers + 2SV as the minimum standard.

Prepare for user education

Explain that passkeys are **safer and easier**, not just “another security thing”.

Review legacy systems

Identify applications blocking passkey adoption and plan remediation.



AI cloud company Vercel breached after employee grants AI tool unrestricted access

<https://www.tomshardware.com/tech-industry/cyber-security/vercel-breached-after-employee-grants-ai-tool-unrestricted-access-to-google-workspace>

Vercel, the cloud platform behind **Next.js**, has confirmed a breach that didn't start with Vercel at all — it started with **an employee granting an AI tool far too much access**.

The incident traces back to **Context.ai**, an enterprise AI platform used to build internal agents trained on company knowledge. A Vercel employee signed up to Context.ai using their **corporate Google Workspace account** and approved “**Allow All**” **OAuth permissions**. That single click effectively handed the AI tool unrestricted access to the employee's Google Workspace. From there, attackers were able to **take over the account and move laterally** into Vercel's internal systems.

The threat actor, linked to the **ShinyHunters** group, is reportedly demanding **\$2 million** for the stolen data. Vercel says the exposed data mainly consisted of **non-sensitive environment variables**. Anything explicitly marked as *sensitive* was encrypted and not accessed — but variables that weren't flagged that way should be assumed exposed.

Digging deeper, the chain gets more uncomfortable. Context.ai acknowledged its own compromise, and threat intelligence firm Hudson Rock claims the original access stemmed from a **Lumma Stealer infection** on a Context.ai employee's machine after downloading **Roblox cheat scripts**. That infostealer allegedly grabbed Google Workspace credentials and various API keys, widening the blast radius. Vercel hasn't independently confirmed all of those details, but it shows how **consumer-grade malware can cascade into enterprise breaches**.

Vercel described the attacker as “**highly sophisticated**”, noting how quickly they moved and how well they understood internal systems. In response, Vercel engaged **Mandiant**, notified law enforcement, and contacted affected customers. They also rolled out new dashboard features to make it easier to see and manage environment variables, particularly those marked as sensitive.

Importantly, Vercel confirmed that its **open-source projects** — including Next.js and Turbopack — were **not compromised**. This wasn’t a code-tampering incident; it was an **identity, OAuth, and third-party risk failure**.

The bigger lesson is uncomfortable but clear: **AI tools with broad OAuth access are now part of your attack surface**. One enthusiastic user, one overly generous permission prompt, and suddenly attackers have a legitimate-looking way in.

Action Points

1 Lock down OAuth permissions

Block “Allow All” scopes and enforce admin approval for third-party AI tools.

Treat AI tools as suppliers

Apply vendor risk checks before staff connect them to corporate identities.

Audit environment variables

Rotate keys and secrets stored outside “sensitive” classifications.

Monitor Entra logs

Look for unusual OAuth grants and rapid lateral movement.



If cyber espionage via HDMI worries you, NCSC built a device to stop it

<https://www.helpnetsecurity.com/2026/04/23/ncsc-silentglass-cybersecurity-tool-hdmi-displayport/>

The UK’s **National Cyber Security Centre (NCSC)** has built and released a new hardware

security device called **SilentGlass**, designed to protect **HDMI and DisplayPort connections** from being abused for cyber espionage. It’s a reminder that not all cyber risk comes through email, the internet, or software — sometimes it comes through **the cables we barely think about**.

Monitors aren’t just “dumb screens”. Modern displays contain processors, memory, firmware, and in some cases networking capability. According to the NCSC, attackers can abuse display connections to **extract sensitive information, pivot into networks**, or support **espionage and disruption**. These risks are higher where devices come from complex

supply chains, are serviced by third parties, or are located in shared or publicly accessible spaces.

SilentGlass is a **plug-and-play hardware device** that sits between a computer and a monitor. Its job is to act as a **security boundary** for the display link, preventing malicious signalling or data leakage over HDMI or DisplayPort. Crucially, it doesn't require complex configuration, specialist skills, or expensive infrastructure changes — something the NCSC explicitly calls out as a weakness of existing mitigations.

The NCSC says SilentGlass has already been **deployed inside government systems** and approved for use in **high-threat environments**. It's now commercially available, manufactured and distributed under licence by **Goldilock**, working with **Sony UK Technology Centre**. The message here is important: this isn't experimental or theoretical — it's already in real-world use.

Industry partners involved in the project argue that hardware interfaces like display cables have been a long-standing blind spot. We treat networks, endpoints, and identities as security boundaries, but **physical interfaces** have often been trusted by default. SilentGlass is aimed at closing that gap, especially in environments like **government, critical national infrastructure (CNI), and high-risk commercial organisations**.

For most organisations, this isn't about panic — it's about **proportionate risk management**. The threat is real, but the control is targeted and relatively low-friction.

Action Points

Re-evaluate physical interfaces as security boundaries

HDMI and DisplayPort should be considered part of your attack surface in high-risk areas.

Identify high-risk environments

Focus on SOCs, exec areas, research labs, CNI-adjacent systems, and shared spaces.

Review supply-chain and servicing risks

Third-party maintenance and off-site manufacturing increase exposure.

Consider hardware controls where software can't help

SilentGlass-type solutions reduce risk without relying on endpoint configuration.

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 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.

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.

In the Loop



What's new, updated, or recently released in Microsoft Security

<https://www.microsoft.com/en-us/security/blog/2026/04/30/whats-new-updated-or-recently-released-in-microsoft-security/>

Microsoft's latest security update is very clearly about one thing: **getting security ready for an "agentic" AI world**, where AI agents don't just suggest things — they **take actions**, access data, and interact across systems on our behalf.

The headline message is that security needs to become **ambient and autonomous**, scaling at the same speed as AI. Microsoft is positioning its platform as an **end-to-end, AI-first security stack**, grounded in **Zero Trust for AI**, fuelled by over **100 trillion daily threat signals**, and shaped by its **Secure Future Initiative**. In short: assume AI agents are everywhere, and secure them by default.

One of the most notable releases is new **Microsoft Defender capabilities for Agent 365**, currently in preview. These are designed to give security teams visibility into **what AI agents are doing**, what they're allowed to do, and whether their behaviour looks risky. Defender can now **detect, block, and investigate anomalous or malicious agent actions** in near real time, using webhooks to stop dangerous activity *before* it executes. This is a big shift — from detecting outcomes to **governing intent and behaviour**.

Another important update is the **general availability of Microsoft Defender for Cloud integration with GitHub Advanced Security**. This closes a long-standing gap between AppSec and SecOps. The integration maps code vulnerabilities directly to **runtime and production environments**, helping teams focus on issues that actually affect live workloads. It also supports **co-ordinated remediation**, so fixes don't get lost between dev and security teams.

On the data side, Microsoft has released a new **hands-on demo for Microsoft Purview Data Security Investigations**. This shows how analysts can use AI-powered content analysis to investigate breaches, leaks, fraud, or insider risk — all from a single interface. The standout feature is the **data risk graph**, which correlates sensitive data, users, and activity to help analysts understand *impact*, not just alerts.

Overall, this update reinforces Microsoft's direction of travel: **security that's continuous, integrated, and capable of operating at machine speed**, rather than relying on human reaction times.

Action Points

Prepare for AI agent governance

Treat AI agents as identities with permissions, not just tools.

Review Defender coverage

Assess whether your current Defender setup can monitor agent behaviour.

Bridge AppSec and SecOps

If you use GitHub, look at enabling the Defender–GitHub integration.

Revisit data investigation workflows

Purview's approach may reduce siloed investigations.

Align with Zero Trust for AI

Assume autonomous action, enforce least privilege, and verify continuously.



Another AI-Assisted Software Scan Yields 9-Year-Old Linux Bug

<https://www.darkreading.com/vulnerabilities-threats/ai-assisted-software-scan-linux-bug>

This article is another example of something slightly uncomfortable for defenders: **AI-assisted analysis has just uncovered a serious Linux vulnerability that's been sitting in plain sight for nine years.**

The flaw lives deep in Linux code that's widely deployed across servers, appliances, containers, and embedded systems. It isn't new, it isn't exotic, and it wasn't introduced by some recent change. It's simply been **missed by humans for nearly a decade**, despite years of audits, code reviews, and security research.

What's changed is *how* the bug was found. Researchers used **AI-assisted static analysis**, allowing tools to reason across large, complex codebases and identify unsafe logic paths that are extremely hard to spot manually. The AI didn't just flag a suspicious line — it traced execution paths and edge cases that rarely occur in normal testing.

That's the real takeaway: **AI is compressing the discovery timeline**, not just for defenders, but inevitably for attackers too. If defenders can use AI to uncover long-standing flaws, so can threat actors — and once a bug is publicly disclosed, exploitation often follows quickly.

The vulnerability itself can lead to **serious impact**, including memory corruption and potential privilege escalation, depending on how the affected code is reached. In practical terms, that means:

- Systems running vulnerable kernels may be exploitable locally
- Containers and shared environments could be higher risk
- Cloud and appliance vendors may inherit exposure without realising

This also challenges a common assumption in risk discussions: *"If it's been around for years, it must be safe."* The opposite is increasingly true. Old code often lacks modern defensive thinking, and AI tools are now very good at interrogating it.

This isn't about panic — it's about **expectation management**. We should assume there are **more long-lived vulnerabilities still undiscovered**, and that AI will surface them in clusters, not one-by-one.

Action Points

Patch Linux kernels promptly

Once fixes land, treat them as priority — age of the bug doesn't reduce risk.

Don't assume maturity equals safety

Legacy code paths deserve *more* scrutiny, not less.

Expect vulnerability "bursts"

AI will uncover multiple issues at once — plan for patching at pace.

Watch your supply chain

Appliances, containers, and SaaS platforms may inherit kernel risk silently.

Re-evaluate threat modelling

Assume attackers will gain access to the same AI-assisted discovery tools.

Communicate clearly to leadership

This isn't failure — it's the reality of AI-accelerated discovery.



Acumen Cyber Threat Intelligence Digest: Week 17

<https://acumencyber.com/cyber-threat-intelligence-digest-april-2026-week-17>

Week 17 reinforces a trend that's now well established: **attackers are succeeding by abusing trust, identity, and patch gaps rather than burning new zero-days**. The threat landscape isn't getting noisier — it's getting *cleaner*.

A major focus this week is **continued exploitation of known vulnerabilities**, particularly in **internet-facing enterprise platforms and collaboration tools**. In multiple cases, patches have been available for weeks, yet large numbers of exposed systems remain unremediated. Where exploitation is confirmed, the attack paths are low-complexity and automated, reinforcing that **speed of remediation now matters more than severity scoring**.

Identity-centric attacks remain dominant. Phishing has evolved into **token theft, OAuth abuse, and adversary-in-the-middle (AiTM) techniques** that bypass MFA entirely. These campaigns rely on legitimate authentication flows and trusted infrastructure, making them difficult to distinguish from normal user activity. Once access is gained, attackers move quickly to persistence and lateral access rather than noisy malware deployment.

Supply-chain risk also features strongly. Threat actors continue to abuse:

- Open-source ecosystems (npm, CI/CD tooling)
- Third-party SaaS integrations
- AI and productivity tools granted broad permissions

These paths allow attackers to **scale across organisations**, even when individual victims believe exposure is limited. The Vercel/AI-tool pattern remains a reference case for how quickly OAuth misuse can translate into internal compromise.

There is also increased visibility of “**living-off-the-land**” techniques. PowerShell, mshta, signed binaries, and native OS tooling are repeatedly used to evade detection. Rather than developing new malware, attackers are optimising **reliability and repeatability**.

From a sector perspective, **education, government, healthcare, and cloud-heavy organisations** continue to be disproportionately affected due to:

- Large, diverse user populations
- Mixed device estates
- Complex identity and access models

The overarching message of Week 17 is pragmatic rather than alarmist: **most of the activity observed succeeds because organisations leave doors open**, not because attackers are doing anything especially novel.

Action Points

Treat patch delay as a primary risk indicator

Once exploitation is public, every unpatched day increases likelihood of compromise.

Harden identity flows, not just endpoints

Focus on MFA bypass, token theft, OAuth scope abuse, and session persistence.

Reassess trust in third-party tools

Especially AI, developer, and productivity tools with broad permissions.

Watch behaviour, not signatures

Living-off-the-land attacks rarely look like “malware”.

Reduce blast radius everywhere

Least privilege, time-bound access, and segmentation limit damage.

Assume attackers will reuse what works

If it's effective once, it will be automated and reused at scale.



Microsoft to roll out Entra passkeys on Windows in late April

<https://www.bleepingcomputer.com/news/microsoft/microsoft-to-roll-out-entra-passkeys-on-windows-in-late-april/>

Microsoft is getting ready to **roll out Entra passkeys on Windows**, with general availability expected in **late April**. This is a meaningful step towards making **passwordless authentication the default**, rather than something security teams have to fight to enable.

Passkeys replace passwords with **cryptographic credentials** stored securely on a user's device and protected by biometrics or a PIN. There's nothing to remember, nothing to reuse, and crucially **nothing to phish**. Even if a user is tricked into visiting a fake login page, there's no secret they can hand over.

With this rollout, users will be able to sign into **Entra-protected resources** (including Microsoft 365 and other enterprise apps) using passkeys on Windows devices. It builds on Microsoft's broader push towards passwordless sign-in, but this time it's much more **mainstream and integrated**, rather than an advanced or niche feature.

From a security perspective, this directly tackles two of the biggest causes of compromise:

- **Phishing and MFA fatigue**
- **Credential reuse following infostealer infections**

Passkeys simply don't work in the attacker's favour here. There's no password database to dump, and no MFA prompt to spam until someone clicks "Approve".

There's also a user-experience benefit, which matters more than we sometimes admit. Logging in with a fingerprint or face scan is faster and less frustrating than juggling passwords and one-time codes. Microsoft has been clear that usability is part of the security strategy: if secure login is easier than insecure login, adoption follows naturally.

For organisations like universities and colleges, this matters because of **scale and diversity**. Large user populations, personal devices, shared machines, and frequent account turnover all make traditional password models painful and risky. Passkeys reduce reliance on training people not to click links — instead, they **remove the thing attackers want to steal in the first place**.

That said, this isn't a magic switch. Legacy applications, shared accounts, and non-Windows platforms will still need passwords and MFA for some time. But this rollout marks a clear direction of travel: **passwords are being pushed to the edges, not the centre**.

Action Points

Plan passkey adoption early

Decide which user groups and services should move first.

Update identity policies

Where passkeys are available, make them the default sign-in method.

Prepare user communications

Explain passkeys as *simpler and safer*, not “another security change”.

Review legacy dependencies

Identify apps and workflows that still require passwords.

Keep MFA where passkeys aren't supported

Password + MFA remains the minimum baseline.

Align with national guidance

This rollout strongly aligns with recent NCSC advice to move away from passwords.



Microsoft Confirms Active Exploitation of Windows Shell CVE-2026-32202

<https://thehackernews.com/2026/04/microsoft-confirms-active-exploitation.html>

Microsoft has confirmed that **CVE-2026-32202**, a vulnerability in **Windows Shell**, is being **actively exploited in the wild**. That's the key takeaway — this isn't theoretical, and it's not something attackers *might* use later. They're using it now.

Windows Shell is a core component responsible for how Windows handles files, shortcuts, and the desktop experience. Because it sits so close to everyday user activity, bugs here tend to be **high-value**: attackers can often trigger them through things users already do, like **opening files or browsing folders**.

In this case, exploitation allows attackers to execute **arbitrary code** in the context of the logged-in user. That may sound limited, but in practice it's often enough to:

- Drop malware
- Steal credentials
- Establish persistence
- Pivot to privilege escalation using follow-on techniques

The vulnerability is particularly concerning because it can be triggered via **specially crafted files**, meaning **email attachments, downloads, or files on network shares** can all act as delivery mechanisms. No exotic exploit chain is required — just a user interacting with something that looks legitimate.

Microsoft has acknowledged active exploitation, which usually means:

- Threat actors already have reliable exploit code
- The exploit is stable enough to be used operationally
- Detection alone is unlikely to be sufficient

This fits a broader pattern we're seeing: attackers increasingly focus on **low-friction, user-adjacent vulnerabilities** rather than noisy network exploits. If you can compromise a system by abusing normal user behaviour, you're more likely to succeed at scale — especially in large, diverse environments like universities and colleges.

The risk is amplified in environments with:

- Large numbers of users
- Shared file systems
- Mixed patch levels
- Personal devices accessing institutional resources

In short, this is the kind of bug that attackers love and defenders can't ignore.

Action Points

Patch as a priority

Apply Microsoft's fix across all supported Windows systems as soon as possible.

Assume exposure until patched

If a device isn't patched, treat it as potentially vulnerable — especially user endpoints.

Reduce risky file handling

Reinforce guidance around opening unexpected attachments or downloaded files.

Monitor post-exploitation behaviour

Look for suspicious child processes, script execution, and unusual persistence mechanisms.

Prepare for follow-on attacks

Shell bugs are often step one — watch for privilege escalation attempts shortly after initial access.



Employee Misuse Is Now a Bigger Cyber Threat Than Hackers

<https://www.digit.fyi/employee-misuse-is-now-a-bigger-cyber-threat-than-hackers/>

from Orange Cyberdefense. Analysing over 139,000 triaged security events, t...

For years, most cyber conversations have focused on **external attackers** — hackers, ransomware gangs, and nation states. This article makes the uncomfortable point that, in many organisations, **employees now pose a bigger cyber risk than outsiders**.

That doesn't mean staff are suddenly malicious. In fact, the majority of "employee misuse" is **unintentional**. People take shortcuts, reuse passwords, forward data to personal accounts, plug in unknown USB devices, or use unsanctioned tools to "get the job done". The problem is scale: when thousands of trusted users have access to systems and data, small mistakes add up quickly.

The article highlights three broad categories of misuse:

1. **Negligent behaviour**

This is the biggest bucket. Things like weak passwords, falling for phishing, oversharing files, bypassing security controls, or ignoring policy because it's inconvenient. These actions often create the initial foothold that attackers later exploit.

2. **Careless or policy-breaking behaviour**

Examples include shadow IT, unauthorised cloud storage, personal email use for work files, or installing browser extensions and tools without approval. These aren't done with bad intent, but they expand the attack surface and bypass visibility.

3. **Malicious insiders**

This is the smallest group but the highest impact. Disgruntled staff, contractors nearing the end of their engagement, or insiders recruited by criminals can deliberately steal data or sabotage systems. Because they already have legitimate access, detection is harder.

What's driving the trend is **access**. Modern organisations give staff broad access to data, SaaS platforms, collaboration tools, and AI services. Identity has become the new perimeter — and every user action is effectively a security decision.

The article's core message is that **traditional security controls don't work well against insiders**. Firewalls and malware detection are designed to stop strangers. Employees already have the keys. That means prevention has to focus on **behaviour, visibility, and culture**, not just technology.

Action Points

Assume mistakes will happen

Design controls that limit blast radius rather than relying on perfect behaviour.

Reduce standing access

Apply least privilege and time-bound access, especially for sensitive systems.

Improve visibility of user behaviour

Look for unusual access patterns, data movement, and risky actions — not just malware.

Tackle shadow IT properly

If staff keep bypassing controls, the problem is often usability, not attitude.

Pay attention to joiners, movers, leavers

Most insider incidents happen around role changes or exits.

Align HR, IT, and Security

Insider risk is as much a people problem as a technical one.

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.

Table of Contents

CISO-Share Office Weekly Newsletter	1
SC3 Daily Threat Summaries and Weekly Report. Ransomware is on the Rise.....	1
King's Speech Outlines Cyber Security and Resilience Bill.....	2
The Canvas Breach: Reframing Higher Ed's SaaS Risk Exposure.....	8
Google Detects First AI-Generated Zero-Day Exploit	10
Is the SOC Obsolete, and We Just Haven't Admitted It Yet?	11
Acumen Cyber Threat Intelligence Digest: Week 19	12

SC3 Daily Threat Summaries and Weekly Report. Ransomware is on the Rise



Find SC3's daily threat summaries for this week here, 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-BuSC3-Daily-Threat-BuSC3-Daily-Threat-BuSC3-Daily-Threat-BuSC3-Daily-Threat-Bu
lletin-11-May-2026.plletin-12-May-2026.plletin-13-May-2026.plletin-14-May-2026.plletin-15-May-2026.p



SC3-Weekly-Vulner
ability-Report-11-M:



King's Speech Outlines Cyber Security and Resilience Bill

The King's Speech confirmed that cyber security and digital resilience are now firmly on the government's legislative priority list. Two linked strands matter most for organisations: the upcoming Cyber Security and Resilience Bill and renewed plans for voluntary digital IDs.

Main Article

The Cyber Security and Resilience Bill is designed to strengthen protections for critical infrastructure and supply chains, following a steady stream of high-profile cyber incidents hitting UK organisations. While the Speech doesn't go into technical detail, the intent is clear: resilience is no longer optional, and contingency planning is expected to be formalised in law. In other words, organisations will be expected not just to prevent attacks, but to prove they can respond and recover when (not if) one succeeds.

The government also signalled that further legislation may follow to:

- Support cyber security research
- Improve the security of enterprise connected devices
- Address risks introduced by AI-enabled technologies

Industry reaction has broadly welcomed the direction of travel, with some voices arguing the proposals still don't go far enough. A recurring theme from commentators is that AI security and national resilience need clearer guardrails, particularly given the risk of hostile state activity and organised cybercrime targeting banks, hospitals, and public services.

Alongside cyber resilience, the Speech reaffirmed plans to introduce voluntary digital IDs, with the ambition of making them available by 2029. These would allow people who choose to opt in to use a digital ID to access public services. While voluntary, this has significant implications for identity assurance, privacy, and trust — especially for public-sector and education bodies that already manage large, diverse identity populations.

More broadly, the Cyber Security and Resilience Bill sits within a wider package of over 35 legislative measures covering the economy, public services, crime, health, and digital governance. The framing is explicit: digital resilience underpins national resilience, and cyber incidents are now seen as a material risk to economic stability and public safety.

The King's Speech outlined a wave of landmark laws, including the Cyber Security and Resilience Bill and a recommitment to Digital IDs.

Action Points

Expect resilience requirements to become mandatory

Review incident response, disaster recovery, and business continuity plans.

Move from “security” to “resilience” language

Boards and regulators will increasingly focus on recovery, not just prevention.

Map supply-chain dependencies

Critical suppliers are likely to fall under increased scrutiny.

Prepare for incident reporting expectations

Future legislation is likely to formalise reporting and response timelines.

Engage leadership early

This is governance-level change, not a technical tweak.

Main implications for Universities and Colleges in Scotland and Across the UK

For universities and colleges, the UK Cyber Security and Resilience Bill is likely to have substantial operational, governance, procurement, and financial implications — even where institutions are not directly designated as “regulated entities” under the legislation.

The higher education sector is already considered a high-risk cyber target because universities hold:

- sensitive research data
- personal and financial information
- intellectual property
- defence- and industry-linked research
- large, decentralised IT estates
- open-access academic networks

The Bill effectively raises the baseline expectations for cyber resilience across the education ecosystem.

1. Universities may become indirectly regulated through supply-chain status

Scottish universities are unlikely to be universally classified as “operators of essential services” in the same way as NHS boards or energy operators. However, many universities:

- provide critical research infrastructure
- host sensitive government-funded research
- operate large data centres or HPC environments
- deliver services to regulated sectors
- rely heavily on managed service providers (MSPs) and cloud providers

The Bill specifically expands regulatory attention to:

- MSPs
- critical suppliers
- digital infrastructure providers
- supply-chain dependencies

This means universities may become subject to contractual or regulatory cyber obligations because they sit inside the supply chain of regulated entities.

For example:

- a university running defence research for government
- a college hosting NHS-linked training systems
- a university data centre supporting public-sector workloads

could face heightened assurance requirements.

2. Procurement and supplier management will become much stricter

One of the largest practical impacts for Scottish HE/FE institutions will be procurement governance.

Institutions will likely need:

- more formal third-party risk management
- stronger supplier due diligence
- contractual cyber-security clauses
- evidence of supplier resilience
- continuous monitoring of critical vendors

This is especially significant because universities often operate highly decentralised procurement models with:

- individual schools/departments buying software
- research groups contracting niche providers
- legacy unmanaged systems
- extensive collaboration with external partners

The Bill pushes organisations toward enterprise-level control of cyber supply-chain risk.

3. Faster incident reporting requirements

Universities and colleges will need significantly more mature incident-response capability.

The Bill is expected to require:

- initial reporting of major incidents within 24 hours
- fuller technical reporting within 72 hours

Even if a university is not directly regulated, MSPs and regulated partners may contractually require this reporting standard.

Implications include:

- 24/7 security monitoring capability
- clearer escalation procedures
- formal incident classification frameworks
- legal/compliance coordination
- improved forensic readiness
- crisis communications planning

Many institutions currently operate relatively lean cyber-security teams and may struggle to meet these timelines without investment.

4. Greater executive and board accountability

The Bill reinforces the idea that cyber risk is a governance issue, not merely an IT issue.

For Scottish universities this means:

- university courts and governing bodies will face greater scrutiny
- principals and executive teams will need clearer cyber oversight
- cyber resilience may become part of institutional risk assurance frameworks
- internal audit committees will likely expand cyber assurance activity

Cyber security may increasingly be treated similarly to:

- health and safety
- financial solvency
- regulatory compliance

This is a major cultural shift for institutions where IT governance has historically been fragmented.

5. Pressure to align with recognised frameworks

Institutions will likely need stronger alignment with frameworks such as:

- the National Cyber Security Centre Cyber Assessment Framework (CAF)
- ISO 27001
- Zero Trust security architectures
- enhanced identity and access management controls

The government has repeatedly referenced CAF alignment in policy materials.

In practice, Scottish universities may need:

- MFA everywhere
- privileged access management
- network segmentation
- asset inventories
- endpoint detection and response
- improved backup resilience
- regular tabletop exercises

6. Research security implications

Research-intensive Scottish universities face particularly important implications.

Areas likely to receive heightened scrutiny include:

- dual-use research
- defence-linked collaborations
- international research partnerships
- sensitive datasets
- AI and advanced computing environments

Institutions may need:

- stronger research system segregation
- tighter access controls
- improved export-control governance
- enhanced monitoring of research environments

The Bill aligns with broader UK concern around state-sponsored cyber espionage targeting universities.

7. Financial and operational burden

This may be one of the most difficult issues for Scottish institutions.

Scottish universities already face:

- constrained public funding
- rising operating costs
- financial pressure from international student volatility

The Bill could require substantial additional spending on:

- SOC/SIEM capability
- cyber staff recruitment
- compliance tooling
- penetration testing
- governance and audit
- supplier assurance
- cyber insurance
- resilience exercises

Smaller colleges may struggle most because they often lack:

- mature security teams
- dedicated CISOs
- specialist incident-response capability

Sector-wide shared services or regional security partnerships may become more common.

8. Increased scrutiny after ransomware attacks

The education sector has experienced repeated ransomware incidents in recent years, making regulators likely to view universities as strategically vulnerable.

The Bill therefore creates pressure for universities and colleges to move away from:

- open legacy architectures
- permissive network models
- fragmented identity management
- decentralised unmanaged infrastructure

toward:

- resilience-focused operating models
- stronger recovery capability
- continuous cyber assurance

Likely actions Scottish institutions should take now

Most universities and colleges should already be preparing by:

1. conducting gap assessments against NCSC CAF
2. mapping critical suppliers and dependencies
3. reviewing incident-response readiness
4. strengthening board cyber governance
5. improving identity and access controls
6. testing backup and recovery capability
7. reviewing cyber clauses in procurement contracts
8. identifying systems supporting critical national infrastructure or sensitive research

Bottom line

For Scottish universities and colleges, the Bill is less about a narrow compliance exercise and more about a structural shift in expectations around operational resilience.

Even where institutions are not directly regulated:

- their suppliers may be
- their partners may be
- their research environments may be
- their contractual obligations almost certainly will be

The practical result is likely to be:

- more formal cyber governance
- stricter supplier management
- faster incident handling
- higher compliance costs
- greater executive accountability
- stronger technical security baselines across the sector.



The Canvas Breach: Reframing Higher Ed's SaaS Risk Exposure

The Canvas breach is being used as a wake-up call for higher education, not because Canvas is uniquely risky, but because it highlights how deeply dependent universities and colleges now are on SaaS platforms they don't directly control.

Main Article

In this incident, attackers gained unauthorised access through a third-party integration / token misuse scenario, rather than exploiting a flaw in Canvas itself. That distinction matters. The core platform wasn't "hacked" in the traditional sense, but the ecosystem around it — integrations, APIs, permissions, and identity trust — created the exposure.

For HE, this is uncomfortable but familiar. Learning platforms like Canvas sit at the centre of teaching, assessment, and student data.

They integrate with:

- Identity providers
- Assessment and proctoring tools
- Collaboration platforms
- Analytics and accessibility services

Each integration expands the attack surface, often invisibly. Once attackers gain access via a trusted token or third-party service, activity can look legitimate, making detection harder and response slower.

The article's key argument is that higher education needs to reframe SaaS risk. Traditional security models focus on "is the vendor secure?" but that's no longer enough.

The real risk lies in:

- How access is delegated
- How long tokens live
- Who can authorise integrations
- How visibility and logging are configured

Another important theme is blast radius. In SaaS environments, a single compromised account or integration can expose data across multiple courses, cohorts, or even institutions, depending on tenancy and permissions. That's particularly sensitive in HE, where personal data, assessment materials, and research may all coexist.

The Canvas incident also reinforces that incident response looks different in SaaS. You can't pull network cables or image servers. You're reliant on vendor logs, vendor timelines, and contractual clarity about responsibilities — all of which need to be agreed *before* something goes wrong.

Ultimately, the article argues this isn't a Canvas problem. It's a sector-wide SaaS governance problem, and HE needs to mature quickly as regulators, students, and partners expect stronger assurances around digital resilience.

Action Points

Inventory SaaS integrations properly

Know every third-party app connected to core platforms like VLEs

Tighten token and API controls

Reduce token lifetimes, scopes, and standing permissions.

Treat SaaS identity as critical infrastructure

MFA, conditional access, and role review are non-negotiable.

Improve SaaS logging and monitoring

Ensure you can actually see anomalous behaviour when it happens.

Rehearse SaaS-specific incident response


Include vendors and integrations in tabletop exercises.

Review contracts and assurances

Clarify logging access, response times, and breach notification duties.

Brief leadership realistically

This is about risk management, not vendor blame.



Google Detects First AI-Generated Zero-Day Exploit

The company published a new report on Monday summarizing its observations on the use of AI in the cyber threat landscape, drawing on data collected recently by Gemini, Google Threat Intelligence Group (GTIG), and Mandiant.

[Main Article](#)

One of the most notable findings is that a prominent cybercrime group leveraged AI to develop a zero-day exploit designed to bypass two-factor authentication (2FA) on an open source web-based system administration tool. The exploit was implemented in a Python script.

The hacker group and the targeted tool have not been named, but Google said it worked with the impacted vendor to prevent mass exploitation, which appeared to be the threat actor's plan.

"Although we do not believe Gemini was used, based on the structure and content of these exploits, we have high confidence that the actor likely leveraged an AI model to support the discovery and weaponization of this vulnerability," Google explained.

It added, "For example, the script contains an abundance of educational docstrings, including a hallucinated CVSS score, and uses a structured, textbook Pythonic format highly characteristic of LLMs training data (e.g., detailed help menus and the clean `_C` ANSI color class)."

Action Point: According to Google, the North Korean group tracked as [APT45](#) sent out thousands of repetitive prompts to recursively analyze CVEs and validate PoC exploits.

"This results in a more robust arsenal of exploit capabilities that would be impractical to manage without AI assistance," Google said in its report.

The [full report](#) also covers autonomous malware operations, AI-augmented defense evasion, supply chain attacks, and threat actors pursuing premium access to LLMs.



Is the SOC Obsolete, and We Just Haven't Admitted It Yet?

For decades, the Security Operations Center (SOC) has been the beating heart of enterprise defense. Analysts monitor dashboards, triage alerts, and investigate incidents around the clock. The SOC is often portrayed as the last line of defense—a place where intelligence meets action. And yet, if we are honest, the SOC as we know it is already obsolete. Not because analysts aren't skilled or diligent, but because the very nature of cyber threats has changed faster than our operational models can keep up.

Main Article

The modern SOC is still largely a human-centric workflow. Analysts pivot between tools, manually enrich alerts, and painstakingly validate detection rules. Security vendors promise a revolution: an AI SOC capable of autonomous investigations, dramatically reduced workloads, and proactive threat response. In practice, these promises remain aspirational.

In fact, Anton Chuvakin and Oliver Rochford recently [co-authored *When Marketing Fails*](#), highlighting the gap between AI SOC marketing claims and reality. Based on vendor interactions, practitioner interviews, and SOC community OSINT, their report shows that while AI can assist SOC analysts, it rarely replaces human effort or autonomously resolves incidents.

Action Point: Humans and AI Operating On Complete Data

The SOC of the future will invert the current model. Signals will feed AI agents continuously, correlations will be drawn automatically, and human analysts will focus on oversight, exception handling, and strategic response. Humans will no longer chase alerts—they will guide autonomous systems, investigate the hardest problems, and make high-stakes decisions with confidence.

The SOC is not failing because it is inherently flawed; it is obsolete because threats have evolved faster than the SOC architecture defending against them. Enterprises that acknowledge this reality and deploy data complete, AI-native security solutions for their SOC will be positioned to survive—and even thrive—against machine-speed adversaries.



Acumen Cyber Threat Intelligence Digest: Week 19

This weeks digest from Acumen Cyber

Recent intelligence highlights several critical vulnerabilities and evolving malware campaigns.

Main Article

Elastic patched CVE-2026-33467, a cryptographic signature verification flaw in Elastic Package Registry that could enable tampered package delivery through man-in-the-middle attacks. Fortinet also addressed two critical flaws (CVE-2026-44277 and CVE-2026-26083) affecting FortiAuthenticator and FortiSandbox, both enabling unauthenticated code execution or privilege abuse. Zoom disclosed three vulnerabilities affecting Zoom Workplace and Zoom Rooms, including privilege escalation issues on Windows systems. Organisations are advised to prioritise patching across all affected platforms.

Threat activity increasingly targets developers and AI tooling. The “InstallFix” campaign uses fake Claude AI installer pages promoted via malicious Google Ads to deploy fileless malware linked to RedLine Stealer, using PowerShell, mshta.exe, and AMSI bypass techniques to harvest credentials and cryptocurrency wallet data. Separately, attackers weaponised a malicious OpenClaw skill called “DeepSeek-Claw” to distribute Remcos RAT and GhostLoader across Windows, macOS, and Linux environments, targeting developers through trusted open-source AI ecosystems.

Researchers also identified PamDOORa, a Linux PAM-based backdoor sold on cybercrime forums that captures SSH credentials, enables covert access, and suppresses forensic evidence after attackers obtain root access.

Broader cyber risk trends include scrutiny of EU surveillance technology exports to authoritarian regimes, a £963,900 ICO fine against South Staffordshire Water after CIOp actors remained undetected for nearly two years, and US congressional investigation into repeated ShinyHunters attacks against Instructure’s Canvas platform affecting educational institutions nationwide.

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-2026-33467 (Elastic) – This vulnerability can be mitigated by updating to version 1.38.0.**
- **CVE-2026-44277, CVE-2026-26083 (FortiSandbox/FortiAuthenticator) – These vulnerabilities can be remediated by updating to version 8.03, 6.6.9, or 6.5.7 or later.**
- **CVE-2026-30904, CVE-2026-30905, CVE-2026-30906 (Zoom) – These vulnerabilities can be remediated by updating Zoom Workplace and Zoom Rooms to the most recent versions.**

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.

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.



What's new in Microsoft Security: May 2026

<https://www.microsoft.com/en-us/security/blog/2026/05/21/whats-new-in-microsoft-security-may-2026/>

Microsoft's May 2026 update continues the same theme we've been seeing: **security is being redesigned around AI, automation, and identity**, not just traditional controls. The shift is subtle but important — it's less about adding new tools, and more about **making security continuous, real-time, and embedded across platforms**.

A key focus area is **AI-driven security operations**. Microsoft is pushing deeper into using AI (including Copilot capabilities) to help security teams:

- Detect threats faster

- Prioritise alerts based on real risk
- Automate investigation and remediation steps

The goal is to reduce the burden on SOC teams who are already overwhelmed with alerts, and instead move towards **decision-assisted security at scale**.

Another major theme is **identity protection and Entra integration**. Microsoft is continuing to strengthen controls around:

- Conditional access
- Token protection
- Privilege management

There's clear alignment with broader industry trends: attackers are increasingly targeting **identity, sessions, and OAuth tokens**, so Microsoft is investing heavily in **identity-centric detection and control**.

On the **Defender side**, updates typically include improved:

- Threat detection coverage (especially for cloud and SaaS workloads)
- Exposure management and attack surface visibility
- Integration across endpoints, cloud, identity, and apps

The emphasis is on breaking down silos — instead of separate tools, Microsoft is pushing a **single, unified security experience**.

There's also continued evolution in **data security with Microsoft Purview**, particularly around:

- Data classification and sensitivity awareness
- Insider risk management
- Investigation workflows powered by AI

This reflects a growing reality: breaches aren't just about access anymore — they're about **what data gets exposed and how quickly you can understand impact**.

Finally, Microsoft is reinforcing its **Secure Future Initiative (SFI)** — focusing on:

- Faster vulnerability discovery
- Secure-by-design product development
- Continuous improvement driven by threat intelligence at scale

The underlying message is simple:

Security needs to operate at machine speed to keep up with machine-driven threats.

Action Points

Lean into automation where possible

SOC teams won't scale manually — start adopting AI-assisted detection and response.

Prioritise identity protection above all else

Focus on:

- Conditional access

- Token/session monitoring
- Privileged access control

Strengthen data visibility and classification

You can't manage breach impact if you don't know what data you hold.

Reframe security as continuous exposure management

Not “alerts and incidents” — but **ongoing risk reduction**.



Microsoft shares mitigation for YellowKey Windows zero-day

<https://www.bleepingcomputer.com/news/microsoft/microsoft-shares-mitigation-for-yellowkey-windows-zero-day/>

Microsoft has shared mitigation guidance for a newly identified Windows zero-day dubbed “**YellowKey**”, and the key point is this: **it's already being actively exploited**, and there's no full fix available yet.

YellowKey targets **core Windows functionality**, giving attackers a way to execute code or escalate privileges once they gain a foothold. Like many modern attacks, it's not a single-click “game over”. Instead, it's used as **part of a wider attack chain** — typically after initial compromise through phishing, malicious attachments, or stolen credentials.

What makes YellowKey concerning is how it fits into the modern attacker playbook:

- Initial access via social engineering or exposed services
- Rapid use of a **privilege escalation exploit** like YellowKey
- Disabling security controls and moving laterally
- Data theft or ransomware deployment

Because it's a **zero-day**, traditional defences like patching aren't immediately available. That shifts the focus onto **mitigations and detection**, rather than prevention alone.

Microsoft's guidance focuses on limiting what the exploit can do, rather than eliminating it outright (for now). That means organisations need to rely on **defence-in-depth** — layering controls so that even if one step succeeds, the overall attack is still disrupted.

Another important point is that exploits like YellowKey are rarely used in isolation. Threat actors will combine them with:

- Credential theft
- Token/session abuse
- Living-off-the-land techniques

So even if the exploit itself is technical, the real-world attack often still depends on **user activity and weak identity controls**.

For environments like universities and colleges, the risk is amplified by:

- Large, diverse user populations
- Mixed device estates (managed and unmanaged)
- High reliance on remote access and shared systems

The takeaway isn't panic — it's realism. This is exactly what modern attacks look like: **fast, multi-stage, and built around known weaknesses plus one "new" exploit to accelerate impact**

Action Points

Apply Microsoft mitigations immediately

Follow official guidance — don't wait for a full patch.

Assume post-compromise use

Focus on detecting what happens *after* initial access.

Harden privilege escalation paths

- Remove unnecessary admin rights
- Monitor unusual elevation behaviour

Strengthen endpoint protection

Ensure Defender/EDR tamper protection and behavioural detection are active.

Prepare for fast patch rollout

Be ready to deploy updates quickly once Microsoft releases a fix.



Cyber Threat Intelligence Digest: Week 20

<https://acumencyber.com/cyber-threat-intelligence-digest-may-2026-week-20>

Attackers are still winning with speed, scale, and simplicity, not necessarily sophistication. The main takeaway is that most successful attacks aren't using cutting-edge techniques — they're exploiting **gaps in patching, identity controls, and user trust**.

A recurring theme this week is the **continued exploitation of known vulnerabilities**, particularly in internet-facing services and enterprise platforms. Despite patches being available, large numbers of systems remain exposed for long periods. Attackers are taking advantage of this using automated scanning and exploitation, meaning if something is visible on the internet and unpatched, it's likely already being targeted.

Identity-based attacks are still front and centre. Rather than relying on malware, attackers are focusing on:

- **Phishing and social engineering**
- **Session/token theft (MFA bypass)**
- **OAuth and API abuse**

This allows them to operate using **legitimate credentials and workflows**, making detection much harder. In many cases, once inside, attackers move quickly to persistence and lateral movement without triggering traditional alerts.

The digest also highlights **supply-chain and ecosystem risk**. Third-party integrations, open-source components, and SaaS tools continue to be abused as entry points. These pathways are attractive because they **bypass perimeter controls entirely** and allow attackers to scale across multiple organisations with minimal effort.

Another notable trend is the continued use of **“living off the land” techniques** — using built-in tools like PowerShell, scripting engines, and native system utilities. This reduces reliance on custom malware and helps attackers blend into normal activity.

From a sector perspective, environments with **large user populations and complex estates** — like higher and further education — remain particularly exposed. The combination of decentralised IT, legacy systems, and heavy SaaS reliance creates multiple weak points attackers can exploit.

The overall message is simple but important: **attackers don’t need new tricks if the old ones still work.**

Action Points

Prioritise patching of internet-facing systems

Assume any exposed vulnerability will be exploited quickly.

Treat identity as your primary attack vector

Focus on MFA resilience, token/session monitoring, and OAuth controls.

Strengthen detection of normal-looking activity

Look for unusual behaviour, not just malware.

Minimise blast radius

Apply least privilege and time-bound access across systems.



Microsoft Takes Down Malware-Signing Service Behind Ransomware Attacks

<https://thehackernews.com/2026/05/microsoft-takes-down-malware-signing.html>

Microsoft has taken legal and technical action to shut down a **malware-signing service** that was helping cybercriminals bypass security controls — particularly in **ransomware campaigns**.

At a high level, attackers were using this service to get their malware **digitally signed with legitimate-looking certificates**. That matters because modern operating systems (like Windows) trust signed software more than unsigned code. In many cases, signed malware can:

- Avoid security warnings
- Bypass application controls
- Slip through endpoint protection more easily

In other words, the attackers weren't just writing malware — they were making it look **legitimate**.

The service operated like a **criminal marketplace**. Threat actors could effectively “rent” access to valid or stolen code-signing certificates, allowing them to quickly package ransomware or other malware in a way that would evade defences. This lowered the barrier to entry — even less sophisticated actors could deploy more effective attacks.

Microsoft's response combined **legal takedown efforts** with technical disruption. By seizing domains, infrastructure, and associated accounts, Microsoft has attempted to break the supply chain behind these signed malware operations. This is part of a broader trend where vendors don't just defend systems — they actively go after **criminal ecosystems**.

The article highlights that these signed malware campaigns were linked to known ransomware activity, reinforcing a wider pattern:

- Initial access (phishing, vulnerability, or credentials)
- Drop malware (often signed to evade detection)
- Escalate privileges, move laterally
- Deploy ransomware or exfiltrate data

The use of trusted certificates is particularly worrying because it **undermines one of the core trust mechanisms** in computing. If attackers can reliably obtain or abuse valid signatures, many traditional security assumptions break down.

It also points to an ongoing problem: **certificate theft and misuse**. Whether through compromised developers, stolen keys, or weak controls around certificate issuance, attackers continue to find ways to weaponise trust.

The disruption is positive, but not a permanent fix. These services tend to **re-emerge quickly** in different forms.

Action Points

Don't trust "signed = safe"

Ensure security tooling doesn't rely solely on code signing trust.

Monitor for anomalous signed binaries

Look for unusual execution of newly signed or rarely seen applications.

Strengthen application control policies

Use allow-listing tied to behaviour and reputation, not just signatures.

Protect your own code-signing keys

If you develop software, treat signing certs as high-value secrets.



Verizon DBIR: Vulnerability exploitation is the dominant initial access vector

<https://www.helpnetsecurity.com/2026/05/20/verizon-2026-dbir-findings/>

The big message from the latest Verizon Data Breach Investigations Report (DBIR) is pretty stark: **exploiting vulnerabilities is now the most common way attackers gain initial access**, overtaking things like phishing and stolen credentials.

That's a shift worth paying attention to. For years, we've told users "don't click links" — and that still matters — but attackers are increasingly bypassing users altogether by targeting **systems that haven't been patched in time**.

The pattern looks like this:

- A new vulnerability is disclosed
- Exploit code appears quickly (often within days)
- Attackers scan the internet for exposed systems
- Anything unpatched becomes a target

The speed is the real issue. The DBIR highlights that attackers are now exploiting vulnerabilities **faster than most organisations can patch them**. In some cases, the time between disclosure and exploitation is measured in **hours or days**, not weeks.

Another key point is that attackers are focusing on **internet-facing systems** — VPNs, web apps, edge devices, and remote access infrastructure. These are often complex to patch, hard to monitor, and easy to overlook in large estates (like universities).

Once inside, the rest of the attack often looks familiar:

- Deploy loaders or backdoors
- Steal credentials
- Move laterally
- Exfiltrate data or deploy ransomware

The DBIR also reinforces that attackers don't need cutting-edge zero-days. Most compromises come from **known vulnerabilities that already have patches available**. In other words, the issue isn't discovery — it's **patching delay and exposure management**.

For sectors like HE/FE, this is particularly relevant. Large estates, decentralised IT, legacy systems, and research environments mean patching isn't always straightforward. But attackers don't distinguish — if it's exposed, it's fair game.

Action Points

Prioritise internet-facing assets

Patch VPNs, web apps, and edge devices first — they're the front door.

Speed up patching cycles

Focus on reducing **time-to-patch**, not just coverage.

Track exposure continuously

You can't patch what you don't know exists — maintain an accurate asset inventory.

Assume "known vuln = active exploit"

Don't wait for confirmation of exploitation in your sector.

Use compensating controls where patching is slow

Network restrictions, WAF rules, and isolation buy you time.

Brief leadership differently

Frame this as **exposure risk**, not just "patch Tuesday hygiene".

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.

Table of Contents

CISO-Share Office Weekly Newsletter	1
SC3 Daily Threat Summaries and Weekly Report. Ransomware is on the Rise.....	1
NCSC Guidance: Designing secure access with ZTNA	2
Over 5,500 GitHub Repositories Infected in 'Megalodon' Supply Chain Attack.....	3
Vulnerabilities Across 1,000 OSS Projects.....	4
Google Launches AI Threat Defense Platform for Enterprise Cybersecurity	5
Acumen Cyber Threat Intelligence Digest: Week 21	6

SC3 Daily Threat Summaries and Weekly Report. Ransomware is on the Rise



Find SC3's daily threat summaries for this week here, 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-BuSC3-Daily-Threat-BuSC3-Daily-threat-bu SC3-Daily-threat-bu
lletin-25-May-2026-1lletin-26-May-2026.plletin-27-May-2026.plletin-28-May-2026.p



SC3-Weekly-Vulner
ability-Report-25-M:



NCSC Guidance: Designing secure access with ZTNA

Zero Trust Network Access (ZTNA) is often introduced to modernise access to applications. However, without changes to the underlying design, these deployments can continue to reflect older models of trust.

Main Article

In many cases, ZTNA tools are deployed in environments that still treat network location as a primary signal of trust. This means the tools may be new, but the underlying approach continues to rely on broad, network-based access rather than more granular, context-driven decisions.

Our [new guidance explains how organisations can design and implement ZTNA](#) in a way that better aligns with zero trust principles, and supports modern network environments.

[The guidance](#) focuses specifically on **network access** within a broader [zero trust architecture](#). It:

- Explains **what ZTNA is**, and how it differs from traditional ‘walled garden’ approaches to access.
- Outlines the **organisational and technical foundations** needed before starting a ZTNA deployment.
- Describes key **design requirements** that ZTNA architectures should meet.
- Provides a **simple reference architecture** showing how ZTNA can be used to access private applications and Software as a Service (SaaS).
- Highlights **common anti-patterns** that can undermine ZTNA.

This guidance does not redefine zero trust, or prescribe a single technical solution, and is not intended to be used as a checklist or compliance framework. ZTNA decisions should always be shaped by an organisation’s users, systems, threats, and operational constraints.

- **Action Point:** This guidance is primarily aimed at architects, security practitioners, and technical decision-makers responsible for designing or evolving access architectures.
 - It will be useful whether you are:
 - exploring ZTNA as part of a broader zero trust strategy
 - replacing or reducing reliance on legacy ‘walled garden’ architecture
 - reviewing an existing ZTNA deployment to check whether it delivers the expected security outcomes
-



Over 5,500 GitHub Repositories Infected in 'Megalodon' Supply Chain Attack

More than 5,500 GitHub repositories were infected with malware in a supply chain attack that relies on automated commits, security researchers warn.

Main Article

The campaign, dubbed Megalodon, relies on GitHub Actions workflows containing a payload designed to steal credentials, keys, tokens, and other secrets.

The workflows, [SafeDep](#) says, were injected through over 5,700 malicious commits pushed to the impacted repositories within a six-hour window, on May 18.

According to the cybersecurity firm, the attackers deployed two payloads as part of the attack. One was designed to add a new workflow that would be triggered on every push and pull request, and another that replaced existing workflows with specific triggers, creating dormant backdoors.


On infected machines, the malware would exfiltrate all CI environment variables, AWS credentials, GCP access tokens, Azure credentials, SSH private keys, Docker and Kubernetes configurations, API keys, database connection strings, GitHub Actions tokens, GitLab CI/CD tokens, and dozens of other types of secrets.

Action Point: Last week, NPM [announced](#) that all NPM granular access tokens with write access that bypass two-factor authentication have been invalidated to prevent supply chain attacks similar to [Mini Shai-Hulud](#).

According to [Ox Security](#), this should prevent account hijacking, but does not resolve the underlying problem, and malicious code will continue to spread through compromised repositories.

"If platforms continue allowing any type of code to be uploaded without serious vetting, the number of attacks will only increase," Ox notes.

"We've entered a new supply chain attack era, and TeamPCP compromising GitHub was only the beginning. What's coming next is an endless wave, a tsunami of cyber attacks on developers worldwide," the cybersecurity firm says.



Vulnerabilities Across 1,000 OSS Projects

Anthropic says its Claude Mythos model discovered thousands of severe vulnerabilities across more than 1,000 open source software (OSS) projects.

Main Article

According to the AI giant, [Mythos](#) Preview has identified more than 23,000 potential vulnerabilities. Of these, 1,900 have been reviewed by external security firms, and 1,726 have been confirmed, including over 1,000 rated 'high' or 'critical' severity.

The findings are still being reviewed, and Anthropic estimates that nearly 3,900 critical and high-severity vulnerabilities will be confirmed based only on current findings. As the scans are ongoing, the company believes the number of severe vulnerabilities may reach 6,200.

Anthropic says more than 1,100 unverified findings have been reported to vendors, and 75 issues with a critical or high severity rating have been patched. Vendors have published 65 security advisories.

"The number of patches is still relatively low for three reasons. First, we're still early in the 90-day window that's set out in our Coordinated Vulnerability Disclosure policy: we expect many more patches to land soon," the AI company [explained](#).

"Second, we are likely to be undercounting patches because some vulnerabilities are patched without a public advisory: in those cases, we're reliant on scanning for the patches ourselves using Claude. Third, the low volume of patches reflects a genuine problem: even at our relatively slow pace of disclosures, Mythos Preview is adding to an already-overloaded security ecosystem," it added.

Action Point: Mythos testing results from Project Glasswing participants

The vulnerabilities described in Anthropic's new report are limited to OSS projects, with much of the scanning conducted by the AI company itself.

Roughly 50 organizations have access to Mythos Preview through [Project Glasswing](#) — Anthropic is concerned that wider access could lead to the model being abused — and several of them have disclosed good results after testing it.



Google Launches AI Threat Defense Platform for Enterprise Cybersecurity

The landscape is changing fast, cyber threats and AI are evolving together, and big tech companies are hustling to build defenses that move at machine speed. Google AI Threat Defense blends cloud security, AI analysis, threat intelligence, and automated remediation into one seamless package.

Main Article

The launch really shows how much pressure businesses are under now. Cybercriminals use AI to accelerate attacks, automate exploit development, and launch sophisticated phishing campaigns. Google's solution aims to ease the strain on security teams by helping them identify critical threats and simplify workflows for fixing and responding.

- Google just rolled out Google AI Threat Defense, a new cybersecurity platform built to help companies fight the growing AI-driven cyberattacks through automation, cloud-based smarts, and AI-powered threat detection.
- The platform combines several security technologies including Gemini AI, Mandiant intelligence, VirusTotal telemetry, Wiz cloud security, and CodeMender remediation.
- The idea is to help organizations prepare, detect, fix, and respond to threats faster, especially as attackers use AI to automate exploits, phishing, and vulnerability discovery.

Action Point: Google AI Threat Defense marks a clear shift toward AI-powered cybersecurity, especially as automated attacks progressively increase. By combining Gemini AI with Mandiant, VirusTotal, Wiz, and CodeMender, Google is aiming to build a unified ecosystem that can detect, prioritize, and react to threats faster and at a bigger scale.

The four-step framework: preparation, detection, remediation, response, signals that cybersecurity is getting more proactive and automated. Still, the company emphasizes that human experts remain central to governance, approvals, and incident handling, even with all these AI-driven processes.



Acumen Cyber Threat Intelligence Digest: Week 21

This weeks digest from Acumen Cyber

Recent intelligence highlights several critical vulnerabilities and evolving malware campaigns.

[Main Article](#)

Elastic patched CVE-2026-33467, a cryptographic signature verification flaw in Elastic Package Registry that could enable tampered package delivery through man-in-the-middle attacks. Fortinet also addressed two critical flaws (CVE-2026-44277 and CVE-2026-26083) affecting FortiAuthenticator and FortiSandbox, both enabling unauthenticated code execution or privilege abuse. Zoom disclosed three vulnerabilities affecting Zoom Workplace and Zoom Rooms, including privilege escalation issues on Windows systems. Organisations are advised to prioritise patching across all affected platforms.

Threat activity increasingly targets developers and AI tooling. The “InstallFix” campaign uses fake Claude AI installer pages promoted via malicious Google Ads to deploy fileless malware linked to RedLine Stealer, using PowerShell, mshta.exe, and AMSI bypass techniques to harvest credentials and cryptocurrency wallet data. Separately, attackers weaponised a malicious OpenClaw skill called “DeepSeek-Claw” to distribute Remcos RAT and GhostLoader across Windows, macOS, and Linux environments, targeting developers through trusted open-source AI ecosystems.

Researchers also identified PamDOORa, a Linux PAM-based backdoor sold on cybercrime forums that captures SSH credentials, enables covert access, and suppresses forensic evidence after attackers obtain root access.

Broader cyber risk trends include scrutiny of EU surveillance technology exports to authoritarian regimes, a £963,900 ICO fine against South Staffordshire Water after CIOp actors remained undetected for nearly two years, and US congressional investigation into repeated ShinyHunters attacks against Instructure’s Canvas platform affecting educational institutions nationwide.

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-2026-33467 (Elastic) – This vulnerability can be mitigated by updating to version 1.38.0.**
- **CVE-2026-44277, CVE-2026-26083 (FortiSandbox/FortiAuthenticator) – These vulnerabilities can be remediated by updating to version 8.03, 6.6.9, or 6.5.7 or later.**
- **CVE-2026-30904, CVE-2026-30905, CVE-2026-30906 (Zoom) – These vulnerabilities can be remediated by updating Zoom Workplace and Zoom Rooms to the most recent versions.**

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.



[FIFA World Cup 2026 Scams Are Already Live: Fake Sites, Banking Malware, and Stolen Logins](#)

The main message from this article is straightforward: cyber-enabled fraud around the 2026 FIFA World Cup is already operating at scale , well before the tournament even begins. Researchers and law enforcement are warning that attackers are actively...

Steve McIntosh

6/5/2026

2 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

6/4/2026

10 views



[Cyber Threat Intelligence Digest – Week 22](#)

Week 22 reinforces a now very consistent picture: attackers are succeeding through speed, automation, and identity abuse , not necessarily advanced techniques. Most activity observed continues to revolve around known weaknesses and operational gaps ...

Steve McIntosh

6/4/2026

12 views



[‘HTTP/2 Bomb’ Exploit Knocks Web Servers Offline in Seconds](#)

Known denial-of-service (DoS) techniques can be chained together in a new exploit that can knock major web servers offline, Calif security researchers warn. Main Article
Dubbed HTTP/2 Bomb and discovered using OpenAI’s Codex, the exploit combines a ...

David Robertson

6/4/2026

10 views



[Google Patches Android Zero-Day Under Active Exploitation](#)

Google has issued a patch for a serious Android zero-day vulnerability that was already being actively exploited in the wild . That’s the key point — this isn’t theoretical risk, it’s something attackers were using before a fix was available. The fl...

Steve McIntosh

6/4/2026

13 views



[Known vulnerabilities behind most application security incidents](#)

The core message of this article is straightforward but important: most application security incidents aren't caused by new, unknown bugs — they're caused by known vulnerabilities that organisations haven't fixed in time . Industry data shows that a...

Steve McIntosh

6/3/2026

17 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.

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.



Cyber Threat Intelligence Digest – Week 22

<https://acumencyber.com/cyber-threat-intelligence-digest-june-2026-week-22>

Week 22 reinforces a now very consistent picture: **attackers are succeeding through speed, automation, and identity abuse**, not necessarily advanced techniques. Most activity observed continues to revolve around **known weaknesses and operational gaps** rather than novel exploits.

A key theme is the **ongoing exploitation of exposed and unpatched systems**. Internet-facing services remain a primary entry point, particularly where patching is delayed or incomplete. Attackers are increasingly automating discovery and exploitation, meaning that as soon as a vulnerability becomes known, **scanning and targeting follow almost**

immediately.

Identity-based attacks also remain dominant. Rather than deploying obvious malware, threat actors are:

- Stealing credentials or session tokens
- Using **adversary-in-the-middle (AiTM)** techniques to bypass MFA
- Abusing OAuth permissions and APIs

This allows attackers to operate using **legitimate access and workflows**, significantly reducing the likelihood of detection. Once inside, activity often looks like normal user behaviour, making traditional security controls less effective.

The digest also continues to highlight **supply-chain and SaaS risk**. Third-party integrations, cloud services, and open-source components are being used as entry points or amplification mechanisms. These routes are attractive because they:

- Sit outside traditional perimeter controls
- Often have **broad permissions once trusted**
- Enable attackers to scale across multiple organisations

Another persistent trend is the use of “**living off the land**” techniques. Attackers rely heavily on built-in tools (e.g. PowerShell, scripting engines, native binaries) rather than custom malware. This approach reduces detection and increases reliability.

From a sector perspective, **education, public sector, and other decentralised environments** remain attractive due to:

- Large user bases
- Diverse and sometimes unmanaged devices
- Complex identity and SaaS ecosystems

The key takeaway is consistent with previous weeks:

Attackers don't need new techniques when known weaknesses remain exploitable at scale.

Action Points

Prioritise exposure over volume

Focus on patching internet-facing systems and high-risk vulnerabilities first.

Treat identity as your primary defence layer

Strengthen MFA resilience, monitor token usage, and restrict OAuth permissions.

Reassess SaaS and third-party risk

Audit integrations, API access, and SaaS configurations regularly.

Monitor behaviour, not just malware

Look for unusual access patterns and activity that deviates from normal.

Assume automation on the attacker side

If something is exposed, expect it to be discovered quickly.



'HTTP/2 Bomb' Exploit Knocks Web Servers Offline in Seconds

<https://www.securityweek.com/http-2-bomb-exploit-knocks-web-servers-offline-in-seconds/>

Known denial-of-service (DoS) techniques can be chained together in a new exploit that can knock major web servers offline, Calif security researchers warn.

Dubbed [HTTP/2 Bomb](#) and discovered using OpenAI's Codex, the exploit combines a compression bomb that targets HTTP/2's header compression scheme (HPACK) with a Slowloris-style hold that prevents the server from freeing memory.

According to California-based cybersecurity firm Calif, the attack potentially affects over 880,000 websites that support HTTP/2 and run default NGINX, Apache HTTPD, Microsoft IIS, Envoy, or Cloudflare Pingora configurations.

Furthermore, the company says, an attack can be launched from a home computer on a 100 Mbps connection and can render any of these servers unavailable within seconds.

The techniques chained by the exploit are not new. In fact, three of the underlying issues were disclosed a decade ago, while another was resolved last year.

Action Point:

These HTTP/2 Slowloris-type issues are abused for memory exhaustion by advertising a zero-byte flow-control window so that the server does not send a response, and then resetting the send timeout to prevent the server from freeing memory allocations.

"What's new here is where the amplification comes from. The classic bomb stuffs a large value into the table and references it repeatedly, so servers learned to cap the total decoded header size," Calif notes.

"Our variant goes the other way: the header is nearly empty, and the amplification comes from the per-entry bookkeeping the server allocates around it. The decoded-size limit never fires because there's almost nothing to decode," the company explains.

Calif also identified a bypass for servers that cap the header-field count, and [released](#) proof-of-concept (PoC) code to demonstrate the attack.

The company says NGINX resolved the bug in April, while Apache rolled out fixes in late May (and issued CVE-2026-49975). Microsoft IIS, Envoy, and Cloudflare Pingora have not been patched at the time of writing.

“The other thing worth noting is how this exploit was found. Both halves have been public for a decade. What Codex did was read the codebases, recognize that the two compose, and build the combined attack. That combination is obvious once you see it, and yet as far as we can tell no human had put it together against these servers,” Calif notes.



Google Patches Android Zero-Day Under Active Exploitation

<https://www.esecurityplanet.com/threats/google-patches-android-zero-day-under-active-exploitation/>

Google has issued a patch for a **serious Android zero-day vulnerability** that was already being **actively exploited in the wild**. That’s the key point — this isn’t theoretical risk, it’s something attackers were using *before* a fix was available.

The flaw allows attackers to gain elevated access on affected Android devices. In practical terms, that can mean:

- Accessing sensitive data (messages, files, credentials)
- Installing additional malicious payloads
- Maintaining persistence on the device

Zero-days like this are particularly dangerous because **defenders have no warning window**. There’s no time to prepare — by the time it’s public, attackers may already be using it.

While details on the exploitation chain are usually limited (to prevent copycat attacks), these types of bugs are commonly used in:

- **Targeted attacks** (e.g. against individuals, researchers, or executives)
- **Mobile spyware campaigns**
- **High-value surveillance operations**

That said, once patched and publicised, the risk often **broadens**, because:

- Threat actors reverse-engineer the patch
- Exploit techniques become more widely available
- Lower-skilled attackers begin to adopt them

So even if initial exploitation is targeted, **mass opportunistic exploitation can follow** if patch adoption is slow.

Another important angle is **device fragmentation**. Unlike tightly controlled environments, Android has:

- Multiple vendors
- Different patch cycles
- Devices that may no longer receive updates

This means many devices will remain **permanently vulnerable**, especially in BYOD environments like universities.

The article reinforces a wider pattern: **mobile devices are now a full part of the attack surface**, not an edge case. They hold credentials, provide access to SaaS systems, and often bypass the same controls applied to corporate laptops.

For HE/FE, where staff and students heavily rely on personal devices, this becomes a **real identity and access risk**, not just a “mobile issue”.

Action Points

Patch immediately

Ensure devices receive the latest Android security updates where available.

Assume some devices won't be patched

Older or unmanaged devices should be treated as higher risk.

Strengthen mobile access controls

Use:

- Conditional access
- Device compliance checks
- App-level protections

Focus on identity protection

Compromised devices often lead to stolen tokens or credentials.



Known vulnerabilities behind most application security incidents

<https://www.helpnetsecurity.com/2026/06/03/csa-application-security-incidents/>

The core message of this article is straightforward but important: **most application security incidents aren't caused by new, unknown bugs — they're caused by known vulnerabilities that organisations haven't fixed in time.**

Industry data shows that attackers don't need zero-days very often. Instead, they rely on a predictable pattern:

- A vulnerability is publicly disclosed
- A patch or fix is released
- Many organisations delay applying it
- Attackers scan for and exploit those gaps at scale

This creates a situation where **the biggest risk isn't discovery — it's remediation speed and execution.**

The findings highlight that many organisations already have the tools and visibility needed to identify vulnerabilities. The problem is more operational:

- Too many alerts to triage
- Weak prioritisation (everything looks urgent)
- Fragmented ownership between Dev, Sec, and Ops
- Slow or risky patching processes

In application environments, this challenge is even worse. Modern apps are built from:

- Open-source libraries
- APIs and microservices
- CI/CD pipelines

This creates a constantly changing attack surface, where vulnerabilities can exist not just in your own code, but in **dependencies and integrations you don't fully control.**

Another key point is that **not all vulnerabilities matter equally**, but many teams treat them that way. Attackers don't. They focus on:

- Internet-exposed systems
- Known, weaponised vulnerabilities
- Weakly protected services

So while organisations are busy fixing low-risk issues, attackers are exploiting the ones that actually lead to compromise.

The article also reinforces that vulnerability management needs to move from:

- **"Find everything and fix eventually"**
- to
- **"Fix what's exploitable, now"**

In short, this is less about discovering more vulnerabilities and more about **acting on the ones that already matter.**

Action Points

Prioritise exploitable vulnerabilities

Focus on:

- Internet-facing services
- Known exploited CVEs
- Systems with sensitive data

Reduce time-to-patch

Speed matters more than completeness — shorten remediation cycles.

Improve Dev–Sec–Ops alignment

Make sure ownership of vulnerabilities is clear and actionable.

Track real exposure, not just scan results

Understand which assets are actually reachable and at risk.

Manage third-party and open-source risk

Keep dependencies updated and monitored continuously.

Use compensating controls when patching is slow

WAF rules, isolation, or access restrictions can buy time.

Change the narrative for leadership

Frame this as **exposure management**, not just “vulnerability numbers”.



FIFA World Cup 2026 Scams Are Already Live: Fake Sites, Banking Malware, and Stolen Logins

<https://thehackernews.com/2026/06/fifa-world-cup-2026-scams-are-already.html>

The main message from this article is straightforward: **cyber-enabled fraud around the 2026 FIFA World Cup is already operating at scale**, well before the tournament even begins.

Researchers and law enforcement are warning that attackers are actively exploiting the hype, scarcity, and urgency around tickets and travel. The result is a **large, organised scam ecosystem** spanning fake websites, malware, and credential theft campaigns.

One of the biggest issues is the sheer volume of **fraudulent domains**. Thousands of World Cup-themed websites have already been registered, many designed to look almost identical to FIFA’s official site. In one case, a coordinated operation created **hundreds of near-perfect phishing pages**, replicating FIFA’s login flow closely enough to steal real user credentials and hijack accounts.

Once attackers capture login details, they can:

- Reset passwords and lock users out
- Access tickets linked to accounts
- Resell those tickets on secondary markets

This turns credentials into **direct financial value**, not just data.

The scams don’t stop at phishing. Attackers are also distributing **banking malware through fake streaming apps** and unofficial downloads. Victims looking for free or cheaper access to matches may install apps that silently steal financial details or credentials.

Another major vector is **payment fraud**. Scam sites accept multiple payment methods — including crypto — which makes recovery almost impossible. A key red flag: FIFA does **not** accept cryptocurrency for official ticket sales.

These campaigns are heavily promoted through:

- Social media ads (especially Facebook)
- Messaging platforms (Telegram, WhatsApp)
- Search results and sponsored links

The scale is significant. Security researchers estimate that ticket-related scams alone could already be responsible for **tens or even hundreds of millions in losses**, with total impact potentially reaching much higher.

The broader takeaway is that global events like the World Cup create the perfect storm: **high demand, emotional urgency, and large financial flows** — exactly what attackers need.

Action Points

Use only official channels

Always go via **fifa.com** — never third-party ticket sites or links.

Be sceptical of urgency

“Limited tickets”, “final chance”, or big discounts are classic scam signals.

Avoid crypto payments for tickets

This is a strong indicator of fraud.

Don't install unofficial streaming apps

They're a common delivery method for banking malware.

Watch for lookalike domains

Small spelling changes or unusual domains = high risk.

Brief users now, not during the event

Scams will peak during the tournament — awareness needs to be early.

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.

Table of Contents

CISO-Share Office Weekly Newsletter	1
SC3 Daily Threat Summaries and Weekly Report. Ransomware is on the Rise.....	1
NCSC Guidance: Designing secure access with ZTNA	2
Over 5,500 GitHub Repositories Infected in 'Megalodon' Supply Chain Attack.....	3
Vulnerabilities Across 1,000 OSS Projects.....	5
Google Launches AI Threat Defense Platform for Enterprise Cybersecurity ... Error! Bookmark not defined.	
Acumen Cyber Threat Intelligence Digest: Week 21	6

SC3 Daily Threat Summaries and Weekly Report.



Find SC3's daily threat summaries for this week here, 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-BuSC3-Daily-Threat-BuSC3-Daily-threat-bu
lletin-10-June-2026.lletin-9-June-2026.plletin-8th-June-2026lletin-11-June-2026.l



SC3-Weekly-Vulner SC3 - Monthly
ability-Report-8-Junransomware report



[SHINYHUNTERS] – Ransomware Victim: nottingham[.]ac[.]uk

AI Generated Summary of the Ransomware Leak Page

The leak post identifies the victim as nottingham.ac.uk, an institution in the Education sector based in the United Kingdom. The entry describes a data-leak scenario rather than a purely encryption-focused incident, indicating that sensitive financial and personal data were exfiltrated. The published material states that the compromised data set comprises billing and payment records, including payer contact information, transaction amounts, and other campus-related financial data, spanning the University of Nottingham and its Malaysia and China campuses. The post notes that the data includes personal identifiers and campus-specific information, underscoring the potential impact on current and former students, staff, and affiliates. The post-date is 2026-06-09, and in the absence of an explicit compromise date, this is treated as the post date for the leak.

Main Article

Students' data taken in major University of Nottingham cyber-attack - BBC News

the incident seems to be related to a wider attack

Oracle PeopleSoft servers hacked in ShinyHunters data theft attacks

Oracle Security Alert Advisory - CVE-2026-35273

The post presents the event as a data-leak scenario with sensitive financial records rather than a purely encrypted environment. The threat actors describe the data as containing payer contact information, transaction amounts, IP addresses, full names, home addresses, postcodes, email addresses, phone numbers, dates of birth, and other internal campus data, suggesting broad exposure of personal and financial records. The material indicates the size of the affected data as at least 40 GB in uncompressed form, with a compressed size of 19 GB plus. The leak page does not provide a ransom figure or a direct download link within the visible content, and there are no accompanying screenshots or images reported on the page. This aligns with a data-leak/extortion pattern where disclosing or threatening exposure of sensitive records is a primary motive, rather than a purely encryption-based disruption. The incident highlights the ongoing risk to higher education institutions handling student and financial information.

Action Point: No files or stolen information are exfiltrated, downloaded, taken, hosted, seen, reposted, or disclosed by RedPacket Security. Any legal issues relating to the content should be directed at the attackers, not RedPacket Security. This blog is an editorial notice informing that a company has fallen victim to a ransomware attack. RedPacket Security is not affiliated with any ransomware threat actors or groups and will not host infringing content. The information on this page is automated and redacted whilst being scraped directly from the SHINYHUNTERS Onion Dark Web Tor Blog page.



Claude Mythos Turns N-Days Into N-Hours With Rapid Exploit Creation

Anthropic says its Claude Mythos Preview model can build working exploits targeting known vulnerabilities within hours, or even minutes.

[Main Article](#)

[Anthropic Launches Claude Fable 5](#)

Announced in early April and promoted as the most capable AI frontier model, Mythos right from the start raised fears regarding its ability to supercharge attacks.

In April and May, Anthropic touted its ability to find vulnerabilities, including [271 Firefox flaws](#) and [thousands of severe security defects](#) across over 1,000 open source software (OSS) projects.

Now, the company says its most advanced model can also weaponize these discoveries, demonstrating that the surge in AI use in cyberattacks increases the threats faced by organizations in the patch gap.

Put to the test, Claude Mythos Preview delivered 16 working exploits targeting Firefox and Windows within hours.

Anthropic's public models were also tested, with safeguards off. While they did not rise to Mythos's level, they too delivered working exploits, proving that LLMs significantly increase the threat posed by N-days that have not been exploited in attacks before.

Action Point: Faster patching amid low exploit costs

“At this speed, Mythos Preview would have finished creating all eight full chain exploits before any of the Windows devices had received the patch as an update. Turning these exploits into a real campaign still requires further work, but Mythos Preview has now collapsed one of the most time-intensive steps into hours,” Anthropic notes.

The cost of building these exploits is not high either, the company says. Each model was given a three-million-token budget for creating the PoCs and exploits targeting Firefox. The cost of creating the full chain exploits targeting Windows was \$15,700 in API credits, or around \$2,000 per privilege escalation.

“The binding constraint to N-days is now just a few thousand dollars and API access, which expands the pool of capable N-day attackers dramatically,” Anthropic says.

The company calls for an updated patching playbook, which should rely on “N-hour” rather than “N-day”, and should no longer assume that weaponizing a patch takes weeks.

“N-days have historically caused most harm to systems that are slow or difficult to patch. Industrial control systems, medical devices, and ‘internet of things’ devices often run on fixed maintenance windows, vendor-locked firmware, or have uptime guarantees. As the cost of weaponizing any given patch falls toward zero, these devices and systems will become even more exposed. And even systems operating on an established, ‘responsible’ patch cadence are now far easier targets than before,” Anthropic notes.



'GreatXML' Zero-Day Exploit Bypasses BitLocker

Security researcher Nightmare Eclipse has released a new Windows BitLocker bypass, only one day after publishing an exploit targeting Microsoft Defender.

[Main Article](#)

[Rogue Planet](#)

Named [GreatXML](#), the fresh exploit allows users to bypass BitLocker and spawn a command prompt with SYSTEM privileges while in Recovery Mode.

The proof-of-concept (PoC) code the researcher released targets a vulnerability in Microsoft Defender's offline scan functionality.

According to Nightmare Eclipse, all systems on which an offline scan was initiated at least once automatically become vulnerable.

The PoC exploit includes an XML file and a Recovery folder (containing another XML) that need to be copied to the root of the computer's recovery partition.

Any Windows machine becomes vulnerable to GreatXML as soon as Defender's offline scanning is initiated. Thus, an attacker simply needs to launch the functionality before executing the exploit.

"If Defender offline scan was never initiated, then you have to either log in and initiate it yourself or figure out a way to boot into WinRE in offline scan state (I believe it should be very possible to do so without logging in)," the researcher says.

Nightmare Eclipse released GreatXML just one day after RoguePlanet, a zero-day flaw in Microsoft Defender that leads to local privilege escalation (LPE) to SYSTEM.

Action Point: Also known as Chaotic Eclipse, Nightmare Eclipse has been dropping exploits for various zero-day vulnerabilities in Windows after expressing discontent with how [Microsoft treats researchers](#) who participate in its vulnerability disclosure programs.

Microsoft has been scrambling to resolve the publicly disclosed flaws, including [BlueHammer](#), [RedSun](#), and [UnDefend](#), which have been exploited in attacks. It also patched [GreenPlasma](#) and [YellowKey](#) with [the June 2026 Patch Tuesday updates](#).



Acumen Cyber Threat Intelligence Digest: Week 23

A rapidly evolving threat landscape driven by critical vulnerabilities, phishing campaigns, malware-as-a-service operations, and increasing geopolitical focus on cybersecurity.

[Main Article](#)

A key concern is the disclosure and patching of several significant vulnerabilities. Google released updates addressing 74 Chrome security flaws, including CVE-2026-11645, an actively exploited vulnerability in the V8 JavaScript engine that could enable arbitrary code execution through malicious web content. TP-Link also patched CVE-2026-34126 affecting several Tapo smart-home devices, where exploitation could allow interception of Bluetooth communications and unauthorized device control. Additionally, HP addressed the critical CVE-2026-0826 vulnerability impacting Poly VVX and Trio conference phones, which could permit unauthenticated remote code execution with root privileges.

The report identifies several active threat campaigns. One notable operation uses fake Amazon security alerts to distribute HarborWatch Agent malware through the ClickFix social-engineering technique. Victims are tricked into executing malicious PowerShell commands that install a remote access trojan capable of collecting extensive system information and communicating with command-and-control infrastructure. Another major threat is WeedHack, a malware-as-a-service platform that has reportedly compromised more than 116,000 systems since January 2026. Distributed through malicious Minecraft-related downloads, YouTube videos, and SEO poisoning, WeedHack steals credentials, cryptocurrency wallet data, browser information, and other sensitive assets while establishing long-term persistence.

The digest also highlights a dependency-confusion campaign involving malicious npm packages designed to mimic internal corporate repositories. These packages execute reconnaissance scripts that harvest development environment data, cloud credentials, API keys, and CI/CD secrets, potentially enabling follow-on attacks against targeted organizations.

Broader cybersecurity developments include upcoming CISA guidance for implementing the U.S. AI executive order, European initiatives to strengthen technological sovereignty, and proposals for a dedicated U.S. Cyber Force. Acumen Cyber recommends immediate patching of affected technologies, proactive vulnerability management, and ongoing threat hunting to reduce exposure to these evolving threats.

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-2026-11645 (Chrome)** – This vulnerability can be addressed by updating Chrome to version 149.0.7827.102 or 149.0.7827.103 for Windows, macOS and Linux.
- **CVE-2026-34126 (TP-Link)** – TP-Link released patches for Tapo L535E versions prior to 1.4.1 Build 251016 Rel.204554, Tapo P300 versions prior to 1.4.2 Build 251219 Rel.142654, and Tapo D100C versions prior to 1.3.1 Build 260421 Rel.031658.
- **CVE-2026-0826 (HP)** – HP have recommended disabling unnecessary ICE connectivity as an additional mitigation measure, as well as updating VVX devices to UCS 6.4.8, Trio 8300 to UCS 8.1.7, and Trio 8500 and 8800 to UCS 7.2.8.

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.

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.



Cyber Threat Intelligence Digest – Week 24

<https://acumen cyber.com/cyber-threat-intelligence-digest-june-2026-week-24>

This week's report paints a familiar picture: **nothing wildly new, but attackers continuing to get better at executing proven techniques** — especially phishing, social engineering, and abusing trusted platforms.

There are three main themes:

- **Steady flow of high-severity vulnerabilities**
- **Ongoing phishing-driven malware campaigns**
- **Real-world impact in the education sector (data breach)**

Vulnerabilities: patching still matters

Several notable vulnerabilities were disclosed and patched:

- **IBM i OS flaw (CVE-2026-7870)** – could allow attackers to run code with admin privileges
- **Apache CXF bug (CVE-2026-50623)** – could enable remote code execution
- **Canon EOS tool vulnerabilities** – multiple issues affecting both Windows and macOS

There's no widespread exploitation yet, but these are exactly the types of issues attackers jump on quickly.

Campaigns: phishing and user trickery still dominate

Two strong examples:

NFCShare (mobile fraud)

- Fake banking apps trick users into installing malware
- Uses NFC to **steal card data and PINs**
- Distributed via phishing sites and even GitHub
-

ClickFix campaigns (very relevant trend)

- Fake “verification” or “error” pages trick users into running commands
- Victims are guided to execute malicious PowerShell
- Payloads steal credentials, email data, crypto wallets

Also worth noting:

- Fake **Amazon security alerts** delivering a remote access trojan
- Heavy abuse of **trusted platforms (Google Sites, GitHub)** to look legitimate

Real-world impact: University breach

- **University of Nottingham** confirmed a breach
- Claimed by ShinyHunters group
- Potentially **40GB of data stolen**
- Around **455,000 individuals affected**

Data includes:

- Personal information
- Contact details
- Potentially financial data

This is a direct reminder of the **sector risk profile** for HE/FE.

Broader trends

- Attackers continue to **abuse trusted services** (Google, GitHub, SaaS platforms)
- **AI-themed lures** (e.g. fake Codex/Claude downloads) are becoming common [[Cyber Threats](#)]
- Threat actors are **incrementally improving**, not reinventing the wheel

This is all about **execution at scale**:

- Users are still the **primary entry point**

- Legitimate tools and services are being weaponised
- Data theft (not just encryption) remains the main goal

For the education sector, the Nottingham breach reinforces:

- Large data estates = attractive targets
- Reputation and regulatory impact are significant

There's no single headline threat here — and that's the point.

Attackers are succeeding by **combining well-known techniques, social engineering, and trusted platforms.**

The organisations that will stay ahead are the ones that:

- Execute the basics consistently
- Focus on user-driven risk
- Improve visibility beyond traditional endpoints

In short: **this is a discipline problem, not a technology problem.**

Actions Points

Double down on patching discipline

- Prioritise **externally exposed systems and middleware**
- Don't delay patching just because exploitation hasn't started yet

Monitor abuse of trusted platforms

- Look for traffic to:
 - Google Sites
 - GitHub
- Combine **domain reputation with behaviour-based detection**

Strengthen endpoint protection

- Detect:
 - In-memory execution
 - Steganography-style payload delivery
- Ensure EDR visibility into **script execution chains**

Prepare for data breach scenarios

- Assume **data exfiltration is inevitable in some attacks**
 - Review:
 - Data classification
 - Monitoring of large outbound transfers
 - Test breach response and communications plans
-



Microsoft Confirms RoguePlanet Defender Zero-Day, Says Patch is in Development

https://thehackernews.com/2026/06/microsoft-confirms-rogueplanet-defender_02022423645.html

Microsoft has confirmed a **zero-day vulnerability in Microsoft Defender**, nicknamed “**RoguePlanet**.”

- It’s now tracked as **CVE-2026-50656** with a CVSS score of 7.8 (high severity).
- The flaw allows **privilege escalation** — essentially letting an attacker gain full control of a system.
- A patch is **not yet available**, but Microsoft is actively working on one.

This is particularly uncomfortable because the weakness sits inside **Defender itself — the thing meant to protect you.**

What makes this interesting (and risky)

The exploit works by abusing a **race condition** in the Microsoft Malware Protection Engine.

In simple terms:

- The attacker triggers a timing issue
- If successful, they get a **SYSTEM-level shell (highest privilege on Windows)**

A few key points stand out:

- It’s a bit unreliable (hit-and-miss), but works consistently on some machines
- It can work even when **Defender real-time protection is enabled**
- It may also work in passive mode (not fully confirmed)

So effectively, if an attacker already has a foothold, this gives them a path to **full system compromise**, bypassing the very control you rely on.

How it fits into real attacks

This isn’t an initial entry vulnerability — it’s a **post-compromise privilege escalation**.

In practice:

1. Attacker gets in (phishing, vulnerability, stolen creds)
2. Uses RoguePlanet to escalate privileges
3. Gains **complete control of the endpoint**
4. Moves laterally or deploys ransomware

Also worth noting: this researcher has found multiple Defender flaws recently, all of which were patched — suggesting **Defender is an increasingly attractive attack surface**.

This is a classic reminder of a few uncomfortable truths:

- **Security tools are not immune from vulnerabilities**
- Endpoint protection can be **turned into an attack path**

- Privilege escalation remains one of the most critical controls to get right

Until a patch lands, this is about **risk mitigation rather than full remediation**.

RoguePlanet isn't about mass exploitation (yet) — it's about **amplifying impact once attackers are in**.

The bigger lesson:

Even your defensive controls can become attack surfaces. The organisations that do well here will be the ones that **layer controls, reduce privilege, and detect behaviour — not just rely on prevention tools**.

Actions Points

Assume “Defender bypass” is possible

- Treat endpoints as **potentially compromisable even with AV active**
- Don't rely solely on Defender for containment

Focus on initial access prevention

Since this is post-compromise:

- Tighten phishing protections
- Patch externally facing systems quickly
- Monitor for use of stolen credentials

Lock down privilege escalation paths

- Remove unnecessary **local admin rights**
- Enforce **least privilege everywhere**
- Monitor for suspicious privilege changes

Strengthen EDR / detection layers

- Ensure behavioural detection is enabled (not just AV scanning)
- Look for:
 - Unusual SYSTEM processes
 - Suspicious process spawning patterns
 - Exploit-like behaviour in Defender components

Increase monitoring on sensitive endpoints

- Prioritise:
 - Admin workstations
 - Servers
 - Backup systems

These are the most valuable targets for privilege escalation.

Test your assumptions

Ask internally:

- *If Defender was partially bypassed, would we still detect an attacker?*
- *How quickly could we spot privilege escalation activity?*



INC Ransomware Emerges as Major RaaS Threat in 2026 with 830+ Victims Since 2023

<https://thehackernews.com/2026/06/inc-ransomware-claims-830-victims-since.html>

A ransomware group called **INC** has quickly grown into one of the most active cybercrime operations around. Since 2023, they've **hit over 830 organisations**, and they're now a major "ransomware-as-a-service" (RaaS) player — meaning they effectively franchise their tools to other attackers.

Their rise has partly come from gaps left by the disruption of other big gangs like LockBit and BlackCat, allowing affiliates to move over to INC.

They're not especially novel — and that's the point. They succeed by **combining well-known techniques at scale**, rather than relying on cutting-edge exploits.

How they operate

- They get in via **phishing, stolen credentials, or unpatched internet-facing systems**.
- Once inside, they **steal credentials**, especially from backup systems like Veeam.
- They move around using legitimate tools (e.g. RDP, admin utilities), blending in with normal activity.
- They often deploy tools like **Cobalt Strike or remote access software** for control.
- Data is **exfiltrated first**, then systems are encrypted (double extortion).

They've also modernised their malware (rewritten in Rust), making it harder to analyse and easier to run across different environments, including Linux and virtual infrastructure.

There's no silver bullet here — this is a **scale and discipline problem, not a sophistication problem**. Attackers are winning with consistency, reuse, and good operational execution.

INC isn't winning because it's clever — it's winning because organisations **still struggle to execute the fundamentals at scale**. Tightening operational discipline will do far more here than chasing the latest shiny security tool.

Actions Point

Fix the basics (properly)

- Patch **internet-facing systems** quickly (Citrix, Fortinet, remote access tools).
- Hunt for **known exploited vulnerabilities**, not just CVSS scores.

Lock down identity and backups

- Protect **backup systems (e.g. Veeam)** as Tier 0 assets.
- Enforce **MFA everywhere**, especially for admin access.
- Monitor for **credential dumping behaviour**.

Reduce lateral movement

- Limit or control use of:
 - RDP
 - PsExec
 - Remote admin tools
- Move towards **Zero Trust** / **least privilege** models.

Detect “living off the land”

- Tune detection for **LOLBins and legitimate tools being abused**.
- Focus on behaviour, not just malware signatures.

Prepare for data theft, not just encryption

- Assume **data exfiltration is standard**.
- Monitor outbound traffic and tools like Rclone.

Test your response

- Run **ransomware tabletop exercises**:
 - Can you operate without key systems?
 - Can you recover backups safely?



DragonForce Hackers Abuse Microsoft Teams Relays to Hide Backdoor.Turn C2 Traffic

<https://thehackernews.com/2026/06/dragonforce-hackers-abuse-microsoft.html>

A ransomware group called **DragonForce** is using a clever new trick to **hide its malicious traffic inside Microsoft Teams**.

They’ve developed a custom backdoor (called *Backdoor.Turn*) that routes attacker communications through legitimate Microsoft infrastructure. To defenders, it looks like **normal Teams traffic**, which makes it extremely hard to detect.

In one real incident, attackers stayed inside a corporate network for **1–2 months without being spotted**, communicating covertly the entire time.

How the attack works

Here’s the simplified version of their approach:

- **Initial access:** Likely through exploiting a vulnerable server or buying access from a broker.
- **Execution:** They drop a malicious file disguised as a “hotfix” using PowerShell.
- **Persistence and defence evasion:**
 - Use DLL sideloading to run malicious code
 - Install vulnerable drivers (BYOVD technique) to disable security tools

- **Stealthy communications:**
 - The malware connects to **Microsoft Teams infrastructure using anonymous tokens**
 - It uses Microsoft's TURN relay servers to set up communication
 - Then establishes a hidden channel to attacker systems
 - Result: defenders only see **legitimate Microsoft traffic**

Once inside, the backdoor can:

- Execute commands
- Move laterally using credentials
- Scan networks and Active Directory
- Steal browser and user credentials

Why this matters

This is a step up from “standard” ransomware tradecraft.

The key issue is **abuse of trusted platforms**. If attackers hide inside tools like Microsoft Teams:

- Traditional network monitoring becomes far less effective
- Blocking the traffic outright isn't realistic
- Attacks can persist undetected for longer

It's also a sign that DragonForce is evolving into a **more structured, capable operation**, moving beyond typical ransomware playbooks.

This is not just about malware — it's about **trust being weaponised**.

Security teams must assume:

- Legitimate SaaS platforms can be abused
- “Good” traffic may actually be malicious
- Visibility gaps in cloud and collaboration tools are now a real risk

DragonForce shows how attackers are shifting tactics — not by breaking systems, but by **hiding inside the systems we trust most**.

If INC ransomware highlighted the importance of basics, this shows the next challenge: **visibility and detection in trusted platforms**.

Action Points

Improve visibility into SaaS traffic

- Monitor **Microsoft 365 / Teams activity logs**, not just network traffic
- Look for unusual token usage, session patterns, or long-lived connections

Strengthen endpoint detection

- Ensure EDR can detect:
 - DLL sideloading
 - Suspicious PowerShell execution
 - Driver abuse (BYOVD techniques)

Control risky drivers and kernel access

- Block or monitor **known vulnerable drivers**
- Use modern protections (e.g. driver block rules, HVCI where possible)

Assume persistence and resale risk

- Attackers may keep access even after ransomware is deployed
- Treat incidents as **long-term compromise**, not one-off events

Test detection gaps

- Ask: *Would we spot malicious activity hidden inside Teams traffic?*
- Run adversary simulations that mimic trusted-service abuse

Member raised issue relating to Microsoft email environments.

Directed towards more technical readers:

1) Direct Send vs sending directly to an Exchange Online tenant

- Direct Send is a feature designed mainly for things like printers or apps inside your company to send emails to your own users without needing a login. [\[mimecastsolutions.com/endsk.com\]](https://mimecastsolutions.com/endsk.com)
- It's simple but less secure, because it doesn't require authentication and can be abused to send fake internal emails. [\[blog.admindroid.com\]](https://blog.admindroid.com)
- Sending directly to a tenant is different. It simply means anyone on the internet can send email straight to your Microsoft 365 mail system (because it has a public address, like all email systems do). [\[windowsforum.com\]](https://windowsforum.com)



Key takeaway:

- Direct Send = a specific “no-login” method mainly for internal devices.
- Direct-to-tenant sending = the normal ability for anyone to send email to your organisation.
- Confusing the two can lead to security gaps, especially if you don't lock down who is allowed to send mail into your environment. [\[linkedin.com\]](https://linkedin.com)

2) Ghost Sender (and the test site)

- Ghost Sender is a security issue caused by misconfigured email setups, especially when companies route mail through a third-party filter but don't fully restrict direct access to Microsoft 365. [\[infoguard.ch\]](https://infoguard.ch)
- In these cases, an attacker can:
 - Bypass your email security tools
 - Send emails directly to your users
 - Pretend to be anyone (internal staff or trusted external senders) [\[labs.infoguard.ch\]](https://labs.infoguard.ch)

- Even protections like SPF, DKIM, and DMARC (which are supposed to stop spoofing) can be bypassed in these setups, and the fake emails can land in inboxes without warnings. [\[blog.gridinsoft.com\]](https://blog.gridinsoft.com)
- The ghost-sender.com tool simply lets you test your own organisation to see if you are vulnerable to this problem. [\[infoguard.ch\]](https://infoguard.ch)

Key takeaway:

- It's not a single bug, but a common configuration weakness.
- If not fixed, attackers can send very convincing fake emails that look completely legitimate.

Simple overall summary

- Microsoft 365 email systems are publicly reachable by design, which is normal.
- Features like Direct Send and open inbound mail paths can be misused if not tightly controlled.
- Ghost Sender shows how poor configuration can let attackers bypass protections and spoof emails, even when standard security checks are in place.

[Direct Send vs sending directly to an Exchange Online tenant | Microsoft Community Hub](#)

Direct Send vs sending directly to an Exchange Online tenant | Microsoft Community Hub

This post discusses what Direct Send is, is not, and how to help control email sent directly to your tenant.

<https://ghost-sender.com/>

<https://labs.infoguard.ch/posts/ghost-sender/>

Test your infrastructure against Ghost-Spoof.

Action Point:

There are some fairly complex arrangements which are recommended to manage this potential exploit which include a third party and/or partnership layer. If any of you have a good solution in place - preferably without the need for additional licensing, please let us know!

June 2026 Cybersecurity Threat Update Summary

Over the past two months (Mid April–June2026), the Scottish higher and further education (HE/FE) sector has experienced a continued high volume of cyber incidents alongside a marked escalation in threat sophistication and impact. Sector-wide data indicates that cyber activity is persistent and normalised, with nearly all Universities and 88% of colleges reporting breaches or attacks over 2025-26, and over a quarter experiencing incidents at least weekly.

Operationally, the period has been characterised by sustained phishing and credential-harvesting campaigns, with human-layer compromise remaining the dominant initial access vector. These attacks increasingly enable downstream ransomware, data exfiltration, and account takeover.

Intelligence reporting highlights a rise in multi-stage intrusions and targeting of identity systems, reflecting the sector's large, transient user populations and federated access models. A significant global incident impacting UK institutions was the late April–May 2026 Canvas LMS breach, which caused service disruption and data compromise across thousands of universities, demonstrating systemic third-party risk and the sector's reliance on shared platforms.

The University of Nottingham disclosed a significant data breach affecting approximately 450,000 current and former students after attackers accessed its Oracle Campus Solutions environment and exfiltrated personal and academic data. Emerging reporting indicates the incident may be part of a broader campaign targeting Oracle PeopleSoft environments, linked to the ShinyHunters extortion group. Threat actors are exploiting the critical zero-day vulnerability CVE-2026-35273, which Oracle patched on 10 June 2026. More than 100 organisations, predominantly in the higher education sector, are believed to have been compromised through similar attacks, highlighting a systemic supply-chain risk to Oracle-based student administration systems.

Beyond the Canvas and Nottingham incidents, reported specific events (all contained) include three instances of insider threat activity, one outage and one phishing email propagation event. Indicators of compromise (IOC's) and guidance, where available and allowable, have been shared anonymously to the sector.

Threat actor intent has broadened, with ransomware, hacktivism, and state-linked espionage converging. Attack volumes have risen sharply (up 63% year-on-year), alongside increased data breach and hacktivist activity. SC3 and NCSC have both shared alerts of increased activity across the sector over the reporting period.

Action Point:

Overall, the threat landscape is intensifying, with faster exploitation cycles, AI-enabled attack techniques, and increasing focus on disruption, extortion, and strategic intelligence acquisition.

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.

Kind Regards

David



Scottish
Cyber
Coordination
Centre

[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.

David Robertson

6/25/2026

11 views



[A new unpatchable flaw in Apple chips opens the door to an iPhone jailbreak](#)


A company that sells spyware and hacking tools to government agencies has published details of a vulnerability in Apple chips that can potentially help hackers unlock older

iPhones. Main Article This release opens the door for other researchers who ...

David Robertson

6/25/2026

8 views



[More Cybersecurity Firms Disclose Impact From Klue Hack](#)

At least nine organizations have publicly acknowledged the impact of the supply chain attack on market intelligence platform Klue. Main Article LastPass The incident occurred on June 11-12 and affected Klue's integration with Salesforce, resulting i...

David Robertson

6/25/2026

9 views



[Cyber-attacks on UK universities underline persistent risk across the sector](#)

Despite ongoing efforts, cyber criminals continue to target UK universities at an alarming rate. The latest figures tell the story – research from the Department for Science, Innovation and Technology has found that 98% of higher education instituti...

David Robertson

6/25/2026

11 views



[Trump Signs Executive Order Accelerating Post-Quantum Cryptography Migration](#)

US President Donald Trump on Monday signed an executive order to strengthen data protection in preparation for the arrival of practical quantum computing. Main Article Executive Order 14409 highlights the threat posed by ‘harvest now, decrypt later...’

David Robertson

6/25/2026

13 views



[Cyber Threat Intelligence Digest – Week 25](#)

Cyber Threat Intelligence Digest – Week 25 (24 June 2026) Summary Main Article Acumen Cyber’s Week 25 threat intelligence report highlights a rapidly evolving cyber threat landscape driven by critical vulnerabilities, AI-themed malware campaigns, an...

David Robertson

6/25/2026

13 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.

Table of Contents

CISO-Share Office Weekly Newsletter	1
SC3 Daily Threat Summaries and Weekly Report.....	1
A new unpatchable flaw in Apple chips opens the door to an iPhone jailbreak.....	2
Trump Signs Executive Order Accelerating Post-Quantum Cryptography Migration.....	4
Cyber-attacks on UK universities underline persistent risk across the sector	5
The evolving nature of cyber risk.....	5
The scale of the challenge.....	5
Acumen Cyber Threat Intelligence Digest: Week 25 (.....	7

SC3 Daily Threat Summaries and Weekly Report.



Find SC3's daily threat summaries for this week here, 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-Bul SC3-Daily-Threat-Bul SC3-Daily-threat-bull SC3-Daily-threat-bull SC3-Daily-threat-bull
letin-19-June-2026.pdletin-22-June-2026.pretin-23-June-2026.pdetin-24-June-2026.pdetin-25-June-2026.pd



SC3-Weekly-Vulnera
bility-Report-22-June



A new unpatchable flaw in Apple chips opens the door to an iPhone jailbreak

A company that sells spyware and hacking tools to government agencies has published details of a vulnerability in Apple chips that can potentially help hackers unlock older iPhones.

Main Article

This release opens the door for other researchers who specialize in finding iOS vulnerabilities, such as those working for governments or their contractors, to develop effective hacks for iPhones, provided they can find additional vulnerabilities to chain together with this one. This could help security researchers develop a so-called iPhone jailbreak, a technique to hack into Apple's mobile operating system and remove all the restrictions the company puts on it.

The release is also a reminder that while Apple has made iPhones extremely hard to hack, there are and will always be vulnerabilities that sophisticated hackers can take advantage of to break in.

On Friday, Paradigm Shift, an offensive cybersecurity company based in Barcelona, published a blog post about the vulnerability, which it dubbed "usbliter8." The company also published a proof of concept that shows how to exploit the vulnerability, which requires physical access to the target phone.

The flaw and related exploit affect iPhones that have Apple-made chips A12 and A13, which were released in 2018 and 2019, and are included in older iPhones such as the XS, XR and up to the iPhone 11.

The release of usbliter8 is significant in the world of security research and spyware and hacking tools' makers, but it does not mean that older iPhones are easily hackable by anyone.

The bug found by Paradigm Shift affects the iPhone's Boot ROM, which is the first piece of code that runs when an iPhone is turned on and, consequently, its first line of defense against hackers. To hack an iPhone with physical access to it — meaning having the ability to connect a cable to it — hackers need to first exploit the Boot ROM. Now, they can do that thanks to usbliter8, which allows them to potentially defeat and bypass further security checks.

Action Point: Public iPhone jailbreaks were relatively widespread in the past, but they have become rarer in the last decade. Jailbreaking an iPhone is often the first step to research other vulnerabilities on the system. Researchers — intent on finding valuable iPhone flaws and ways to exploit them — have few incentives to release that information publicly, because that would lead to Apple fixing the flaws and setting the researchers back.

Paradigm Shift did not respond to a series of questions related to usbliter8.



More Cybersecurity Firms Disclose Impact From Klue Hack

At least nine organizations have publicly acknowledged the impact of the supply chain attack on market intelligence platform Klue.

Main Article

LastPass

The incident occurred on June 11-12 and affected Klue's integration with Salesforce, resulting in data being exfiltrated from the Salesforce instances of multiple Klue customers, including several cybersecurity firms.

To date, at least nine Klue customers have disclosed impact from the incident, including cybersecurity firms [HackerOne](#), [Huntress](#), [Jamf](#), [OneTrust](#), [Recorded Future](#), [Snyk](#), and [Tanium](#). [Insurity](#) and [Sprout Social](#) also notified their customers of the incident.

On Friday, Klue confirmed previous security reports that the attackers used compromised legacy credentials to access its systems and compromise Salesforce integrations.

"The attacker used that access to obtain OAuth tokens used to connect Klue with certain third-party platforms, including Salesforce, and subsequently accessed data within a number of connected customer environments," [Klue](#) said.

The company revoked the affected credentials and tokens, disabled the integrations across multiple services, and has been investigating the attack together with CrowdStrike and law enforcement.

Action Point:

All the affected companies pointed out that the intrusion was limited to the Salesforce instances and did not involve their systems, as Klue said in its incident notice.

Across the board, the hackers stole business information from the affected organizations' Salesforce CRMs, including sales account data and business contact information, such as names, email addresses, job titles, phone numbers, and business addresses.

Salesforce disabled the Klue integration in the wake of the incident, and revenue intelligence platform Gong did the same on Friday, warning that the hackers exploited its Klue integration to access internal licensed user data.

"We can confirm no direct impact on call recordings or customer transcripts. Examples of data accessed included user names, user business titles, and user emails," [Gong](#) said.



Trump Signs Executive Order Accelerating Post-Quantum Cryptography Migration

US President Donald Trump on Monday signed an executive order to strengthen data protection in preparation for the arrival of practical quantum computing.

Main Article

Executive Order 14409 highlights the threat posed by ‘harvest now, decrypt later’, in which threat actors exfiltrate encrypted data now, intending to crack the encryption later using quantum computers.

The private sector has been taking major steps toward post-quantum cryptography (PQC), including companies such as Google, Dell, and HP. And now the US government is also attempting to establish a clear roadmap and accelerate PQC migration.

The new EO tasks OMB, NIST, NSA, DHS, and CISA with working together to develop and oversee the implementation of comprehensive technical guidance to help federal agencies adopt PQC.

Agencies will have to inventory their high-value assets and high-impact systems and transition them to PQC for key establishment by December 31, 2030, and for digital signatures by December 31, 2031.

The order instructs agencies to designate a PQC migration lead, and the Department of Commerce will run a pilot project until the end of 2027.

Action Point: Federal contractors will also be required to comply with NIST standards regarding the use of PQC-compliant algorithms by the end of 2030.

“This executive order sends an unambiguous signal to every organization doing business with the federal government: the clock is ticking and maybe running out for some,” commented Garfield Jones, EVP Strategy and Research at QuSecure.

“The 2030 deadline for key establishment is a tangible compliance deadline, and the gap between where most organizations are today and where they need to be is significant,” Jones added. “Agencies and contractors that haven’t started a cryptographic inventory are already behind. The organizations that move now will have options. The ones that wait will find themselves managing a crisis.”

This comes shortly after President Trump signed an executive order establishing a voluntary framework for federal vetting of frontier AI models before their public release.



Cyber-attacks on UK universities underline persistent risk across the sector

Despite ongoing efforts, cyber criminals continue to target UK universities at an alarming rate. The latest figures tell the story – research from the Department for Science, Innovation and Technology has found that 98% of higher education institutions in the UK have been subject to a cyber-attack or data breach in the past year.

The question, then, is straightforward: what can be done to strengthen defences, and how can universities respond more effectively to an increasingly persistent threat?

Main Article

The evolving nature of cyber risk

Cyber incidents affecting universities vary significantly in both scale and complexity, but they tend to expose the same underlying challenges.

In practice, **attackers use a range of entry points**. In some cases, access is gained through compromised credentials or phishing activity, allowing attackers to make use of existing user permissions. In others, weaknesses in third-party platforms or integrations create opportunities to access data indirectly. Vulnerabilities in core systems can also be exploited, particularly where updates are delayed or visibility is limited.

These entry points differ in method, but the outcome is often the same: once initial access is established, it becomes significantly easier to move through the environment and reach sensitive financial, operational or personal information.

The scale of the challenge

The frequency of cyberattacks on UK universities tells part of the story, but not all of it. The real issue lies in how these institutions are leaving their technology estates open to exploitation.

Universities now operate across a mix of internal systems and external platforms that support everything from finance and student records to learning environments and assessment. While there is often some level of integration between these systems, they largely remain standalone, with **data spread across multiple platforms** and no clear, consistent view of access or activity.

As digital estates expand, so too does the surface available to exploit, particularly where responsibility for security is shared across multiple platforms and providers.

Many universities are still working with this kind of fragmented setup. Security in this environment tends to develop in stages, with new tools introduced alongside existing systems rather than as part of a single, coherent approach. The result is a collection of controls that do not always connect, making it harder to see where risks sit or how access is being used.

Findings from the [OneAdvanced Education Trends Report 2026](#) reinforce the picture. Nearly one fifth of organisations cite weak data protection measures as a barrier to improvement, while almost a quarter report a lack of regular audits or review processes.

Action Point: The shift to identity-led, integrated security

A consistent pattern is emerging across the sector. In many cases, attackers are not forcing their way into systems, but making use of access that already exists – whether through compromised credentials, poorly managed integrations, or vulnerabilities in platforms that institutions rely on as part of day-to-day operations.

Once that access is established, it becomes significantly easier to move further into the environment.

This places **identity at the centre of the problem**. Understanding who has access to which systems, and how that access is controlled, is fundamental. Measures such as Single Sign On (SSO), Multi Factor Authentication (MFA), and clearly defined permissions do not remove risk entirely, but they do limit how far an attacker can move once inside.



Acumen Cyber Threat Intelligence Digest: Week 25 (24 June 2026) Summary

A rapidly evolving threat landscape driven by critical vulnerabilities, phishing campaigns, malware-as-a-service operations, and increasing geopolitical focus on cybersecurity.

Main Article

Acumen Cyber's Week 25 threat intelligence report highlights a rapidly evolving cyber threat landscape driven by critical vulnerabilities, AI-themed malware campaigns, and growing concerns over nation-state cyber activity. Key vulnerabilities included Microsoft Defender's "RoguePlanet" privilege-escalation flaw (CVE-2026-50656), a critical Drupal Core PHP object injection vulnerability (CVE-2026-55803), and three severe vulnerabilities in NVIDIA's NeMo AI framework that could enable code execution and privilege escalation if left unpatched.

A major trend identified was the abuse of artificial intelligence branding to distribute malware. Threat actors used fake repositories offering free access to Claude and ChatGPT to deploy the Nova Stealer malware on macOS systems. Other campaigns leveraged compromised WordPress sites, blockchain-hosted command-and-control infrastructure, and "ClickFix" social-engineering techniques to distribute malware such as Vidar, DanaBot, HijackLoader, and MacSync Stealer. Several campaigns specifically targeted developers, AI researchers, and cryptocurrency users.

The report also noted warnings from the Five Eyes intelligence alliance that AI is accelerating cyberattack capabilities, reducing the time between vulnerability discovery and exploitation. Meanwhile, the UK's National Cyber Security Centre warned that hostile states are increasingly pre-positioning within critical infrastructure to enable future disruptive cyber operations.

Acumen Cyber recommends immediate patching of affected systems, strengthening access controls, and maintaining layered security defenses to mitigate these emerging threats.

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-2026-50656 (Defender)** – Microsoft released hardening updates in May 2026 which will remediate this vulnerability.
- **CVE-2026-55803 (Drupal Core)** – Drupal addressed this vulnerability in versions 11.3.12, 11.2.14, 10.6.11, and 10.5.12.
- **CVE-2026-24155, CVE-2026-24252, CVE-2026-24228 (Nemo)** – Upgrading to NVIDIA NeMo version 2.7.3 will remediate this vulnerability.

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.

Table of Contents

CISO-Share Office Weekly Newsletter	1
SC3 Daily Threat Summaries and Weekly Report.	1
FortiBleed Credential Theft Linked to INC and Lynx Ransomware Operations.....	2
Anthropic Launches Claude Fable 5: Mythos-Class AI With Cybersecurity Guardrails and other AI news	3
Security Breach at Tata Electronics Affects Apple, Tesla, and Other Technology Giants.....	4
Threat Analysis: BlueHammer Ransomware Escalation and SimpleHelp RMM Exploitation	5
Acumen Cyber Threat Intelligence Digest: Week 26 (.....	7

SC3 Daily Threat Summaries and Weekly Report.



Find SC3's daily threat summaries for this week here, 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-bull SC3-Daily-Threat-Bul SC3-Daily-Threat-Bul SC3-Daily-threat-bull SC3-Daily-threat-bull
etin-26-June-2026.pdletin-29-June-2026.pdletin-30-June-2026.pc etin-1-July-2026.pdf etin-2-July-2026.pdf



SC3-Daily-threat-bull
etin-3-July-2026.pdf



FortiBleed Credential Theft Linked to INC and Lynx Ransomware Operations

The recently discovered financially-motivated **FortiBleed** campaign has been attributed to INC and Lynx ransomware operations, indicating that the verified, stolen credentials were intended for follow-on intrusions.

Main Article

"An operator tied to FortiBleed's infrastructure was found actively working negotiation panels for both groups, tying mass FortiGate credential theft directly to ransomware deployment for the first time," SOCRadar said in a new report published Wednesday.

The company said it tracked scanning activity against approximately 11,250 FortiGate portals in more than 150 countries, followed by confirmed admin-level access on 409 targets and successful completion of the full attack chain on 354 of them. In all, at least 12 ransomware deployments have resulted from this access, causing hundreds of endpoints to be encrypted across affected organizations.

Action Point: The disclosure comes as eSentire said it observed threat actors exploiting a flaw in Fortinet FortiClient EMS ([CVE-2026-35616](#), CVSS score: 9.1) to deploy an information stealer called **EKZ Stealer** against a customer in the energy, utilities, and waste sector with the end goal of harvesting credentials from Chromium-based browsers and Firefox and exfiltrating them via PowerShell.



Anthropic Launches Claude Fable 5: Mythos-Class AI With Cybersecurity Guardrails and other AI news

Anthropic on Tuesday announced the general availability of **Claude Fable 5**, a powerful Mythos-class AI model engineered with new safeguards that specifically restrict its use in high-risk domains, including cybersecurity.

[Main Article](#)

[New Claude Code Attack](#)

[Chat GPT Update](#)

The AI giant says this marks the first time a model of this capability class has been deemed safe enough for widespread public and developer access.

While **Fable 5** demonstrates good performance — surpassing prior models in software engineering, knowledge work, vision, and long-running tasks — the company prioritized safety by implementing targeted blocks.

In sensitive areas such as cybersecurity and biology, the model automatically falls back to the less capable Claude Opus 4.8 to prevent potential misuse. Early usage data indicates that at least 95% of sessions run entirely on Fable 5's capabilities without triggering any fallback.

“The uplift from Mythos-level capabilities is valuable to many adversaries — for instance, those who could financially gain from cyberattacks — and we therefore expect them to be motivated to try to circumvent our safety measures,” Anthropic noted.

The company emphasized the rigor of its safety measures. It conducted extensive internal red-teaming of its classifiers, followed by an external bug bounty program spanning over 1,000 hours that yielded no universal jailbreaks.

Action Point: AI is developing ever faster and not always used for the common good

The AI giant has not listed the new additions, but several cybersecurity and tech companies have since announced their participation in the project, including Dragos, Tenable, TrendAI (Trend Micro), Netskope, BeyondTrust, Rubrik, BT, Intercontinental Exchange, and Hitachi.

Both Fable 5 and Mythos 5 are priced at \$10 per million input tokens and \$50 per million output tokens, with the former available immediately via the Claude API for developers.



Security Breach at Tata Electronics Affects Apple, Tesla, and Other Technology Giants

Tata Electronics is investigating a security breach following claims by the World Leaks hacking group that it had stolen company data and listed it for sale on an underground hacking forum.

Main Article

Tata Electronics accounts for nearly a third of all iPhone production in India, while Foxconn makes up the rest. It supplies circuit boards, back panels, and enclosures.

The company confirmed the data breach nearly two weeks after the threat actor listed the stolen data for sale on the dark web.

“A few weeks ago, Tata Electronics identified a cybersecurity incident on some of our systems,” the company told Reuters.

Tata security breach leaked proprietary data from Apple, Tesla, and other companies

The 630.4 GB archive reportedly contained about 204,341 files, including Apple’s and Tesla’s component designs and specifications, including clearly marked confidential proprietary data. It also includes emails, employee passports, and event logs, which could expose the IT infrastructure to more attacks.

Since June 10, the stolen data was available for free download. Its exposure could enable cybercriminals to discover security vulnerabilities and competitors and counterfeiters to reverse-engineer products.

The leaked folders contain hundreds of references to Apple and Tesla. The security breach reportedly also affected Pegatron, Foxconn, and Qualcomm, though those claims have not been independently confirmed.

“What stands out here is not just the volume of data, but the amount of strategic information concentrated inside one supplier,” said Michael Centrella, Head of Public Policy at [SecurityScorecard](#). “Tata Electronics sits inside the operations of companies such as Apple and Tesla, so a breach can expose far more than ordinary corporate records.”

Action Point: Tata has experienced security breaches in the past. In September 2025, Tata-owned British automotive company Jaguar Land Rover suffered a ransomware attack that disrupted operations for 6 weeks.

In May 2026, Foxconn confirmed a cyber attack that leaked 8 terabytes of data from major technology companies, including Intel, Apple, Google, Dell, Nvidia, and others. The ransomware attack disrupted operations, forcing the company to suspend operations and order employees to stay at home.

In December 2025, Apple, Nvidia, Qualcomm, LG, Meta, Geely, Tesla, and other tech giants reportedly experienced a data breach after the Russian ransomware group RansomHouse compromised a third-party supplier, Luxshare. The cyber attack leaked proprietary data, including CAD models, circuit board designs, and engineering documentation.



Threat Analysis: BlueHammer Ransomware Escalation and SimpleHelp RMM Exploitation

Although this article is aimed at SMB's, it also relates to University and College supply chains and any outsourced IT support.

[Main Article](#)

BlueHammer: From Zero-Day to Ransomware Commodity

CISA added the BlueHammer vulnerability to its Known Exploited Vulnerabilities catalogue today, confirming that ransomware gangs are actively using it in the wild. BlueHammer is a local privilege escalation flaw in Microsoft Defender, the security software that comes installed on virtually every Windows machine in the country.

Here is what the advisory does not say plainly enough: this vulnerability was already being exploited before today's confirmation. It was used in targeted zero-day attacks. CISA's KEV addition means it has now crossed into commodity ransomware operations. That transition, from targeted to commodity, is the inflection point that matters. It means the attack no longer requires sophisticated operators. It means volume increases. It means UK SMBs are now squarely in the blast radius.

The mechanics are worth understanding briefly. A local privilege escalation vulnerability does not give an attacker initial access to your system. What it does is allow someone who has already obtained a foothold, through phishing, a compromised credential, a vulnerable remote access tool, to escalate their permissions to SYSTEM level. From there, they can disable security software, move laterally across your network, and deploy ransomware. It is the step between 'we are in' and 'we own everything.'

The patch for BlueHammer has been available. The question is whether it has been applied. If your endpoints are managed by an IT provider, ask them directly. If you manage your own Windows estate, check your patch status today. There is no nuance here: patch it.

The advisory attributes active exploitation to the group operating under the 'Nightmare Eclipse' campaign identifier. What the advisory does not say is that the ransomware-as-a-service landscape has expanded significantly in the first half of 2026. Securelist's recent analysis of the Gentlemen RaaS group, which emerged as a top-ten threat actor in H1 2026, shows a pattern of exploiting internet-facing devices before deploying privilege escalation tools for lateral movement. BlueHammer fits neatly into that playbook. The attribution to one group should not give anyone comfort that only one group is using it.

Action Point: On BlueHammer:

- If your IT provider manages your Windows endpoints, ask them to confirm that the BlueHammer patch has been deployed across all managed devices
- If you manage your own systems, check Windows Update status and verify the Defender patch is applied
- Do not wait for your provider to contact you. Ask proactively.

On SimpleHelp [CVE-2026-48558](#):

- Ask your IT provider whether they use SimpleHelp in their tooling
- If they do, ask specifically whether [CVE-2026-48558](#) has been patched and when
- Ask whether they have reviewed their RMM access logs for anomalous activity since the vulnerability was disclosed
- If they cannot answer these questions promptly and specifically, that is information worth having

On the broader pattern:

- Ensure your IT provider has a documented process for disclosing security incidents that affect their tooling to clients
- If no such process exists or they cannot describe it, that is a contractual and operational gap worth addressing



Acumen Cyber Threat Intelligence Digest: Week 26 (02 July 2026) Summary

A rapidly evolving threat landscape driven by critical vulnerabilities, phishing campaigns, malware-as-a-service operations, and increasing geopolitical focus on cybersecurity.

Main Article

The Cyber Threat Intelligence Digest highlights several critical vulnerabilities, emerging attack techniques, and notable cybersecurity developments observed during the reporting period. Google disclosed a critical vulnerability (CVE-2026-12537) affecting Gemini CLI that could enable host-level remote code execution in CI/CD environments, while Synology and Dell released patches for multiple high-severity vulnerabilities impacting MailPlus Server and Wyse Management Suite respectively. Organisations are advised to prioritise these updates to reduce exposure.

The report identifies several active threat campaigns, including callback phishing attacks abusing Shopify's Shop application through fake purchase receipts impersonating trusted brands to steal credentials and financial information. Another significant campaign targeted hospitality organisations across Europe and Asia using phishing emails that leveraged Calendly notifications and Google URL redirection to deliver the TonRAT malware, demonstrating increasingly sophisticated evasion techniques. Separately, LastPass disclosed a third-party breach after attackers compromised OAuth tokens associated with the Klue platform, exposing customer CRM information but not password vault data.

Broader industry developments include Microsoft's accelerated transition towards post-quantum cryptography, WhatsApp's introduction of usernames to enhance user privacy, and the CIA's renewed focus on AI and cybersecurity capabilities. The digest also highlights evolving threat actor activity, with ShinyHunters remaining a moderate threat and several new actors emerging. Global threat trends show increased activity from groups including RedDelta, Anonymous, and Akira Ransomware, alongside growing exploitation of newly disclosed vulnerabilities and advanced malware families. The report concludes that timely patching, phishing awareness, identity protection, and proactive threat monitoring remain essential for defending against an increasingly sophisticated cyber threat landscape.

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-2026-12537 (Gemini CLI)** – This vulnerability can be remediated by updating to Gemini CLI to version 0.39.1 or 0.40.0-preview.3, and the run-gemini-cli GitHub Action to version 0.1.22 or later.
- **CVE-2026-13136, CVE-2026-13135, CVE-2025-15660 (Synology)** – These vulnerabilities can be addressed by patching to DSM versions 4.0.1-31663 and 4.0.1-21663.
- **CVE-2026-41120, CVE-2026-49506 (Dell)** – Updating Wyse Management Suite to version 5.5 HF1 or later will remediate this vulnerability.

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

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.

Kind Regards

David



Scottish
Cyber
Coordination
Centre

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

Find SC3's daily threat summaries for this week here, for those who do not receive this information directly. Action Point: All Daily, Weekly and Monthly Reports:

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

David Robertson

7/3/2026




[Anthropic Launches Claude Fable 5: Mythos-Class AI With Cybersecurity Guardrails and other AI news](#)

Anthropic on Tuesday announced the general availability of Claude Fable 5, a powerful Mythos-class AI model engineered with new safeguards that specifically restrict its use in high-risk domains, including cybersecurity. Main Article New Claude Code...

David Robertson

7/2/2026

9 views



[FortiBleed Credential Theft Linked to INC and Lynx Ransomware Operations](#)

The recently discovered financially-motivated FortiBleed campaign has been attributed to INC and Lynx ransomware operations, indicating that the verified, stolen credentials were intended for follow-on intrusions. Main Article "An operator tied to F...

David Robertson

7/2/2026

6 views



[Security Breach at Tata Electronics Affects Apple, Tesla, and Other Technology Giants](#)

Tata Electronics is investigating a security breach following claims by the World Leaks hacking group that it had stolen company data and listed it for sale on an underground hacking forum. Main Article Tata Electronics accounts for nearly a third o...

David Robertson

7/2/2026

11 views



[Threat Analysis: BlueHammer Ransomware Escalation and SimpleHelp RMM Exploitation](#)

Although this article is aimed at SMB's, it also relates to University and College supply chains and any outsourced IT support. Main Article BlueHammer: From Zero-Day to Ransomware Commodity CISA added the BlueHammer vulnerability to its Known Explo...

David Robertson

7/2/2026

9 views



[Acumen: Cyber Threat Intelligence Digest: Week 26](#)

Highlights of Cyber Threat Intelligence Digest
Main Article The Cyber Threat Intelligence Digest highlights several critical vulnerabilities, emerging attack techniques,

and notable cybersecurity developments
observed during the reporting period. Go...

David Robertson

7/2/2026

9 views

[Go To Site](#)



CISO Share

ThreatScape Newsletter

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

While no specific, new cybersecurity incidents have been publicly disclosed for UK Higher Education (HE) and Further Education (FE) institutions in the past 7 days, the sector is currently navigating the fallout from the massive cyberattack on the University of Nottingham.

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.



Scottish
Cyber
Coordination
Centre

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

Find SC3's daily threat summaries for this week here, for those who do not receive this information directly. Action Point: All Daily, Weekly and Monthly Reports:

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

David Robertson

7/9/2026

4 views



[China-Aligned Hackers Exploit Roundcube Servers at Universities](#)

Proofpoint has identified a suspected China-aligned threat actor exploiting vulnerable Roundcube servers at U.S. and Canadian universities to steal credentials and gain persistent access. Rather than targeting email content alone, the campaign appea...

David Robertson

7/10/2026

4 views



[CISA Reportedly Using Anthropic's Mythos to Scan Government Software for Flaws](#)

The US Cybersecurity and Infrastructure Security Agency (CISA) is using Anthropic's powerful Mythos AI model to scan and audit federal government software for security vulnerabilities, according to a report from Reuters. Main Article Citing three ...

David Robertson

7/9/2026

6 views



[Attackers exploit critical Adobe ColdFusion vulnerability \(CVE-2026-48282\)](#)

CVE-2026-48282, one of the maximum severity vulnerabilities patched in Adobe ColdFusion on June 30, 2026, has been targeted by attackers in the wild. Main Article

Exploitation attempts were detected on July 2, through the honeypot sensors of cyberse...

David Robertson

7/9/2026

4 views



[Januscape Linux VM Escape Flaw Affects Intel and AMD Systems](#)

A newly disclosed Linux kernel vulnerability can be exploited to escape virtual machines (VMs) and execute code on the underlying host, security researchers warn. Main Article Tracked as CVE-2026-53359 and referred to as Januscape , the security def...

David Robertson

7/9/2026

5 views



[Acumen Cyber Threat Intelligence Digest: Week 27 \(09 July 2026\) Summary](#)

A rapidly evolving threat landscape driven by critical vulnerabilities, phishing campaigns, malware-as-a-service operations, and increasing geopolitical focus on cybersecurity. Main Article This cyber threat intelligence report highlights several im...

David Robertson

7/9/2026

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.

While no specific, new cybersecurity incidents have been publicly disclosed for UK Higher Education (HE) and Further Education (FE) institutions in the past 7 days, the sector is currently navigating the fallout from the massive cyberattack on the University of Nottingham.

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.

Table of Contents

CISO-Share Office Weekly Newsletter	1
SC3 Daily Threat Summaries and Weekly Report.....	1
China-Aligned Hackers Exploit Roundcube Servers at Universities	2
CISA Reportedly Using Anthropic's Mythos to Scan Government Software for Flaws	3
Attackers exploit critical Adobe ColdFusion vulnerability (CVE-2026-48282).....	4
Januscape Linux VM Escape Flaw Affects Intel and AMD Systems	5
Acumen Cyber Threat Intelligence Digest: Week 27 (.....	6

SC3 Daily Threat Summaries and Weekly Report.



Find SC3's daily threat summaries for this week here, 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-Bul
letin-7-July-2026.pdf



SC3-Daily-Threat-Bul
letin-6-July-2026.pdf



SC3-Weekly-Vulnera
bility-Report-6-July-2026



SC3 - Monthly
Ransomware Report -



China-Aligned Hackers Exploit Roundcube Servers at Universities

Proofpoint has identified a suspected China-aligned threat actor exploiting vulnerable Roundcube servers at U.S. and Canadian universities to steal credentials and gain persistent access.

Rather than targeting email content alone, the campaign appears designed to use compromised mail servers as an entry point into broader institutional networks.

Main Article

Proofpoint tracks the activity cluster as UNK_MassTraction, a temporary designation used for emerging threat clusters that have not yet been assigned a formal threat actor identifier.

Since May 2026, the group has targeted physics and engineering departments, particularly those involved in astrophysics, particle physics, and research with national security connections.

Action Point: How Roundcube vulnerabilities enabled multi-stage attacks

The campaign exploits multiple previously disclosed Roundcube vulnerabilities to compromise webmail servers.

The attack begins with [CVE-2024-42009](#), a cross-site scripting (XSS) flaw that allows malicious JavaScript to execute when a victim simply opens a specially crafted email in a vulnerable Roundcube webmail client.

Once executed, the JavaScript loader deploys a credential-stealing payload that Proofpoint calls IceCube.

The malware collects usernames, passwords, session cookies, two-factor authentication material, and browser reconnaissance data before transmitting the information to command-and-control (C2) infrastructure.

IceCube then leverages a second vulnerability, [CVE-2025-49113](#), to gain server-side access by exploiting a deserialization flaw.

Depending on the environment, the malware either installs a webshell known as SquareShell or loads the VShell backdoor directly into memory, providing attackers with persistent access while minimizing forensic evidence.

Proofpoint noted that the malware appears intentionally engineered to evade detection by cleaning up browser artifacts, monitoring user activity, retrying failed exploitation attempts, and destroying active sessions after compromise.



CISA Reportedly Using Anthropic's Mythos to Scan Government Software for Flaws

The US Cybersecurity and Infrastructure Security Agency ([CISA](#)) is using Anthropic's powerful Mythos AI model to scan and audit federal government software for security vulnerabilities, according to a report from Reuters.

Main Article

Citing three sources familiar with the matter, Reuters reported that CISA is utilizing Mythos to scan code repositories across federal agencies. The operation aims to proactively discover and patch security bugs that could otherwise be exploited by foreign intelligence agencies and cybercriminals.

The audits are reportedly being spearheaded by CISA's Attack Surface Evaluation team, a specialized unit tasked with conducting digital defense assessments and simulated hacking exercises across the federal landscape. Two sources stated that the AI-driven initiative has already uncovered a "large number" of software vulnerabilities. However, specific details regarding the severity of the flaws, the impacted agencies, or the volume of software reviewed have not been disclosed.

Neither Anthropic nor CISA provided formal on-the-record comments to Reuters regarding the operation.

Tensions between Anthropic and federal officials spiked dramatically earlier this year after the company refused administration demands to remove built-in safeguards restricting its models from being used for autonomous weaponry or domestic surveillance. In response, the Pentagon designated Anthropic as a supply-chain risk, a classification typically reserved for foreign firms suspected of espionage.

The National Security Agency (NSA) is also believed to be using Mythos in its operations.

Action Point:

Late last month, a US official told the Associated Press (AP) that one of Anthropic's artificial intelligence models had identified vulnerabilities in highly sensitive and secure US government computer systems during a testing exercise.

While the private application of Mythos has accelerated within the US intelligence and defense communities, Anthropic's public-facing rollouts have triggered separate regulatory battles. When the company launched its public version of the model in early June, called Fable, concerns from the White House regarding foreign nationals accessing the tool prompted an abrupt administrative demand to restrict access. The ensuing standoff led to a temporary global shutdown of the Fable model, which was only lifted last week.



Attackers exploit critical Adobe ColdFusion vulnerability (CVE-2026-48282)

CVE-2026-48282, one of the maximum severity vulnerabilities patched in Adobe ColdFusion on June 30, 2026, has been targeted by attackers in the wild.

Main Article

Exploitation attempts were detected on July 2, through the honeypot sensors of cybersecurity threat-intelligence service KEVIntel, mere minutes after watchTower researchers published a technical analysis of this and other ColdFusion flaws recently fixed by Adobe.

What makes CVE-2026-48282 dangerous

Adobe ColdFusion is a widely used development platform for building and deploying enterprise-grade websites and web applications. It's usually run on Windows or Linux servers.

Seven days ago, Adobe disclosed that it had released patches for CVE-2026-48282 and nine other critical vulnerabilities affecting ColdFusion.

CVE-2026-48282's description, CVSS score and vector indicate that this path traversal vulnerability may allow remote, unauthenticated attackers to achieve arbitrary code execution by sending a specially crafted HTTP request to upload a malicious file to a web-accessible location.

"The attacker accesses the uploaded file directly via the web server, triggering execution of arbitrary code in the context of the current user, and can then escalate to further compromise the host," the Centre for Cybersecurity Belgium noted.

WatchTower's write-up explained that the vulnerability is found in ColdFusion's Remote Development Services (RDS) feature "that allows a developer's IDE, historically ColdFusion Builder, Dreamweaver, or the Eclipse plugin, to interact with a running ColdFusion server."

According to them, the IDE "can browse the filesystem, execute database queries, and assist with debugging, all over HTTP."

Resecurity researchers said they are also tracking the exploitation of this vulnerability, and they have released additional details that can help both defenders and attackers.

Action Point: AI is developing ever faster and not always used for the common good

Vulnerabilities in Adobe ColdFusion are regularly exploited by attackers.

To successfully leverage CVE-2026-48282, attackers must target ColdFusion servers on which RDS is enabled (and it's not, by default), and authentication for it is disabled.

The Shadowserver Foundation is currently tracking around 750 internet-facing ColdFusion servers, but it's unknown how many of these still run a vulnerable ColdFusion version or whether they have RDS enabled.

Admins are advised to upgrade to ColdFusion 2025 update 10 or ColdFusion 2023 Update 21, and if their servers are or were internet-facing in the last week, to hunt for indicators of compromise such as unauthorized files within ColdFusion's web root and /CFIDE/ directories.

UPDATE (July 8, 2026, 07:20 a.m. ET): The US Cybersecurity and Infrastructure Security Agency has added CVE-2026-48282 to its Known Exploited Vulnerabilities catalog and ordered US federal civilian agencies to address the issue by July 10, 2026.



Januscape Linux VM Escape Flaw Affects Intel and AMD Systems

A newly disclosed Linux kernel vulnerability can be exploited to escape virtual machines (VMs) and execute code on the underlying host, security researchers warn.

Main Article

Tracked as [CVE-2026-53359](#) and referred to as [Januscape](#), the security defect impacts the shadow MMU code in Linux Kernel-based Virtual Machine (KVM) hypervisor.

Although Januscape requires specific conditions for exploitation, it highlights how vulnerabilities in foundational infrastructure can affect far more than a single workload.

With cloud and virtualized environments continuing to expand, maintaining visibility into hypervisor technologies and incorporating them into vulnerability prioritization efforts can help reduce enterprise-wide risk.

Adopting a [Zero Trust](#) architecture can help further reduce risk by continuously verifying users, devices, workloads, and access requests across critical infrastructure.

Action Point: Key Takeaways

- Januscape (CVE-2026-53359) is a Linux KVM vulnerability that enables guest-to-host VM escape on both Intel and AMD systems.
- Successful exploitation could allow attackers with root access in a guest VM to execute code on the host or disrupt other virtual machines sharing the same server.
- The vulnerability poses the greatest risk to multi-tenant cloud and virtualized environments that rely on KVM for workload isolation.



Acumen Cyber Threat Intelligence Digest: Week 27 (09 July 2026) Summary

A rapidly evolving threat landscape driven by critical vulnerabilities, phishing campaigns, malware-as-a-service operations, and increasing geopolitical focus on cybersecurity.

Main Article

This cyber threat intelligence report highlights several important security risks and developments from early July 2026. Several software vendors, including WatchGuard, Apache and Seiko Solutions, released security updates to fix critical vulnerabilities. Organisations using affected products should apply these patches promptly to reduce the risk of attack.

A reported campaign called "FortiBleed" allegedly targeted UK government organisations by exploiting Fortinet firewalls and using stolen credentials to access email accounts. The UK National Cyber Security Centre has advised organisations to review firewall security, isolate compromised systems and strengthen credential management.

Researchers also uncovered a campaign distributing malicious Python packages through fake GitHub proof-of-concept repositories. The malware, known as ChocoPoC RAT, targets security researchers and developers, stealing credentials, documents and other sensitive information.

Another significant threat involves ARToken, a phishing-as-a-service platform targeting Microsoft 365 users. Rather than stealing passwords, it captures authentication tokens, allowing attackers to gain long-term access to accounts, even after passwords are changed.

The report also notes wider cybersecurity developments, including the UK's plans to develop an AI-powered "Cyber Shield" to strengthen national cyber defence, an investigation into Pegasus spyware targeting a former European Parliament member, and an ongoing cybersecurity incident affecting the US Department of Homeland Security's unclassified information-sharing systems. Overall, organisations should prioritise patching, monitor emerging threats, strengthen identity security and remain vigilant against phishing and supply chain attacks.

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-2026-13368 (WatchGuard) – This vulnerability can be remediated by updating the 2025.1 branch to Fireware OS 2026.2.1, and the 12.x branch to 12.12.1.

CVE-2026-54428, CVE-2026-54399 (Apache) – These vulnerabilities can be addressed by patching to Apache HttpComponents Core version 5.4.3 or later.

CVE-2026-50043 (Seiko) – Updating SkyBridge to the most recent version will remediate this vulnerability.

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.

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.



Cyber Threat Intelligence Digest – Week 28

<https://acumencyber.com/cyber-threat-intelligence-digest-july-2026-week-28>

This week's Acumen Cyber report is a good reminder that cyber threats continue to be driven by a mix of **unpatched vulnerabilities, increasingly sophisticated phishing campaigns, and the theft of sensitive data**. None of the individual stories are revolutionary, but together they highlight how attackers are consistently exploiting known weaknesses and human behaviour. For security leaders, the message is simple: the fundamentals still matter.

Vulnerabilities continue to create opportunities

Several high-severity vulnerabilities were disclosed and patched during the week.

IBM released fixes for a vulnerability within IBM i operating systems that could allow an attacker to execute code with administrative privileges. Apache also patched a critical flaw affecting Apache CXF that could potentially lead to remote code execution, while Canon addressed multiple vulnerabilities in its EOS Network Setting Tool.

The good news is that none of these vulnerabilities are known to be actively exploited at present. The bad news is that threat actors often move quickly once technical details become public, meaning patching windows can close rapidly.

Social engineering remains the attacker's weapon of choice

The report highlights several campaigns that rely heavily on tricking users rather than exploiting technical weaknesses.

One of the most notable is **ClickFix**, a growing social engineering technique where users are presented with fake verification or error messages and then instructed to execute commands themselves. In the campaigns observed this week, attackers impersonated popular AI tools such as Claude and Codex, directing victims to convincing fake installation pages hosted on legitimate Google Sites infrastructure.

Once users followed the instructions, malware was executed entirely in memory, stealing browser credentials, email data and cryptocurrency wallet information.

The lesson here is clear: attackers are increasingly persuading users to become part of the attack chain.

Mobile banking customers remain a target

Researchers also highlighted the continued evolution of **NFCShare**, an Android malware family designed to steal payment card details.

Victims are directed to phishing pages impersonating well-known banks and encouraged to install a fake mobile application. The malware uses the phone's NFC capability to read payment card information and capture PIN details before transmitting the data back to attackers.

While this threat primarily targets consumers, it reinforces a broader trend of mobile devices becoming an increasingly attractive attack surface.

Education sector breach highlights data protection challenges

The cyber incident affecting the **University of Nottingham** is particularly relevant to those working within Higher and Further Education.

The university confirmed that data belonging to current and former students was accessed by an unauthorised third party. The threat actor group ShinyHunters claims to have obtained a significant volume of information, potentially including personal and financial data.

While the exact scope is still being investigated, the incident serves as another reminder of the value of educational data and the reputational, regulatory and operational consequences that can result from a breach.

Key themes emerging

Across all the stories, several common themes stand out:

- Attackers continue to exploit trusted platforms such as Google Sites and GitHub.

- Social engineering remains highly effective.
- Data theft is often the primary objective.
- Threat actors are refining existing techniques rather than creating entirely new ones.
- Universities and other data-rich organisations remain attractive targets.

This week's intelligence doesn't point to a single critical threat, instead, it reinforces a familiar reality: organisations are still most likely to be compromised through a combination of **unpatched systems, successful phishing attacks and inadequate visibility over data movement**.

For security teams, success continues to come down to disciplined execution of the basics rather than chasing the latest cyber trend.

Action Points

Prioritise patch management

- Apply updates for IBM, Apache and Canon products where relevant.
- Focus patching efforts on externally accessible services and critical systems.

Increase user awareness

- Educate staff about ClickFix-style attacks.
- Reinforce that users should never run PowerShell or command-line instructions provided by websites or unsolicited emails.

Improve monitoring of trusted platforms

- Review controls around access to GitHub, Google Sites and other commonly abused services.
- Monitor for unusual downloads and script execution activity.

Strengthen endpoint detection

- Ensure EDR solutions can detect:
 - In-memory malware execution
 - Suspicious PowerShell activity
 - Credential theft attempts

Enhance data protection controls

- Review data classification and retention practices.
- Monitor for unusual outbound data transfers.
- Validate breach response and communications plans.

Review risks to mobile users

- Remind users not to install applications from unofficial sources.
 - Raise awareness of banking-themed phishing attacks and fake mobile apps.
-



Microsoft Fixes Record 570 Security Flaws in Biggest Patch Tuesday Ever

<https://www.techrepublic.com/article/news-microsoft-july-2026-patch-tuesday/>

Microsoft's July 2026 Patch Tuesday wasn't just another monthly update—it was the largest security release in the company's history.

Microsoft fixed a record **570 vulnerabilities**, including **three zero-days**, two of which were already being actively exploited before patches became available. The release also included **59 critical vulnerabilities**, many capable of enabling remote code execution.

Why this matters

The scale of this release is significant, but the bigger story is what it says about today's threat landscape.

Microsoft has now patched more than **1,300 vulnerabilities during the first seven months of 2026**, almost double the same period last year. According to Microsoft, this increase is partly being driven by AI-assisted vulnerability discovery, enabling security researchers and vendors to identify flaws faster than ever before.

While finding more vulnerabilities is ultimately positive, it means organisations now face an increasing volume of security updates and a shrinking timeframe in which to apply them.

The most urgent vulnerabilities

Security teams should focus first on the two vulnerabilities that were already under active attack.

Active Directory Federation Services (AD FS)

CVE-2026-56155 allows an authenticated attacker to elevate privileges to administrator level.

For organisations relying on AD FS for identity and authentication services, this creates a significant risk because successful exploitation could provide broad control over identity infrastructure.

Microsoft SharePoint Server

CVE-2026-56164 affects Microsoft SharePoint Server and allows attackers to elevate privileges due to missing authentication controls within a critical function.

Given the widespread use of SharePoint for collaboration and document management, this vulnerability should be considered high priority—particularly for internet-facing environments.

Microsoft also recommended additional mitigations for organisations unable to patch immediately.

BitLocker Security Bypass

Microsoft also fixed **CVE-2026-50661**, a vulnerability affecting BitLocker.

The issue could potentially allow someone with physical access to a device to bypass protections and gain access to encrypted information.

The wider impact

Beyond the headline vulnerabilities, the update addresses critical flaws across a broad range of Microsoft technologies, including:

- Windows
- SharePoint
- Exchange Server
- Microsoft Defender
- Dynamics
- Media Foundation
- Various enterprise services and components

For many organisations, this means a significant proportion of the Microsoft estate may require patching and testing.

AI is changing the game

One of the more interesting themes emerging from this release is Microsoft's acknowledgement that AI is now influencing both attackers and defenders.

Attackers are increasingly using AI to identify weaknesses more quickly, while Microsoft is using AI-driven systems to discover and remediate vulnerabilities before they can be exploited.

The result is an accelerating race between vulnerability discovery, disclosure, patching and exploitation.

Microsoft has already advised organisations to reduce the amount of time they defer Windows updates because the gap between disclosure and active exploitation continues to shrink.

What this means

The traditional approach of waiting several weeks before deploying critical updates is becoming harder to defend.

Historically, organisations have balanced patching speed against operational stability and change management requirements. While testing remains essential, attackers are becoming much faster at weaponising newly disclosed vulnerabilities.

The challenge is no longer simply patching vulnerable systems.
The challenge is **patching them quickly enough**.

This is particularly important for organisations running:

- SharePoint
- Exchange
- AD FS
- Internet-facing Windows services
- Identity platforms

These remain attractive targets because they often provide direct access to business-critical systems and data.

This month's Patch Tuesday is about more than a record-breaking **570 vulnerabilities**.

It highlights a broader shift in cyber security: AI is accelerating vulnerability discovery, attackers are moving faster to exploit weaknesses, and organisations have less time than ever to respond.

The message is clear: **patch management is no longer just an IT process—it is a critical business capability**. Organisations that can identify, prioritise and deploy fixes quickly will be far better placed to manage the risks created by this rapidly evolving threat landscape.

Action Points

Prioritise the actively exploited vulnerabilities

Focus immediate attention on:

- CVE-2026-56155 (AD FS)
- CVE-2026-56164 (SharePoint)

Any vulnerability already under active exploitation should move straight to the top of the patching queue.

Review internet-facing systems

Identify externally exposed services, particularly:

- SharePoint
- Exchange
- Authentication platforms
- Remote access services

These systems are most likely to be targeted first.

Accelerate patch deployment timelines

Review current patching SLAs and determine whether they remain appropriate given the shrinking window between disclosure and exploitation.

Improve vulnerability prioritisation

Don't rely solely on CVSS scores.

Consider:

- Active exploitation status
- External exposure

- Business criticality
- Availability of mitigating controls

Implement mitigations where patching isn't immediate

If operational constraints delay deployment, ensure recommended compensating controls are applied as quickly as possible.

Plan for increased patch volumes

Security teams should expect vulnerability volumes to continue increasing as AI-assisted discovery becomes more common.

Ensure there is sufficient capacity for:

- Patch testing
- Change management
- Deployment
- Validation
- Rollback planning

Review patch management governance

This release demonstrates that patch management is no longer a routine operational task. It has become a core security capability that requires executive visibility and support.



Identity Attacks Overtake Exploits as Top Ransomware Cause

<https://www.darkreading.com/identity-access-management-security/identity-attacks-overtake-exploits-top-ransomware-cause>

For years, cyber security teams have focused heavily on patching vulnerabilities to reduce ransomware risk. While patching remains critical, new research suggests attackers have changed tactics.

The latest findings show that **identity attacks have now overtaken vulnerability exploitation as the leading cause of ransomware incidents**. Put simply, attackers are increasingly logging in rather than hacking in.

What's changed?

For the first time in several years, exploited vulnerabilities are no longer the primary cause of ransomware attacks.

The leading causes are now:

- Malicious email – 26%
- Phishing – 24%

- Compromised credentials – 23%
- Exploited vulnerabilities – 18%

Combined, email, phishing and credential theft account for the vast majority of ransomware incidents.

This marks a significant shift away from technical exploitation and towards identity compromise as the preferred route into organisations.

Attackers are targeting identities

One of the most striking findings is that nearly **80% of ransomware attacks now originate from identity-related compromises**.

In addition, around two-thirds of victims reported that their ransomware incident was also the most significant identity attack they experienced during the year.

This reinforces a trend many security teams have witnessed firsthand: cyber criminals increasingly prefer to abuse legitimate accounts rather than exploit software flaws.

If an attacker can obtain valid credentials, they can often bypass many of the traditional controls designed to detect hacking attempts.

The MFA surprise

Perhaps the most interesting finding is the role of multi-factor authentication (MFA).

Among organisations where compromised credentials were identified as the root cause of ransomware, **97% had MFA deployed in some form**.

At first glance, this appears worrying. If MFA is in place, why are attacks still succeeding?

The likely answer is that MFA coverage is incomplete. Many organisations have strong protection around Microsoft 365 and cloud services but still have gaps in areas such as:

- VPNs
- Legacy applications
- Firewall administration portals
- Service accounts
- Third-party systems

Attackers only need to find one weakness in identity controls to gain access.

Why this matters

Many organisations have invested heavily in improving vulnerability management, reducing patching times and addressing externally exposed weaknesses.

Whilst that work has undoubtedly reduced risk, attackers have adapted.

Rather than spending time developing exploits, it is often easier to:

- Steal credentials through phishing
- Buy credentials from criminal marketplaces
- Exploit weak authentication controls

- Abuse trusted user accounts

From an attacker's perspective, using a legitimate account is often more reliable and significantly less noisy than exploiting a vulnerability.

What this means

The message from this report is simple:

Identity has become the new security perimeter.

Historically, organisations built security programmes around networks, endpoints and perimeter defences. Increasingly, attackers are bypassing those controls altogether by targeting identities.

This doesn't mean vulnerability management becomes less important. Attackers still exploit vulnerabilities and unpatched systems. However, patching alone is no longer sufficient to prevent ransomware.

Security leaders need to treat identity security as a strategic priority rather than simply another control domain.

The most important takeaway is that ransomware is increasingly becoming an **identity problem rather than a vulnerability problem**.

Attackers aren't always breaking in through sophisticated exploits. More often, they're persuading users to hand over credentials, abusing weak identity controls, or finding gaps in MFA deployment and simply logging in.

The strategic lesson is clear: **if identity is the new perimeter, identity security needs to sit alongside patching and vulnerability management as a top security priority.**

Action Points

Review MFA coverage

Don't just ask whether MFA is enabled.

Ask:

- Is MFA applied to every critical service?
- Are administrative accounts protected?
- Are legacy systems covered?
- Are users relying on phishing-resistant authentication methods where appropriate?

Focus on identifying gaps rather than celebrating deployment percentages.

Strengthen phishing defences

Given the growing importance of phishing, organisations should:

- Improve email filtering
- Implement SPF, DKIM and DMARC
- Conduct regular phishing simulations
- Deliver targeted awareness training

Users remain the first line of defence.

Protect credentials

Review controls around:

- Password reuse
- Privileged accounts
- Credential theft monitoring
- Password spraying
- Brute-force attacks

Credential security should be treated as a critical business risk.

Expand identity monitoring

Ensure monitoring covers:

- Unusual login behaviour
- Impossible travel alerts
- MFA fatigue attacks
- Suspicious privilege changes
- Abnormal authentication patterns

The goal is to identify compromised accounts before attackers can deploy ransomware.

Strengthen privileged access management

Attackers rarely stop at a standard user account.

Review:

- Administrative privileges
- Standing privileged access
- Tier 0 accounts
- Just-in-time administration

Reducing privilege reduces organisational impact.

Continue patching—but don't rely on it alone

Vulnerability management remains essential, but it must be balanced with:

- Identity security
- Email security
- Endpoint detection
- User awareness
- Threat monitoring

Modern ransomware defence requires all of these capabilities working



Companies keep getting breached by vulnerabilities they already knew about

<https://www.helpnetsecurity.com/2026/07/16/ciso-vulnerability-remediation-gap/>

This article highlights a challenge that many will find familiar: finding vulnerabilities is no longer the difficult part. Most organisations have invested heavily in scanning tools, vulnerability management platforms and security dashboards. The problem is that identifying weaknesses doesn't automatically mean they get fixed.

The vulnerability management gap

According to the research, organisations are discovering more vulnerabilities than ever before, largely thanks to improved automation and scanning capabilities.

However, on average, **58% of remediation activities still require direct human involvement**. While discovery is largely automated, decisions around prioritisation, approvals, deployment and verification continue to rely heavily on people.

This creates a significant gap between knowing a vulnerability exists and actually removing the risk.

Ownership is often unclear

One of the most interesting findings is that **82% of organisations cannot consistently remediate vulnerabilities within the same team that identifies them**.

In most organisations, security teams find the issue, but another team is responsible for fixing it. That means vulnerabilities are handed between infrastructure teams, application owners, cloud teams and service managers.

Every handoff introduces delay.

Even more concerning, over a third of organisations reported that remediation ownership is either unclear or varies depending on the situation. When nobody clearly owns a problem, it often sits unresolved while teams debate responsibility rather than taking action. This is frequently how known vulnerabilities end up becoming breach headlines months later.

Discovery has become a ticketing exercise

The report found that the most common response to identifying a critical vulnerability is raising a ticket in platforms such as Jira or ServiceNow.

While tracking work is important, creating a ticket doesn't reduce risk. After a ticket is raised, organisations still need to:

- Prioritise the issue
- Assign ownership
- Obtain approvals
- Deploy fixes

- Validate that the vulnerability has been successfully remediated

Every additional step introduces friction and delay.

Many organisations have become excellent at finding vulnerabilities but have not built equally mature processes for fixing them.

What the best performers are doing differently

The report identified a small group of organisations that have largely automated the remediation process.

What stood out was how similar their approach was:

- A single platform managed the process from discovery through to verification
- Tool sprawl was reduced
- Remediation was automated wherever possible
- Ownership was clearly defined
- Validation occurred as part of the workflow

Rather than focusing on generating more findings, they focused on reducing the time between discovery and remediation.

What this means

This article reinforces an important truth:

Most organisations don't have a vulnerability discovery problem — they have a remediation problem.

Attackers aren't interested in how many vulnerabilities you've found. They're interested in the ones that remain unpatched.

Many security programmes still measure success through:

- Number of scans completed
- Number of findings identified
- Number of critical vulnerabilities discovered

However, these are activity metrics, not outcome metrics.

The real measure of maturity is how quickly and consistently risk is removed.

The biggest takeaway from this article is that organisations continue to suffer breaches from vulnerabilities they already knew about because remediation processes are often fragmented, manual and poorly owned.

The security industry has become very good at identifying weaknesses. The organisations that will reduce risk most effectively over the next few years will be those that focus less on finding more vulnerabilities and more on **fixing them faster**.

In short: **the future of vulnerability management is not better discovery — it's better remediation.**

Action Points

Clarify ownership

Ask a simple question:

Who owns the remediation of a critical vulnerability?

If the answer isn't immediate and clear, there is likely a governance issue that needs addressing.

Measure remediation performance

Focus on metrics such as:

- Mean Time to Remediate (MTTR)
- Critical vulnerability SLA compliance
- Percentage of overdue vulnerabilities
- Age of unresolved critical findings

Reduce handoffs

Review the workflow between discovery and remediation.

Every approval stage, queue and escalation point adds delay and increases risk exposure.

Automate wherever possible

Look for opportunities to automate:

- Patch deployment
- Configuration updates
- Ticket assignment
- Validation and verification activities

Automation should focus on removing operational bottlenecks.

Align security and IT teams

Vulnerability management should be a shared objective across:

- Security
- Infrastructure
- Cloud operations
- Application owners
- Service delivery teams

Finding the issue and fixing the issue must be part of the same risk management process.

Report risk, not volume

Boards and executive teams generally care about:

- How much risk exists today
- How quickly critical issues are being addressed
- Whether exposure is increasing or decreasing

The number of vulnerabilities found is far less important than the number successfully remediated.



Zoom Patches Critical Windows Flaw That Could Enable Account Takeover

<https://thehackernews.com/2026/07/zoom-patches-critical-windows-flaw-that.html>

This week's Zoom security update is a reminder that even the collaboration tools we use every day can become attractive targets for attackers. The headline issue is a **critical vulnerability affecting Zoom Workplace for Windows** that could potentially allow an attacker to take over a user's account without authentication.

What happened?

Zoom has released patches for **CVE-2026-53412**, a critical vulnerability with a CVSS score of **9.8**.

The flaw affects:

- Zoom Desktop Client for Windows
- Zoom VDI Client for Windows
- Zoom Meeting SDK for Windows

According to Zoom, an improper input validation vulnerability could allow an unauthenticated attacker to conduct an account takeover over the network. In simple terms, an attacker may not need valid credentials to compromise a user's Zoom account if the system is running a vulnerable version.

For organisations that rely heavily on Zoom for meetings, webinars, teaching, or collaboration, that makes this a high-priority patching issue.

It's not just one vulnerability

Alongside the critical flaw, Zoom also fixed three high-severity vulnerabilities affecting Windows deployments. These include privilege escalation weaknesses that could allow a user with local access to gain higher levels of access on a device.

The vulnerabilities include:

- A privilege escalation vulnerability in the Zoom Workplace VDI Plugin.
- A race-condition vulnerability during installation and uninstallation processes.
- A privilege management flaw affecting Zoom Rooms for Windows.

While these require some level of local access, they could be valuable to attackers who have already gained a foothold on a system and are looking to move further into an environment.

Why this matters

Collaboration platforms sit at the centre of modern organisations. They are trusted, widely deployed, and often integrated with:

- Corporate identities
- Calendars
- Email systems

- Contact lists
- Meeting recordings
- AI assistants and productivity tools

A compromised collaboration account can provide significant intelligence about an organisation and potentially open the door to wider attacks. Even if the initial impact appears limited to Zoom, the downstream consequences could be much broader.

The good news is that, at the time of reporting, there is **no evidence of active exploitation in the wild**. However, public disclosure often increases attacker interest, meaning organisations should move quickly to reduce exposure.

What this means

The key lesson here is that collaboration tools should be treated as **business-critical applications**, not simple productivity software.

We often focus security attention on servers, firewalls and cloud platforms while overlooking the tools staff use every day. Attackers do not make the same distinction.

This isn't currently an active incident, but it is exactly the kind of vulnerability that attackers look to weaponise once details become public. The combination of a **critical account takeover vulnerability** and several **privilege escalation flaws** makes this a patching priority for any organisation using Zoom on Windows.

The takeaway is simple: **collaboration platforms are now part of the attack surface and should be managed with the same rigour as any other critical business system.**

Action Points

Patch Zoom immediately

Review deployed Zoom versions and update affected clients as a priority. Focus particularly on Windows endpoints running:

- Zoom Workplace
- Zoom VDI Client
- Zoom Meeting SDK
- Zoom Rooms deployments

Review software update processes

Ensure collaboration tools are included in routine vulnerability management programmes rather than being treated separately from core enterprise software.

Strengthen endpoint visibility

Monitor for:

- Unusual Zoom-related processes
- Unexpected privilege escalation activity
- Suspicious software installation behaviour

Verify identity controls

Because the primary risk involves account compromise:

- Enforce MFA

- Review conditional access policies
- Monitor for unusual authentication activity

Assess Zoom Rooms and VDI environments

These platforms are often overlooked during patching cycles but are specifically mentioned in the advisory.

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

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.



[Ransomware in 2026: More groups, more victims, no slowdown](#)

Despite years of investment in cyber security, ransomware continues to grow. The latest data shows that 2026 is seeing more ransomware groups, more victims and no meaningful reduction in attacks . While some major ransomware gangs have been disrupte...

Steve McIntosh

7/24/2026



[Google's Gemini 3.5 Flash Cyber becomes a vulnerability hunter](#)

Google has unveiled Gemini 3.5 Flash Cyber , a specialist AI model designed to find, validate and patch software vulnerabilities. Unlike general-purpose AI models, this version has been specifically tuned for cyber security tasks and is intended to ...

Steve McIntosh

7/24/2026

3 views



[Operational Cyberpsychology: Applying Behavioral Science to Harden Enterprise Cyber Defense](#)

This article explores a simple but often overlooked reality: cyber security is fundamentally a human problem as much as a technology problem . Most security programmes are built around technology controls, processes and compliance requirements. Howe...

Steve McIntosh

7/24/2026

3 views



[Shadow AI is becoming enterprise security's biggest blind spot](#)

Artificial intelligence is rapidly becoming part of everyday work. Staff are using AI to draft emails, analyse spreadsheets, summarise documents, write code and automate

routine tasks. The problem is that many of these tools are being adopted faster...

Steve McIntosh

7/24/2026

5 views



[Cyber Threat Intelligence Digest: Week 29](#)

This week's Acumen Cyber Threat Intelligence Digest highlights a familiar but important trend: attackers are becoming increasingly effective at exploiting trusted technologies, user behaviour and known weaknesses rather than relying on sophisticated...

Steve McIntosh

7/23/2026

9 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

7/23/2026

5 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.

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.



Cyber Threat Intelligence Digest – Week 29

<https://acumencyber.com/cyber-threat-intelligence-digest-july-2026-week-29>

This week's Acumen Cyber Threat Intelligence Digest highlights a familiar but important trend: attackers are becoming increasingly effective at exploiting trusted technologies, user behaviour and known weaknesses rather than relying on sophisticated new techniques.

A key theme throughout the report is the continued rise of **social engineering and identity-based attacks**. Phishing campaigns remain highly effective, with threat actors increasingly impersonating trusted brands, software providers and cloud services to steal credentials and gain access to corporate environments. Once inside, attackers are focusing on data theft, privilege escalation and long-term persistence rather than immediate disruption.

The digest also draws attention to several newly disclosed vulnerabilities affecting widely used technologies. As always, the risk is less about the vulnerability itself and more about how quickly organisations respond. Threat actors continue to weaponise public disclosures rapidly, reducing the window available for patching and remediation.

Another notable trend is the growing abuse of legitimate platforms and services. Attackers are increasingly hiding malicious activity within trusted infrastructure, making traditional detection methods less effective. This reinforces the need for organisations to focus on behavioural monitoring rather than relying solely on signatures or reputation-based controls.

From an education sector perspective, the report serves as a reminder that universities and colleges remain attractive targets due to the volume of personal data, research information and financial records they hold. Credential theft, ransomware and supply chain compromise continue to present significant risks.

Bottom Line

The biggest takeaway from this week's intelligence is that attackers continue to succeed by exploiting **people, identities and operational gaps** rather than groundbreaking technical exploits. Organisations that maintain strong patching discipline, robust identity controls and effective monitoring will be best placed to reduce risk and respond quickly when threats emerge.

Action Points

- Prioritise patching of critical and internet-facing systems.
- Review identity security controls, particularly MFA coverage and privileged access.
- Increase monitoring for suspicious account activity and anomalous user behaviour.
- Strengthen phishing awareness training, particularly around brand impersonation and fake software updates.
- Review detection capabilities for abuse of legitimate tools and cloud services.
- Validate backup, recovery and incident response processes through regular testing.
- Monitor third-party and supply chain dependencies for emerging risks.



Shadow AI is becoming enterprise security's biggest blind spot

<https://www.helpnetsecurity.com/2026/07/23/shadow-ai-security-risks/>

Artificial intelligence is rapidly becoming part of everyday work. Staff are using AI to draft emails, analyse spreadsheets, summarise documents, write code and automate routine tasks. The problem is that many of these tools are being adopted faster than organisations can govern them. As a result, security teams are losing visibility over how AI is being used and what data is being shared.

The article argues that **Shadow AI**—the use of unapproved or unmanaged AI tools—is quickly becoming one of the biggest security blind spots facing organisations today.

Why it matters

Unlike traditional software deployments, AI adoption often happens without formal approval. Employees can sign up for a free AI service, enable AI features built into existing applications, or use personal accounts and devices within minutes. What starts as a quick productivity boost can gradually become embedded in day-to-day business processes.

The real risk isn't the AI tool itself—it's the data being shared.

Employees may upload:

- Internal business documents
- Customer information
- Financial data
- Research materials
- Confidential emails

Often without understanding where that information is stored, processed or reused.

Technical controls alone are unlikely to solve the problem. Blocking one AI tool often simply pushes users towards another service, a personal device or an AI feature already embedded within approved software.

What this means

Shadow AI is starting to look very similar to the Shadow IT problem security teams faced years ago—except with potentially greater consequences.

Many organisations believe they have approved only a handful of AI tools. The reality is that there may be dozens already in use across departments without security, legal or compliance oversight. Visibility is becoming the biggest challenge.

The issue is not that employees are deliberately bypassing controls. In most cases, they are simply trying to work more efficiently and deliver results faster. Unfortunately, those well-intentioned actions can create significant risks relating to data protection, intellectual property, regulatory compliance and information security.

Bottom Line

Shadow AI is being driven primarily by employees trying to be more productive rather than malicious insiders. However, every unapproved AI tool introduces potential data protection, compliance and security risks.

The organisations that will benefit most from AI won't be those that try to ban it. They will be the ones that establish visibility, governance and sensible guardrails, allowing staff to innovate safely while protecting organisational data and maintaining regulatory compliance.

Action Points

- Develop and publish a clear AI governance framework.
- Identify and inventory AI usage across the organisation.
- Educate staff on what information can and cannot be entered into AI tools.
- Review AI capabilities already embedded within existing business applications.
- Monitor for unauthorised AI usage where practical.
- Update data protection, acceptable use and information handling policies to cover AI.

- Provide approved AI services so employees have safe alternatives.
- Work with legal, compliance and data protection teams to assess AI-related risks.



Operational Cyberpsychology: Applying Behavioral Science to Harden Enterprise Cyber Defense

<https://www.cyberdefensemagazine.com/operational-cyberpsychology-applying-behavioral-science-to-harden-enterprise-cyber-defense/>

This article explores a simple but often overlooked reality: cyber security is fundamentally a **human problem as much as a technology problem**.

Most security programmes are built around technology controls, processes and compliance requirements. However, attackers frequently succeed because they understand human behaviour better than we do. They exploit trust, urgency, curiosity, fear and routine to influence decisions and gain access to systems.

The concept of **operational cyberpsychology** focuses on applying behavioural science to improve cyber resilience. Rather than simply asking people to follow policies, it seeks to understand why individuals behave the way they do and how organisations can design environments that encourage secure decisions.

Traditional awareness programmes often assume that if people know the rules, they will follow them.

In reality, employees face competing pressures every day:

- Deadlines
- Workload
- Convenience
- Productivity targets
- Collaboration demands

When security controls are perceived as obstacles, people naturally find workarounds.

Attackers understand this. Phishing emails, social engineering calls and business email compromise attacks all rely on influencing human behaviour rather than bypassing technical controls.

The article suggests that organisations should spend less time blaming users and more time understanding the conditions that lead to risky behaviour.

The biggest takeaway is that security culture is not created through annual training courses.

A truly resilient organisation understands how people make decisions under pressure and designs processes around those realities.

Instead of asking:

"Why did that employee click the link?"

A more useful question is:

"What conditions made that action seem reasonable at the time?"

This shift moves security from a culture of blame towards a culture of learning and continuous improvement.

For many organisations, the weakest link is not a lack of awareness. It is poorly designed processes, excessive complexity and security controls that don't align with how people actually work.

Operational cyberpsychology challenges the idea that technology alone will solve cyber risk. The most effective security programmes recognise that people are not the problem to be managed—they are part of the solution.

Organisations that understand behaviour, reduce friction and build a positive security culture will generally be more resilient than those that focus purely on technical controls. In short, if attackers are using psychology to compromise organisations, defenders should be using psychology to protect them.

Action Points

- Review security controls through a behavioural lens.
- Simplify processes that encourage users to find workarounds.
- Incorporate behavioural science into awareness and training programmes.
- Measure security culture, not just policy compliance.
- Use phishing simulations to understand behaviour patterns rather than assign blame.
- Encourage reporting of mistakes without fear of punishment.
- Design security controls that support productivity rather than obstruct it.
- Work with HR, communications and organisational development teams to strengthen security culture.



Google's Gemini 3.5 Flash Cyber becomes a vulnerability hunter

<https://www.helpnetsecurity.com/2026/07/22/google-gemini-3-5-flash-cyber-model/>

Google has unveiled **Gemini 3.5 Flash Cyber**, a specialist AI model designed to find, validate and patch software vulnerabilities. Unlike general-purpose AI models, this version has been specifically tuned for cyber security tasks and is intended to help defenders identify weaknesses before attackers can exploit them. It will initially be available through Google's CodeMender platform to governments and trusted partners, with broader access expected over time.

Why it matters

One of the biggest challenges in cyber security is the sheer scale of modern software. Large applications contain millions of lines of code and countless possible execution paths. Traditional scanning tools often struggle to keep pace.

Gemini 3.5 Flash Cyber is designed to tackle this problem by analysing multiple code paths simultaneously. Rather than relying on one large AI query, CodeMender can invoke the model repeatedly, allowing it to examine far more potential vulnerabilities at speed and at lower cost.

During testing, the model demonstrated impressive results. In analysis of Google's V8 JavaScript engine, it identified **55 confirmed vulnerabilities**, including several issues that other leading AI models missed entirely. Google also reported that the model discovered remote-code-execution vulnerabilities and memory corruption flaws in production environments within hours.

The significance of this announcement goes well beyond Google's latest AI product.

We're entering an era where AI is becoming a force multiplier for both attackers and defenders. Historically, vulnerability discovery required significant time, expertise and manual effort. AI is now dramatically reducing that effort.

For defenders, this creates an opportunity to identify and remediate weaknesses earlier. However, it also raises concerns that attackers may use similar technologies to discover vulnerabilities faster than organisations can patch them.

The race between discovery and remediation is accelerating.

Gemini 3.5 Flash Cyber highlights how quickly AI is transforming cyber security. The ability to automatically discover and help remediate vulnerabilities at scale could significantly improve defensive capabilities.

However, the technology also reinforces a growing reality: organisations are likely to face a future where vulnerabilities are found faster than ever before. Success will depend not just on finding weaknesses, but on fixing them quickly enough to stay ahead of attackers.

Action Points

- Prepare for AI-assisted vulnerability management becoming mainstream.
 - Review how security teams prioritise remediation once vulnerabilities are identified.
 - Continue investing in secure development practices and code scanning.
 - Strengthen patch management processes to reduce remediation timelines.
 - Assess how AI could improve internal security testing and secure software development.
 - Review governance arrangements for the use of AI in security operations.
-



Ransomware in 2026: More groups, more victims, no slowdown

<https://www.helpnetsecurity.com/2026/07/24/ransomware-attack-trends-2026-report/>

Despite years of investment in cyber security, ransomware continues to grow. The latest data shows that 2026 is seeing **more ransomware groups, more victims and no meaningful reduction in attacks**. While some major ransomware gangs have been disrupted by law enforcement, new groups have quickly emerged to fill the gap.

What's happening?

One of the key findings is that the ransomware ecosystem has become increasingly fragmented. Rather than a handful of dominant groups controlling the market, there are now dozens of active operators competing for victims.

Ransomware-as-a-Service (RaaS) continues to fuel this growth. Criminal groups develop the tools and infrastructure, while affiliates carry out attacks. This lowers the barrier to entry and allows less-skilled attackers to launch sophisticated campaigns.

At the same time, attackers are becoming more focused on stealing data before encrypting systems. Double-extortion remains the dominant model, where organisations face both operational disruption and the threat of public data exposure.

The report highlights a difficult reality: law enforcement takedowns are important, but they are not slowing the overall ransomware market.

When one major group disappears, affiliates often migrate to another platform within days or weeks. The criminal ecosystem has proved highly resilient and adaptable.

Attackers are also becoming more opportunistic. They are exploiting:

- Unpatched vulnerabilities
- Stolen credentials
- Phishing campaigns
- Exposed remote access systems
- Weak identity controls

Most successful attacks still rely on known techniques rather than advanced nation-state level capabilities.

The biggest takeaway is that ransomware should no longer be viewed as a purely technical threat.

Ransomware is now a business resilience issue.

The question is no longer whether your organisation can prevent every attack. The more important question is:

Can you continue operating if prevention fails?

Security leaders should assume that attackers may eventually gain access and focus equal attention on detection, containment, recovery and communications.

For universities and colleges, the risk is particularly significant due to large user populations, complex environments, valuable research data and extensive third-party relationships.

The ransomware landscape continues to expand despite increased defensive investment and high-profile law enforcement actions. More groups are entering the market, attacks remain highly profitable, and victims continue to pay.

Success in 2026 will depend less on trying to stop every attack and more on building organisational resilience. The organisations that recover fastest, contain incidents effectively and minimise operational disruption will be the ones best positioned to withstand the continuing ransomware threat.

Action Points

- Strengthen identity security and MFA coverage.
- Prioritise patching of internet-facing systems.
- Review remote access and privileged account controls.
- Monitor for signs of lateral movement and credential theft.
- Implement robust backup and recovery processes.
- Test ransomware response plans through regular exercises.
- Review third-party and supply chain risks.
- Enhance monitoring for data exfiltration activity.
- Ensure executive leadership understands ransomware response responsibilities.

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

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.



[Ghost Credentials Expose Cloud Systems to Hidden Identity Risks](#)

A growing cyber security concern is emerging from something many organisations barely track: non-human identities (NHIs) . These include service accounts, API keys, automation tools, AI agents, tokens and machine identities that operate behind the s...

Steve McIntosh

7/30/2026

2 views



[Three Critical VMware Flaws Allow Auth Bypass, Code Execution, and VM Escape](#)

Broadcom has released security updates for multiple VMware products, including vCenter, ESX/ESXi, Workstation and Fusion, addressing several serious vulnerabilities. Three of these flaws have been rated as critical, with potential impacts ranging f...

Steve McIntosh

7/30/2026

3 views



[Vulnerability management needs an update for the AI era](#)

The traditional approach to vulnerability management is coming under increasing pressure as artificial intelligence accelerates both the discovery of vulnerabilities and the speed at which attackers can identify potential weaknesses. Security teams ...

Steve McIntosh

7/30/2026

12 views



[NCSC Warns of Russian Zero-click Exploit Campaign](#)

The NCSC has warned that a Russian state-backed threat group known as LAUNDRY BEAR has been carrying out a sophisticated cyber espionage campaign against Western

organisations. Active since at least July 2025, the group has targeted organisations ac...

Steve McIntosh

7/30/2026

11 views



[Cyber Threat Intelligence Digest: Week 30](#)

This week's Acumen Digest demonstrates that while technical vulnerabilities remain a concern, attackers are increasingly succeeding by exploiting trust and user behaviour rather than relying solely on sophisticated exploits. Several significant vuln...

Steve McIntosh

7/30/2026

9 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

7/30/2026

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.

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.



Cyber Threat Intelligence Digest: Week 30

<https://acumencyber.com/cyber-threat-intelligence-digest-july-2026-week-30>

This week's Acumen Digest demonstrates that while technical vulnerabilities remain a concern, attackers are increasingly succeeding by exploiting trust and user behaviour rather than relying solely on sophisticated exploits.

Several significant vulnerabilities were disclosed during the week. Adobe patched a flaw in its Acrobat Chrome extension that could allow attackers to access WhatsApp Web content simply by persuading a user to visit a malicious website.

A serious Linux XFS filesystem vulnerability was also disclosed, potentially allowing local users to gain root privileges on affected systems.

Additionally, a critical remote code execution vulnerability in vBulletin has become more concerning following the release of public proof-of-concept exploit code.

Social engineering continues to be one of the most effective attack methods. Cyber criminals were observed using Steam discussion forums to trick users into running malicious PowerShell commands, resulting in the installation of cryptocurrency mining malware.

Elsewhere, the FakeAgent campaign leveraged malicious advertising and fake Claude AI download pages to distribute SectopRAT, a remote access trojan capable of stealing credentials, financial information and other sensitive data.

Another notable incident involved malware hidden within Steam Workshop content. This ultimately led to the compromise of a Discord administrator account and the takeover of a community server, highlighting the risks associated with downloading untrusted content, even from well-known platforms.

There were also broader industry discussions around the rise of AI-enabled scams, increasing concern over caller ID spoofing, and the unusual report of AI systems being linked to a breach of AI platform Hugging Face.

The common theme this week is trust exploitation. Whether through fake AI downloads, malicious forum advice, or compromised community platforms, attackers are relying on users taking actions that appear legitimate. Strong patch management, user awareness and a healthy culture of verification remain some of the most effective defences available.

Action Points

- Apply the latest patches for Adobe Acrobat Chrome Extension, Linux XFS systems and vBulletin platforms as a priority.
 - Review Linux servers running XFS filesystems, particularly Red Hat, Rocky Linux, AlmaLinux and Amazon Linux environments.
 - Reinforce user awareness training around fake software downloads, sponsored search results and social engineering attacks.
 - Monitor for signs of unauthorised remote access tools, cryptocurrency miners and suspicious scheduled tasks.
 - Review controls around gaming, collaboration and community platforms where malicious content is increasingly being used as an attack vector.
-



NCSC Warns of Russian Zero-click Exploit Campaign

<https://www.digit.fyi/ncsc-warns-of-russian-zero-click-exploit-campaign/>

The NCSC has warned that a Russian state-backed threat group known as **LAUNDRY BEAR** has been carrying out a sophisticated cyber espionage campaign against Western organisations. Active since at least July 2025, the group has targeted organisations across government, defence, education, energy, law enforcement, media and technology sectors with the primary objective of stealing sensitive email data.

What makes this campaign particularly concerning is the use of a **zero-click exploit** known as **Beehive (Ulej)**. Unlike traditional phishing attacks, victims do not need to click a link, open an attachment or enter credentials. Simply viewing a specially crafted email within a vulnerable version of Zimbra Collaboration Suite (ZCS) can be enough to compromise an account.

Once access is gained, attackers can steal large volumes of email data, organisational address books, browser-saved credentials and authentication-related information. The campaign is also designed to establish persistent access, allowing attackers to maintain visibility of communications over an extended period.

The NCSC has assessed the activity as being consistent with Russian state-sponsored espionage. Intelligence agencies noted that the techniques were first tested against Ukrainian targets before being deployed against organisations across NATO-aligned countries and other Western nations.

What This Means for Higher and Further Education

For universities and colleges, email remains one of the most critical business systems and often contains sensitive research, student data, commercial partnerships and executive communications. The fact that compromise can occur without any user interaction means that traditional phishing awareness training alone is no longer sufficient protection.

This campaign reinforces the importance of patching, monitoring and detection capabilities around email platforms. If attackers gain access to mailboxes, they can often gather intelligence, harvest credentials and position themselves for further compromise elsewhere in the environment.

The most important lesson from this campaign is that attackers are actively looking for ways to bypass user awareness controls. When simply viewing an email can lead to compromise, organisations must rely on strong patch management, effective monitoring and rapid detection capabilities. Email remains a prime target for nation-state actors, and any unpatched collaboration platform should be treated as a significant business risk.

Action Points

- Confirm all Zimbra Collaboration Suite instances are fully patched.

- Review email and authentication logs for unusual activity.
- Monitor for unexpected mailbox exports or access patterns.
- Review how MFA recovery codes are stored and protected.
- Ensure incident response plans cover email platform compromise scenarios.
- Strengthen monitoring around privileged accounts and email administration functions.
- Consider additional monitoring for organisations that handle sensitive research, government-funded projects or regulated data.
- Register for and utilise the NCSC Early Warning service where appropriate.

Vulnerability management needs an update for the AI era

<https://www.techtarget.com/searchsecurity/feature/Vulnerability-management-needs-an-update-for-the-AI-era>

The traditional approach to vulnerability management is coming under increasing pressure as artificial intelligence accelerates both the discovery of vulnerabilities and the speed at which attackers can identify potential weaknesses.



Security teams are facing a growing challenge. AI-powered tools are enabling researchers and vendors to find more vulnerabilities across larger codebases and at a much faster rate than ever before. While this is helping vendors identify and fix flaws earlier, it is also contributing to a significant increase in the number of reported vulnerabilities that organisations must assess and manage.

The key message is that organisations can no longer rely solely on traditional patching cycles or attempt to patch every vulnerability equally. With vulnerability volumes continuing to rise, security teams need to become more risk-focused and prioritise vulnerabilities based on exploitability, business impact and exposure rather than CVSS scores alone.

The article highlights that AI is changing the economics of vulnerability discovery. Defenders can use AI to identify weaknesses before attackers find them, but threat actors also have access to increasingly capable tools that can accelerate reconnaissance, exploit development and vulnerability research. This creates a situation where the window between vulnerability discovery and exploitation may continue to shrink.

For many organisations, particularly those with limited security resources, the answer is not necessarily more patching. Instead, it is better prioritisation, stronger asset visibility and a more mature vulnerability exposure management programme. Understanding which systems are critical to the business, which assets are internet-facing and which vulnerabilities are actively exploitable is becoming far more important than maintaining lengthy patch lists.

Universities and colleges often operate large, decentralised environments with thousands of assets, research systems and legacy platforms. As vulnerability volumes increase, attempting to patch everything immediately is neither realistic nor sustainable. A risk-based approach is essential to focus efforts where they reduce the greatest amount of risk.

The challenge is no longer finding vulnerabilities; it is deciding which ones matter most. In the AI era, successful organisations will not be those that patch the most vulnerabilities, but those that can identify and remediate the vulnerabilities that present the greatest risk to the business.

Action Points

- Move from vulnerability management to vulnerability exposure management.
- Prioritise vulnerabilities based on exploitability and business impact.
- Maintain an accurate and up-to-date asset inventory.
- Focus remediation efforts on internet-facing and critical systems first.
- Use threat intelligence to identify vulnerabilities that are actively being exploited.
- Review patching SLAs to ensure they are aligned to risk rather than severity scores alone.
- Consider how AI-powered security tools can enhance vulnerability discovery and prioritisation.
- Ensure security teams are prepared for increasing vulnerability volumes over the next few years.



Three Critical VMware Flaws Allow Auth Bypass, Code Execution, and VM Escape

<https://thehackernews.com/2026/07/three-critical-vmware-flaws-allow-auth.html>

Broadcom has released security updates for multiple VMware products, including vCenter, ESX/ESXi, Workstation and Fusion, addressing several serious vulnerabilities. Three of these flaws have been rated as **critical**, with potential

impacts ranging from authentication bypass and remote code execution to virtual machine escape.

The most significant vulnerabilities affect **VMware vCenter**, the central management platform used to administer virtual infrastructure. The first flaw, **CVE-2026-59309**, allows an attacker with network access to bypass authentication and gain unauthorised access to vCenter. The second, **CVE-2026-59310**, is a directory traversal vulnerability that could enable remote code execution on affected systems.

A third critical vulnerability, **CVE-2026-47876**, affects the VMXNET3 virtual network adapter in VMware ESX. An attacker who already has administrative privileges inside a virtual machine could potentially break out of the guest environment and execute code on the underlying hypervisor host. This type of attack, known as a **VM escape**, is particularly concerning because it undermines one of the fundamental security assumptions of virtualisation.

Although Broadcom has stated there is currently no evidence of active exploitation, VMware vulnerabilities have historically attracted significant attention from both cyber criminals and nation-state actors because of the access they can provide to enterprise infrastructure. Organisations should expect these flaws to become a priority target for attackers in the coming weeks.

Many universities and colleges rely heavily on VMware to support critical business services, student systems, research platforms and teaching environments. A successful compromise of vCenter could provide attackers with visibility and control across large parts of the virtual estate from a single point of access.

The VM escape vulnerability is particularly relevant where multiple workloads share the same host infrastructure, including research, development and teaching environments.

These vulnerabilities highlight the importance of protecting the virtualisation layer. Compromise of vCenter or the hypervisor is not simply a server issue; it can provide attackers with a pathway into large portions of the organisation's infrastructure. Even without evidence of active exploitation, the severity and potential impact justify urgent patching and review of VMware management security controls.

Action Points

- Prioritise patching all affected VMware vCenter systems.
- Update ESX/ESXi hosts to address the VM escape vulnerability.
- Review whether vCenter management interfaces are accessible from untrusted or user networks.
- Ensure management networks are properly separated from production and user environments.
- Review and restrict privileged access to VMware administration platforms.

- Validate logging, monitoring and alerting for vCenter and ESXi activity.
- Review incident response plans for virtual infrastructure compromise scenarios.
- Monitor threat intelligence feeds for reports of active exploitation.



Ghost Credentials Expose Cloud Systems to Hidden Identity Risks

<https://www.darkreading.com/cloud-security/non-human-identity-sprawl-creates-a-new-cloud-attack-path>

A growing cyber security concern is emerging from something many organisations barely track: **non-human identities (NHIs)**. These include service accounts, API keys, automation tools, AI agents, tokens and machine identities that operate behind the scenes across cloud environments.

The article highlights the discovery of so-called "**ghost credentials**" following an investigation into unusual activity from a dormant AI-powered workflow agent. After sitting inactive for more than a month, the agent unexpectedly began making API calls, leading researchers to uncover a web of forgotten permissions, service accounts and trust relationships that could potentially be abused by attackers.

The concern is that many organisations have focused heavily on protecting human identities through MFA, privileged access controls and conditional access policies, while machine identities have grown rapidly with far less governance. These non-human identities often retain permissions long after they are needed and can provide attackers with a path to move laterally, escalate privileges and gain access to critical systems.

Researchers found that some organisations have hundreds of non-human identities per developer, making it extremely difficult to understand who or what has access to sensitive resources. As cloud adoption, automation and AI accelerate, these hidden trust relationships are creating a new attack surface that many security teams struggle to monitor effectively.

The article argues that attackers are increasingly likely to target exposed tokens, API keys and machine credentials because they are often easier to obtain than compromising a user account through phishing. In many environments, a single leaked credential can provide a route to much broader access.

Universities and colleges are increasingly dependent on cloud services, automation platforms, research environments and AI-enabled tools. These environments often generate large numbers of service accounts, application identities and API integrations.

Over time, identities created for projects, research initiatives, integrations and temporary services can remain active long after their original purpose has ended. Without proper governance, these forgotten accounts can become hidden pathways for attackers.

Identity is no longer just about people. In modern cloud environments, machines often outnumber human users by a significant margin, and attackers know it. Organisations that focus only on protecting staff accounts while ignoring service accounts, tokens and AI agents risk leaving a large portion of their attack surface unmanaged. Effective identity security must now cover both humans and machines.

Action Points

- Create and maintain a complete inventory of non-human identities across cloud platforms.
 - Review service accounts, API keys and automation accounts for excessive permissions.
 - Identify and remove dormant, unused or orphaned credentials.
 - Monitor for unusual activity from service accounts and automated processes.
 - Include non-human identities within IAM, PAM and security monitoring programmes.
-

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