Cyber Brief (October 2026)

November 6, 2026 - Version 1

TLP:CLEAR

Executive summary

- We analysed 415 open source reports for this Cyber Brief[1].
- Relating to **cyber policy and law enforcement**, Europol reported the seizure of KillSec ransomware infrastructure and three provisional arrests. In the United States (US) authorities seized tools linked to China-linked Flax Typhoon, while Microsoft and partners disrupted a Chinese-language cybercrime ecosystem known as YAM store.
- Regarding **cyberespionage & prepositioning**, Citizen Lab described a zero-click mercenary spyware campaign using targeted ads dubbed AdFection. Russia-aligned threat actors targeted organisations supporting Ukraine and entities in two NATO states.
- On the **cybercrime** front, Palo Alto Networks reported attempted exploitation of PAN-OS by ransomware operators.
- Regarding **data exposure and leaks**, Danish organisations disclosed separate incidents involving the population register and two universities. Moldova reported the extraction of personal data through government services.
- As for **disruption and opportunistic exploitation**, a number of critical vulnerabilities affecting products like PHP (PHP2Shell), Gitlab, Cisco Secure E-mail Gateway, NodeJS (Node2Shell), Fortimail and Atlassian products were reported to be exploited in the wild.
- In **artificial intelligence**, Black Lotus Labs reported malware spreading through exposed AI services. Microsoft assessed that China-linked Storm-0940 used ARTEX to accelerate intrusions involving ConnectWise cloud applications. South Korean police investigated data theft and extortion that separate reporting linked to ARTEX.

For more information regarding CERT-EU's analytical and operational standards to classify, assess, and prioritise malicious cyber activities, please see our Cyber Threat Intelligence Framework here.

Europe

Cyber policy and law enforcement

KillSec ransomware operation dismantled in Operation KillSwitch On October 1, Europol reported that its Operation KillSwitch disrupted the KillSec cybercrime group by seizing its leak site and servers. Police provisionally arrested three suspects. Investigators identified a 16-year-old as the suspected main operator. Authorities linked KillSec to approximately 1.000 suspected

attacks worldwide and identified around 500 as successful. The group is accused of data theft and extortion. link

German authorities announced the arrest of Qilin suspect extradited from Japan On October 7, authorities in Germany's North Rhine-Westphalia announced a ministerial statement concerning the arrest of a Russian national suspected of membership in the Qilin cybercrime group. Reporting said Japanese authorities had detained the suspect in Osaka before his transfer to Germany. They are suspected of involvement in a ransomware attack against a German company. link link

Cyberespionage & prepositioning

Russia-linked Star Blizzard targeted organisations supporting Ukraine On September 29, Microsoft reported that Russia-linked Star Blizzard had expanded its phishing campaigns, used accounts on compromised websites and adopted a malware-delivery technique it calls RedFlick to deploy the CosmicPulse backdoor. The activity targeted Ukrainian individuals and organisations associated with international policy. Microsoft observed the activity affecting more than 100 organisations, primarily in the US and UK; though this is not a count of confirmed compromises. `public administration` `finance` `europe` `russia` link

Russia-linked Swallowtail targeted organisations in two NATO states On October 20, Symantec reported that Russia-linked Swallowtail (APT28) targeted organisations in two NATO member states: a naval shipbuilder, a military aircraft engine manufacturer and a law enforcement agency. Symantec observed Hookedge activity at these targets; Hookedge was delivered through macro-enabled Word documents. At the shipbuilder, an intrusion lasting at least two months involved credential theft, a persistent SSH tunnel and hands-on reconnaissance. `public administration` `defence` `europe` `russia` link

Polish prosecutors report potential NSO Group access to Pegasus surveillance data On October 12, Poland's National Prosecutor's Office published findings from Internal Security Agency experts concerning Pegasus's operation under the former government. According to prosecutors, the Central Anti-Corruption Bureau did not fully control the system; its producer retained control of some components and could potentially access infrastructure on which surveillance material was processed or stored. Data transmitted by the system left Poland. `public administration` link link

Artificial intelligence

PoeLLM cryptomining and botnet expansion via exposed AI services On October 7, Black Lotus Labs reported that PoeLLM, a financially motivated malware campaign, abused exposed artificial intelligence (AI) services and other internet-facing services to compromise servers. The activity uses command-and-control infrastructure and deploys cryptocurrency miners, while reusing infected hosts to scan for and compromise additional victims. Black Lotus Labs reported more than 3.400 affected servers, mainly in the US and Western Europe. `artificial intelligence` `technology` link

Data exposure and leaks

Denmark reported unauthorised access to Central Person Register data On October 5, Denmark's Central Person Register reported that an unauthorised party abused a company's legitimate search access to obtain names, addresses, reference numbers and other information concerning

approximately 8,8 million people registered in the system, including emigrated and deceased people rather than only current residents. The authority said its review found that names and addresses of people with protected registrations were not included. `public administration` link link

Technical University of Denmark reported unauthorised access to user data On October 2, the Technical University of Denmark (DTU) reported that a malicious actor had compromised DTU user profiles and used them to access DTUBasen, its identity and access management system, extracting a substantial amount of data. Up to 200.000 current and former users could be affected, but DTU could not establish precisely which records were taken or how many people were affected. Potentially exposed information includes civil registration numbers and names; for active users it may also include home addresses and, where registered, next-of-kin details. `education` link

University of Copenhagen reported a test environment data breach On October 23, the University of Copenhagen disclosed that threat actors had accessed a test environment containing personal data relating to approximately 525.000 people; it detected the attack on October 20. Potentially affected data included names, contact details and Danish civil registration (CPR) numbers. The university subsequently reported that access-card PINs were also compromised. Its later notice describes the CPR numbers as encrypted but warns that information needed to decrypt them was also compromised. `universities` `education` link link

Moldova reported extraction of personal data through government services On October 27, Moldova's Electronic Governance Agency reported preliminary estimates that an unidentified actor had made roughly one million queries using personal identification numbers (IDNPs) against government services between October 19 and October 23 and extracted datasets relating to approximately 750.000 IDNPs. An IDNP count is not a confirmed count of distinct affected people. Authorities suspended and reviewed digital services, blocked further attempts and warned of elevated fraud and social-engineering risks. `public administration` link

AfD e-mail messages published after reported provider breach On October 20, Bild reported that an unidentified actor breached an AfD hosting service provider and copied contents from some party e-mail mailboxes, later publishing AI-generated summaries and large volumes of messages online. The published archive was reported to contain approximately 421.000 e-mails (about 281 GB). Reported contents included ID photos, invoices, bank statements and membership lists from North Rhine-Westphalia district associations. The AfD said current federal mailboxes were not affected. `political parties` `civil society` link link

Unauthorised access to ENGIE customer data via partner sales tool In October 2026, Engie, a French multinational electric utility company, notified customers of unauthorised access to personal data following an incident affecting a prospecting and sales tool used by some commercial partners. Exposed data reportedly included identity and contact details and contract identifiers; Engie said banking data and its internal information systems were not affected. `energy` link

World

Cyber policy and law enforcement

US seizes vulnerability-scanning and spearphishing tools linked to Flax Typhoon On October 8, the US Department of Justice and the Federal Bureau of Investigation (FBI) announced court-authorised seizures of Microscan and FishHub, tools operated by China-based Integrity

Technology Group and used by actors linked to Flax Typhoon. The tools enabled vulnerability scanning and spearphishing, and were allegedly used to target US and foreign critical infrastructure and other networks. link

Microsoft and partners disrupt phishing marketplace and kits On October 14, Microsoft and partners announced the seizure and redirection of domains supporting YAM Store, a Chinese-language marketplace for cloud accounts and hosting, and the CoGUI and Iron Man phishing kits. Microsoft said the ecosystem enabled over one billion malicious messages targeting more than 75.000 organisations since January 2026. Proofpoint observed CoGUI-linked credential phishing focused on Japan. Microsoft also disabled over 20.000 YAM Store-linked accounts. link link link

FBI arrests suspected co-conspirator of ShinyHunters On October 9, the FBI director announced the arrest of a suspected ShinyHunters associate. ShinyHunters had claimed to breach an FBI recruitment website and exfiltrate between two TB and three TB of data. The nature and volume of any stolen data have not been independently confirmed. link

Cyberespionage & prepositioning

AdFection spyware campaign used targeted online advertising On October 15, Citizen Lab reported that the Israeli company Discover-Tech used targeted online advertisements to deliver Candiru spyware to devices at Al Jazeera without requiring a click. Its Controlium platform enabled targeting and device fingerprinting. Analysis of four billion bidding records showed that its advertisements appeared on nearly 200.000 distinct websites and mobile apps across dozens of countries over four months and directed millions of devices to contact Candiru servers, likely for tracking and reconnaissance. Citizen Lab confirmed infections at Al Jazeera. `israel` link

Cybercrime

Cybercrime groups attempted to exploit PAN-OS CVE-2026-0257 On October 20, Palo Alto Networks Unit 42 reported exploitation attempts against PAN-OS CVE-2026-0257 by ransomware operators including Qilin and Inc. Separately, it had observed an unidentified actor exploiting the authentication-bypass flaw in June to seek GlobalProtect VPN access; a small proportion of probed devices established VPN sessions. Unit 42 identified no post-access activity or lateral movement in that earlier investigation. The October attempts should not be treated as evidence of ransomware deployment through this flaw. `transport` `education` `health` `finance` `public administration` `technology` `energy` link

Artificial intelligence

China-linked TA419 targeted AI policy experts with credential phishing On October 1, Proofpoint reported that China-linked TA419 conducted credential phishing against AI experts at US think tanks, universities and law firms. The actor impersonated prominent economists and AI policymakers, using staged outreach e-mails to build rapport before directing targets to credential-harvesting pages. Proofpoint assessed that the activity likely sought cloud account access and insight into US AI policy and regulation. `artificial intelligence` `education` `china` link

China-linked Storm-0940 used ARTEX in ConnectWise supply-chain compromise On October 21, Microsoft Threat Intelligence assessed that Storm-0940, a China-linked actor, used the ARTEX agentic AI framework to scale automated attacks against hundreds of organisations globally. Activity included rapid abuse of compromised ConnectWise multi-tenant applications to access

Microsoft 365 customer environments without interactive sign-in, followed by cloud reconnaissance and changes intended to expand or persist access. Microsoft observed several cases in which the actor gained privileged cloud access in under three hours. Microsoft said it found no compromise or vulnerability in its own products or services. `artificial intelligence` `public administration` `energy` `transport` `defence` `technology` `finance` `china` link

Threat actor targeted government organisations with AI-assisted web shells On October 15, Palo Alto Networks Unit 42 reported a campaign tracked as CL-UNK-1211 targeting government entities and state-affiliated infrastructure globally. The actor used AI-assisted payloads to gain and maintain access, harvest credentials and attempt to dump sensitive databases and exfiltrate government mailboxes to access social media accounts. Activity spanned October 2025 to April 2026 and included efforts to evade detection and hinder investigation. `artificial intelligence` `public administration` link

South Korean police investigated ARTEX-enabled data theft On October 26, South Korean police reported an investigation into seven suspects allegedly involved in attacks on more than 90 organisations. The alleged leader was arrested on October 15 and referred to prosecutors on October 23; six others were being investigated without detention. The suspects allegedly stole about 8,2 million records containing personal and internal data from organisations in education, healthcare and other sectors. Extortion e-mails sent to 38 affected organisations demanded an aggregate 360 million won to prevent disclosure. Separate reporting identified ARTEX as a tool allegedly used in the activity. `artificial intelligence` `education` `health` link link

Hugging Face model hub cryptomining and backdoor campaign via remote code loading On October 29, Pluto Research reported a malware campaign distributing trojanised AI model repositories on Hugging Face. Loading an affected repository with remote model code enabled could give the actor backdoor access and enable misuse of graphics processing unit resources for cryptomining. The campaign affected self-hosted inference and shared GPU environments and spanned multiple accounts and tens of thousands of downloads, which are not counts of confirmed infections or unique victims. Pluto assessed that the operator was financially motivated and likely Chinese-speaking. `artificial intelligence` `technology` link

Backdoored Claude Code extension attempted SSH access On October 30, Phorion reported that the official Claude Code for VS Code extension on the Cursor Marketplace contained an SSH backdoor and persistence logic. Cursor removed affected versions 2.1.312 and 2.1.313. The implant attempted to establish persistent SSH access and profile hosts on macOS, Linux and Windows; Phorion identified a macOS implementation error that could prevent inbound SSH access if SSH was not already running. Phorion did not disclose information on the actor. `artificial intelligence` link

Digital foreign interference

Russia-aligned Operation Overload impersonation campaign targeting US midterm candidates On September 21, Graphika reported a likely coordinated Russia-aligned digital foreign interference operation, attributed to Operation Overload and the Foundation to Battle Injustice, targeting US midterm candidates. Inauthentic accounts across major social platforms seeded and amplified manipulated videos and posts impersonating trusted news outlets to push disparaging narratives against Democratic candidates. The activity aimed to undermine confidence in the US election and inflame political divisions. `russia` link

Disruption & destruction

Cisco reported exploitation of SSL VPN denial-of-service flaw On October 28, Cisco reported active exploitation of a Cisco Secure Firewall ASA and Secure Firewall Threat Defense remote access SSL VPN vulnerability. An unauthenticated threat actor can trigger unexpected device reloads, causing denial of service and disrupting remote access. link

Opportunistic

Cisco reported exploitation of Secure E-mail Gateway flaw CVE-2026-76519 On October 14, Cisco released an advisory for CVE-2026-76519, a critical arbitrary file write vulnerability in Cisco Secure E-mail Gateway. The vulnerability is due to improper handling of e-mail attachments when the file reputation service is enabled. A crafted attachment can enable arbitrary file writes and command execution with root privileges. Cisco reported active exploitation and released software updates. link

Active exploitation of critical FortiMail zero-day vulnerability CVE-2026-104286 On October 1, Fortinet released an advisory for CVE-2026-104286, a critical path traversal and NULL byte injection vulnerability (CVSS 9.8) affecting FortiMail versions 7.0 through 8.0. Fortinet reported exploitation in the wild and listed fixed releases 7.4.9, 7.6.7 and 8.0.2. Its advisory listed a forthcoming fix for the 7.2 branch, recommended migration for the 7.0 branch and provided workarounds. Fortinet said it remediated FortiMail Cloud on October 5. link

CISA listed GitLab CVE-2026-97229 as known exploited On October 29, the US Cybersecurity and Infrastructure Security Agency (CISA) added CVE-2026-97229, affecting self-managed GitLab instances, to its Known Exploited Vulnerabilities catalogue. GitLab issued security fixes and detections and reported exploitation attempts. Attempts reported against Alpine Linux on October 20 and 21 were not established as successful in that case. Administrators of affected instances should consult GitLab's advisory for remediation. link link

Known exploitation of Brotli decoder flaw CVE-2026-111497 On October 22, Imperva published a proof-of-concept demonstration of code execution using a Brotli decoder flaw, CVE-2026-111497, in Node.js. On October 27, CISA added the flaw to its Known Exploited Vulnerabilities catalogue based on evidence of active exploitation. The vulnerability is in the Brotli decoder, not solely Node.js; exploitation may affect implementations that decompress untrusted Brotli content. Imperva's demonstration alone is not evidence of an attack in the wild, and the cited sources do not identify a specific threat actor. link link

Exploitation attempts against Atlassian CVE-2026-21589 after proof of concept release On October 7, BleepingComputer reported probing and exploitation attempts against Atlassian CVE-2026-21589 shortly after publication of a proof of concept. The flaw can allow unauthenticated access to specific files within the web root of affected self-hosted Jira, Confluence and Bitbucket environments, provided the threat actor knows the exact file name and path. Researchers demonstrated a possible route to administrator access in certain Crowd-integrated deployments; observed probing does not establish successful escalation. link link link

Researchers reported PHP2Shell and Linux kernel flaw chain On October 26, Dreamgroup reported a chain of PHP2Shell and a Linux kernel vulnerability it calls Fragnesia (CVE-2026-46300), enabling root access and the implantation of PHP web shells. Dream attributed compromises of at least 25 PAN-OS firewall or management servers in the US, Japan, Taiwan and South Korea, including a US government-affiliated environment, to the activity. On October 28, CISA added the kernel flaw to its Known Exploited Vulnerabilities catalogue. The 25-server count is Dreamgroup's research finding, not a CISA estimate. `public administration` link link

OVHcloud command-line interface release assets replaced with malware On October 19, a threat actor replaced GitHub release assets for OVHcloud's command-line interface (CLI; ovh/ovhcloud-cli) versions v0.11.0 to v0.15.0 with a malicious build, using a compromised account. The exposure window lasted about 94 minutes, with OVHcloud removing the replaced assets on the same day and releasing a clean v0.16.0. OVHcloud stated the malicious binary was designed to retrieve secrets from disk; any execution during the window should be treated as credential exposure. `technology` link link

Browser vendors disclosed weak eMudhra intermediate CA keys On October 20, Mozilla and Google disclosed that private keys associated with eight eMudhra intermediate certificate authority (CA) certificates were weak or compromised following flawed key generation in April 2025. A technical disclosure said the keys could be recovered by testing fewer than 256 values. Six of the intermediates could issue TLS certificates, creating a risk of fraudulent website certificates. Firefox and Chrome blocked affected keys on October 17; protection on other platforms followed different schedules. link link link

Power BI phishing delivered unauthorised ScreenConnect software On October 7, Huntress reported a phishing campaign that used Microsoft Power BI links to lure victims into downloading unauthorised ScreenConnect remote monitoring and management tools. The threat actor used cloud hosting to increase the lures' credibility, then deployed multiple instances of the software and established persistent remote access. Huntress observed activity on multiple endpoints and identified additional affected systems through retrospective analysis. link

1. Conclusions or attributions made in this document merely reflect what publicly available sources report. They do not reflect our stance.

TLP definition

TLP	Disclosure	Message
RED	Not for disclosure, restricted to participants only.	Recipients may not share TLP:RED information with any parties outside of the specific exchange, meeting, or conversation in which it was originally disclosed.
AMBER+STRICT	Limited disclosure, restricted to participants' organisations.	Recipients may share TLP:AMBER+STRICT information only with members of their own organisation.
AMBER	Limited disclosure, restricted to participants' organisations and their clients.	Recipients may share TLP:AMBER information only with members of their own organisation and its clients.
GREEN	Limited disclosure, restricted to the community.	Subject to standard copyright rules, TLP:GREEN information may be distributed with peers and partner organisations within their sector or community, but not via publicly accessible channels.
CLEAR	Disclosure is not limited.	TLP:CLEAR information may be distributed freely.