Weekly Cybersecurity Briefing (9 November – 15 November 2026)

CyberSecBrief Team

CyberSecBrief provides clear, concise cybersecurity briefings that summarise the most relevant developments of the week. The report highlights notable incidents, key themes, and significant activity across the threat landscape, supported by daily updates sourced from trusted feeds and advisories. Its purpose is to offer a practical, time-efficient overview for readers who need a reliable snapshot of current cybersecurity events.
For more briefings, visit CYBERSECBRIEF.COM.



November 16, 2026

Editorial Analysis

The **Check Point gateway intrusions** centred on a flaw that was easy to underrate. CVE-2026-123039 scored 7.8 because it is a local vulnerability: the DynamicCLI process creates a world-writable socket at `/tmp/dynamic_socket` and executes any command it receives as root. Attackers obtained initial code execution through PHP2Shell in the Mobile Access portal, then exploited the flaw to place webshells and the TSGA and TSProxy backdoors. They moved over SSH to a VMware ESXi host and ran the `destroy_v6` wiper. Exploitation began on 30 October, about a week after a fix became available. It initially affected Israeli organisations, some destructively, while a wider wave on 2 November showed no destructive activity. Check Point says it observed exploitation against a couple of dozen customers, and CISA has since added the vulnerability to its KEV catalogue. Attribution remains unresolved. TSGA and TSProxy are PhiliKit variants previously linked by ESET to UNC5221, and the scanning hours align with Chinese working time. The wiper, Hebrew ransom note and Dream Security branding do not fit that profile. Check Point has not ruled out an initial access broker, a false flag or a smokescreen.

The **ScreenConnect campaign** has two competing accounts of how attackers gained access, and the sources have yet to agree. Huntress traced the activity to the RMM+ extension, whose versions 1.0.31 and 1.0.32 accepted a blank access key when none had been configured. It counted at least 288 hosts across 48 organisations affected within 90 minutes. Blackpoint says the recovered SCKit toolkit exploits CVE-2026-84869, a file-transfer vulnerability ConnectWise fixed in client version 26.6.5 in September. ConnectWise says that fix prevents the toolkit from spreading from endpoints to technicians, and Huntress has not observed that spread succeeding. ConnectWise's investigation into the initial entry method is ongoing. Microsoft's analysis highlights two practical details: updating the cloud console does not update technician machines, which receive the fix only through the client; and securing the legitimate console does not remove the second ScreenConnect client already installed on customer endpoints.

ARTEX was covered by four sources this week, including SentinelOne, the FBI and Datadog. Much of the reporting concerns activity after its developer withdrew the project. The repository was taken down on 8 October, but forks, archived copies and a commercial ARTEX PRO version remain available. More than 1,600 exposed consoles had been identified by early November. SentinelOne found campaigns costing between $145 and $1,500 in model time. Of the ten instances it analysed, three used American models. Guardrails triggered only on Opus 5.5, after which the operator switched to more permissive Chinese models. Figures for the ConnectWise activity vary by source. SentinelOne describes attempted sign-ins to more than 2,250 tenants through compromised ConnectWise applications, while the FBI cites a compromise that reached more than 2,250 downstream tenants. Huntress's earlier count of confirmed affected tenants was 162, with other reports citing 173.

© CyberSecBrief 2025—2026

Cursor's account of its breach sets out a clear chain of access. An access key issued to Fireworks was exposed and used to copy coding agent traces from a read-only storage bucket shared for model training. The traces came mostly from 2025 and from sessions outside Privacy Mode, with a small number containing unredacted credentials. Those credentials exposed API access for fewer than 0.01% of customers, allowing the attacker to reach archived non-production repositories and limited support data. On 29 October, a stolen credential was used to place malicious copies of the Claude Code extension on the cache serving Cursor's marketplace. The upstream Open VSX source never contained the malware, and up to 5,000 users may have received it over approximately 11 hours. Fireworks separately attributes its own breach to a third-party cloud provider hosting its R&D systems and links the activity to LLMJacking. Railway reported approximately $740,000 in fraudulent AI usage, while Notion confirmed that chat transcripts in a small number of workspaces had been exposed through Fireworks.

France's attribution of the October refrigeration attacks provides a firmer account of the incident. The C4 attributes the attacks on 11 and 12 October to Shahid Kaveh, a unit within the IRGC's electronic and cyber centre. It cites shared infrastructure, use of the FOFA scanning tool and Farsi-language traces in the attack infrastructure. The same infrastructure was used for reconnaissance against Schneider Electric equipment at small hydroelectric sites. France says it will respond using all available means, but has not specified what that will involve.

VulnCheck's seven-month review of Project Glasswing offers a critical assessment from a vulnerability-intelligence vendor. Anthropic has not commented in the material available. VulnCheck tracks 442 CVEs attributed to Anthropic, approximately 0.5% of those published this year, with two exploited in the wild. Anthropic's ledger lists 37,966 findings, of which 7,349 are public and 217 are marked fixed. The company's claimed 94.2% true-positive rate sits alongside 1,654 recorded duplicates. Curl has classified 19 of 34 published scanner reports as not representing vulnerabilities, confirmed six as low severity and is still assessing nine.

Highlights of the Week

France Blames Iran's IRGC for Supermarket Refrigeration Attacks

France's C4 attributed the 11 and 12 October attacks on refrigeration controllers at about 700 retail sites to Shahid Kaveh, a unit of Iran's IRGC electronic and cyber centre. Attackers raised defrost temperatures, causing major stock losses and at least 15 fires. The same infrastructure was used for reconnaissance against Schneider Electric equipment at small hydroelectric sites. France promised a response without specifying it.

Sources:

- ANSSI CERT-FR

Attackers Exploit Check Point Gateways to Deploy Backdoors and a Wiper

Attackers used PHP2Shell on Check Point Mobile Access portals, then gained root through CVE-2026-123039, a local flaw in the DynamicCLI interface. They installed the TSGA and TSProxy backdoors, pivoted over SSH to a VMware ESXi host and ran the destroy_v6 wiper. Activity began on 30 October against Israeli organisations; a wider wave on 2 November showed no destructive activity. Check Point observed exploitation against a couple of dozen customers.

Sources:

- Check Point Research
- Check Point

ScreenConnect Campaign Installs Rogue Clients and SCKit Toolkit at MSPs

Attackers used cloud-hosted ScreenConnect instances at managed service providers to push rogue clients and the SCKit toolkit, affecting at least 288 hosts across 48 organisations within 90 minutes, according to Huntress. RMM+ extension versions 1.0.31 and 1.0.32 accepted a blank access key. Microsoft reports SCKit builds RelayTrap and steals .ASPXAUTH cookies; ConnectWise's investigation of the entry method continues.

Sources:

- ConnectWise
- Huntress
- Microsoft Security Research

ARTEX AI Penetration Framework Used in Espionage and Cybercrime

SentinelOne reports at least five actors have used the ARTEX framework since late August, including VAPOR PANDA, which targeted over 200 organisations in 14 countries. The FBI links it to the theft of data on about 67,000 people from South Korean financial institutions and a ConnectWise Cloud Backup compromise reaching over 2,250 tenants. The developer removed the repository on 8 October; forks and a commercial version remain available.

Sources:

- SentinelOne
- Federal Bureau of Investigation
- Datadog Security Labs

Fireworks AI and Cursor Breaches Expose Credentials, Billing Data and Model Files

Attackers reached Fireworks AI's development machines on 24 October through a compromised third-party cloud provider, accessing account and billing data and, for some customers, secrets,

datasets and incomplete model weights. An exposed Fireworks-issued key let an attacker copy Cursor coding agent traces. On 29 October, malicious Claude Code extension copies reached Cursor's marketplace cache, with up to 5,000 users affected. Railway reported about $740,000 in fraudulent AI usage.

Sources:

- Fireworks AI
- Cursor

Threats

Turla Cereal Toolset Turns Exchange Servers Into Command Hubs

ESET documented Turla's Cereal toolset, which compromised Exchange servers at government bodies in Kazakhstan and Romania. The CerealNeuron backdoor listens for encrypted requests disguised as Exchange traffic, and CerealBox hijacks transport agents to copy emails, including NATO and EU correspondence in Romania.

Sources:

- ESET Research

Stolen GitHub Tokens Push Credential Stealer to Thousands of Repositories

Attackers using stolen GitHub tokens pushed malicious commits into more than 8,000 repositories, adding workflow code or npm postinstall hooks that steal secrets and create AWS and Azure access designed to survive rotation. Commits are reverted within about a minute, and poisoned versions reached one Packagist package and two Go modules.

Sources:

- SafeDep
- Aikido Security

FBI Disrupts North Korean FaceStealer Crypto Theft Scheme

The FBI and DOJ disrupted FaceStealer infrastructure that the bureau credits with more than $400 million in cryptocurrency theft over 18 months. Investigators seized over 900 victim webcam videos. The FBI attributes the activity to a North Korean group tracked as Sapphire Sleet and BlueNoroff, among other names.

Sources:

- Federal Bureau of Investigation

Infrastructure & Exploits

Microsoft Patches Record 1,235 Flaws, Including Exploited Outlook Zero-Day

Microsoft's November Patch Tuesday addressed 1,235 vulnerabilities, 66 of them Critical. An actively exploited Outlook spoofing zero-day lets a crafted message make Outlook connect to an attacker-controlled location without the usual warning, potentially exposing the user's NTLM hash.

Sources:

- BleepingComputer

Xen Orchestra Authentication Bypass Exploited for Admin Access

Vates disclosed VSA-2026-062, where a crafted empty-object login token authenticates as the first stored token's owner. Attackers can gain admin access, run code on managed pools and take over XCP-ng nodes as root. Exploitation is active, and fixes are in versions 6.10.1 and 6.9.2.

Sources:

- Vates

Apache HTTP Server HTTP/2 Flaw Gets Public Proof of Concept

CVE-2026-119618 is a critical race condition in mod_http2 affecting Apache HTTP Server 2.4.55 through 2.4.70, fixed in 2.4.71. A third-party proof of concept published on 12 November reproduces heap corruption but not code execution, and no in-the-wild exploitation has been confirmed.

Sources:

- Wiz
- Horizon3.ai

Tools & Techniques

Attackers Abuse GSocket Relay Tool to Keep Access After Patching

Threat actors deploy the open source GSocket tool on compromised Linux web servers, connecting outbound to a shared public relay network with no attacker-owned IP or domain. Devices with GSocket-related detections rose 356.2% between March and August 2026.

Sources:

- Microsoft

Editing Fewer Than Ten Neurons Can Backdoor AI Agent Models

Microsoft researchers showed that editing fewer than ten neurons in an open-weight model plants a backdoor without retraining. On Nemotron-Terminal-8B, seven edited neurons produced the attacker's command in at least 98.8% of triggered sessions; the attacker needs write access to the weights.

Sources:

- The Hacker News
- Microsoft Research

Policy & Legal

Europol Dismantles Gootloader, CastleRAT and StealC Infrastructure

Europol and partners seized 235 servers and 257 domains used to distribute Gootloader, CastleRAT and StealC. Authorities identified more than 6 million stolen credentials and remediated 681 Gootloader-infected websites, as part of Operation Endgame.

Sources:

- Europol

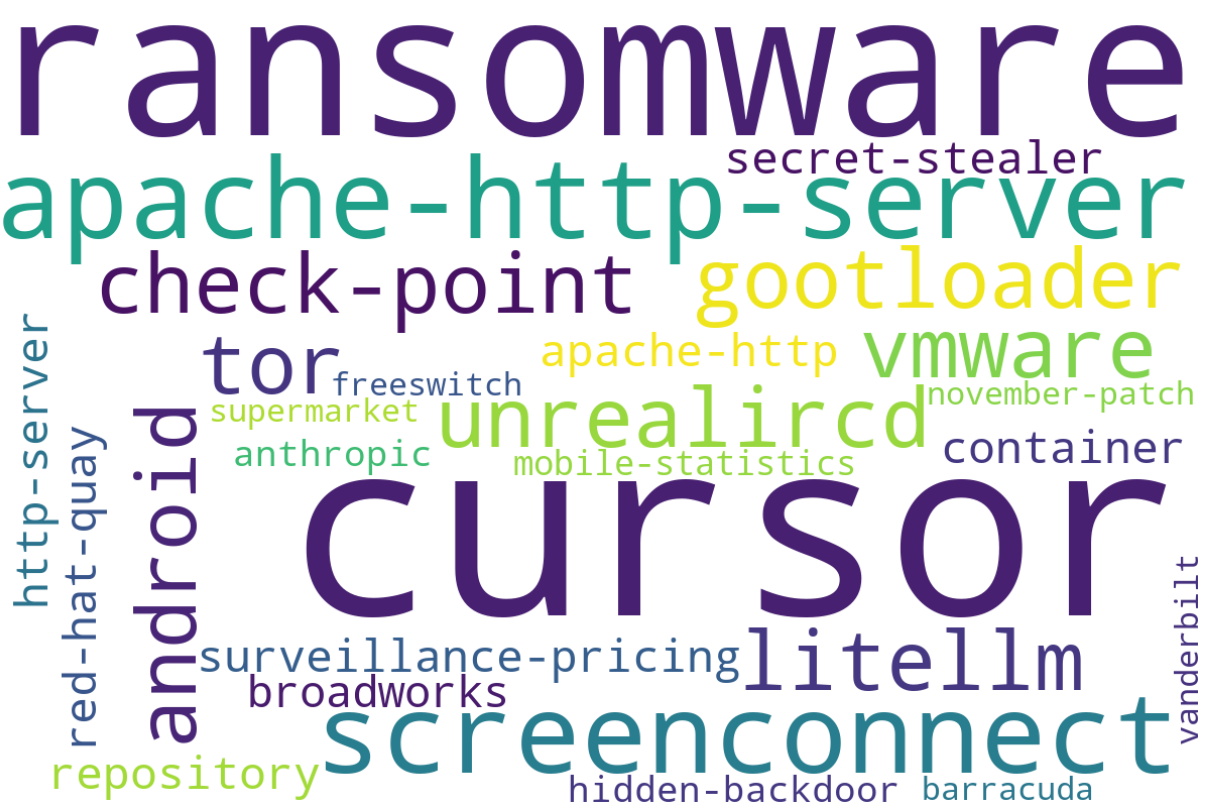
Ukrainian Police Dismantle Call Centres Defrauding German Investors of €4.3 Million

Ukrainian and German authorities dismantled call centres that posed as brokers and persuaded 135 German victims to invest through fake trading platforms, with losses above €4.3 million. Four members were detained, and the suspects face up to 12 years in prison.

Sources:

- Cyber Police Department of Ukraine

Weekly Topic Distribution



Copyright and Legal Notice

© CyberSecBrief. All rights reserved.

Reproduction, redistribution, or modification of this document or its contents without explicit written permission is strictly prohibited.

Disclaimer & Terms of Use

CyberSecBrief is a personal, non-commercial project operated from Germany and is not affiliated with any company or organisation. The information provided is intended solely for general informational and educational purposes and does not constitute professional cybersecurity, legal, or financial advice. Security-related information changes rapidly; although accuracy and timeliness are sought, no guarantee is given regarding completeness, correctness, or applicability. Users must verify details with official sources or qualified professionals before making decisions.

The owner disclaims liability for any indirect, incidental, or consequential damages arising from reliance on this material. External links and third-party feeds are offered for convenience; their content, accuracy, availability, and privacy practices are outside our control and are not endorsed. Sponsored or promoted material may appear, yet inclusion does not imply endorsement.

Original content is protected by copyright. Where external material is used, appropriate attribution is provided. Reproduction of original content without prior written consent is prohibited. Content may be updated, modified, or removed at any time without notice.

This disclaimer and the Terms of Use are governed by the laws of Germany and applicable European Union regulations. Should any provision be deemed invalid, the remaining provisions shall continue in full force and effect.

Credits and Sources

CVE entries are sourced from the CVE® Program, operated by The MITRE Corporation, and from the National Vulnerability Database (NVD), maintained by the National Institute of Standards and Technology (NIST). The NVD is provided “as is” without warranty of any kind.

Definitions for APT groups are derived from the MITRE ATT&CK® knowledge base, reproduced and distributed with the permission of The MITRE Corporation. © 2025 The MITRE Corporation.

All feeds and advisories are properly acknowledged and linked to their original sources where referenced, including within briefings.

Logo icon created by Ida Desi Mariana — Flaticon.

Powered by CyberFeedBites.

Source

This briefing is published by *CyberSecBrief* –
https://www.cybersecbrief.com/