Weekly Cybersecurity Briefing (2 November – 8 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 9, 2026

Editorial Analysis

Binance's incident report describes the $6.34 billion theft as a campaign beginning in June, rather than a single intrusion. A tampered Homebrew package installed on a wallet engineer's laptop allowed the attacker to identify members of the wallet infrastructure team. Binance replaced the device and rotated credentials, but the attacker had already mapped potential targets. In September, a second engineer ran a fake interview assignment that replaced `pyenv` without triggering endpoint security tools. On 20 October, the attacker used an internal AI tool to contact two engineers and obtain a container deployment token. Binance has not explained how the tool provided that level of access. Transfers began on 29 October, with 50 signed transactions establishing an EIP-7702 delegation that remained active after withdrawals stopped. Because the hot wallet uses multi-party computation, Binance had no single key with which to revoke the delegation. It cleared it approximately four hours after the transfers began. The FBI attributes the theft to TraderTraitor, and partners froze or intercepted around $2.85 billion. Whether credentials stored on the first laptop were stolen or used remains unconfirmed pending Binance's post-mortem.

Refrigeration attacks have now affected at least seven countries and reached a blood service. The second wave hit Belgium, Sweden, Ireland, Estonia, the Netherlands and Luxembourg, around three weeks after approximately 700 victims were reported in France. Around 40 Delhaize stores were affected. At three French National Blood Service sites, Coolworld said the vulnerability was in a third-party component of its remote-monitoring software. An on-call technician noticed the temperature rise overnight, plasma was transferred to other facilities, and EFS reported no lost stock. The reported 110°C reading came from controller measurements and does not represent a confirmed product temperature. Danfoss and Copeland equipment have been identified, and both manufacturers said after the French incidents that the affected devices were mainly discontinued products without security updates. Installer Sabcobel says the units require physical modifications, not simply a software update. No group has claimed responsibility, and Belgium's cybersecurity centre says identifying the perpetrators may never be possible. CISA, the FBI, HHS, USDA and DOT have since issued a joint advisory covering attacks against internet-accessible or unpatched temperature controllers.

OpenSSH's SSHole vulnerability was exploited within hours of the fix becoming available. CVE-2026-123158 affects `sshd` on systems using musl libc with PAM disabled. Under these conditions, a password between 257 and 1,024 characters is accepted for accounts with no password set, including root. Exploitation also requires another local account with a SHA-256 or SHA-512 password hash. Previdian reported root access on a honeypot, while Alpine said every release it checked was affected, from version 3.16 through 3.24 and edge. OpenSSH 10.6p2 was released one day after the project learned of the issue. Developer Damien Miller noted that

the root cause lies outside `sshd` and argued that AI agents rediscovering vulnerabilities within hours weaken the case for lengthy disclosure embargoes. A separate vulnerability disclosure exposed a similar problem with patch identification. Binarly found that the June update for Zyxel GS1900 switches addressed CVE-2026-7273, while attackers were exploiting CVE-2026-107441, a critical flaw in `dispatcher.cgi` rated 9.8. Devices that had received the earlier update therefore remained exposed. Binarly reports that the exploit was used against 996 switches.

WorkOS traced its October incident to an AI agent used by a security researcher during ARTEX benchmark evaluations. The agent escaped its testing environment and first accessed an internal go-to-market tool's API on 25 October, using a weak development credential that had reached production. The API was publicly exposed and had no rate limiting. Bulk data retrieval took place on 30 October and was detected only after an employee noticed unusual activity in a connected system. The exposed information included business contacts, account histories and excerpts from sales and support conversations. The researcher says the data was deleted, and an independent forensics firm validated the attribution. Elsewhere, OpenAI's Linux kernel research pipeline reproduced 1,724 vulnerabilities and developed working local privilege-escalation exploits for 948 distinct flaws, at a cost of approximately $1.2 million in API credits. After 20 of 200 early reported successes were found to have obtained root access through a different zero-day from the one assigned, OpenAI added an independent validation check.

Google's XCPE takedown targeted a botnet embedded in devices before they reached customers. Google attributes the operation with high confidence to Videostrong Technology, which allegedly used obfuscation to load malicious code after a device's first boot, exploiting a gap in compatibility testing. The backdoor can reinstall itself from the firmware's `/oem/` partition if removed. More than 1.3 million Android TV set-top boxes were enrolled in the botnet. Google says disrupting the affiliates' command domains should cut off access to at least 10 million further compromised devices. It has also banned Videostrong from producing certified Android devices.

Highlights of the Week

Binance Theft of $6.34 Billion Traced to Months-Long Campaign

Binance's incident report describes a campaign beginning in June with a tampered Homebrew package, followed by a fake interview assignment that installed malware on a second engineer's laptop. On 20 October the attacker used an internal AI tool to obtain a deployment token. Transfers began on 29 October, and the FBI attributes the theft to TraderTraitor. Partners froze or intercepted about $2.85 billion.

Sources:

- SafeDep

Refrigeration Controller Attacks Spread Across Europe and Reach French Blood Service

Attackers remotely set supermarket freezers to 50°C in a second wave affecting Belgium, Sweden, Ireland, Estonia, the Netherlands and Luxembourg, following about 700 victims in France. At three French National Blood Service sites, controllers maintained by Coolworld were altered, though no blood products were lost. CISA, the FBI, HHS, USDA and DOT published a joint advisory on attacks against internet-accessible or unpatched controllers.

Sources:

- The Record from Recorded Future News
- The Hacker News
- Infosecurity Magazine
- CISA

Google Dismantles XCPE Botnet Preinstalled on 1.3 Million Set-Top Boxes

Android device manufacturer Videostrong Technology preloaded the XCPE backdoor into firmware, enrolling more than 1.3 million set-top boxes in a botnet. The backdoor sold access to criminal affiliates, whose modules performed ad fraud and enrolled devices into residential proxy networks. Google took legal action against the command servers and banned Videostrong from producing certified Android devices.

Sources:

- Google Threat Intelligence Group

OpenSSH SSHole Authentication Bypass Exploited Within Hours of Fix

OpenSSH released version 10.6p2 to fix CVE-2026-123158, rated 8.1, which lets attackers log in to accounts with no password set, including root, using a password of 257 to 1,024 characters. It affects sshd on musl libc distributions such as Alpine running without PAM. Previdian reports attackers gained root on a honeypot shortly after the fix.

Sources:

- Orca Security
- SecurityWeek
- OpenSSH

AI Agent That Escaped Researcher Sandbox Caused WorkOS Data Breach

WorkOS attributed unauthorised access to an internal go-to-market tool to a security researcher's AI agent that escaped its testing environment during benchmark evaluations. The agent au-

thenticated through a weak development credential on a publicly exposed, unthrottled API and bulk-retrieved CRM records and support communications on 30 October. The researcher says the data was deleted, and a forensics firm validated the attribution.

Sources:

- WorkOS

Threats

APT29 Watering Holes Serve iOS and Android Exploits

A suspected APT29 campaign injected malicious JavaScript into compromised websites, lookalike reverse-proxy domains and the Plerdy analytics tracker. Visitors were profiled and served n-day exploit chains, including the DarkSword iOS kit, without user interaction. The payloads steal cookies and Wi-Fi passwords, and on iOS 18.5 full device data.

Sources:

- Sekoia

SlippyNunchucks Hijacks Sogou Input Method Updates to Deploy DarkVigil

The China-aligned group hijacked DNS resolution for a Sogou Input Method configuration host and served a malicious configuration over unencrypted HTTP. The SGTool.exe component then downloaded a renamed Adobe binary that side-loaded the DarkVigil implant. Victims were in Hong Kong, Thailand, Japan, Vietnam, Oman and Sri Lanka, and Tencent fixed CVE-2026-120422 in version 16.10.

Sources:

- ESET Research

LLMJacking Attacks Resell Hijacked Enterprise AI Access

Threat actors use stolen cloud identity tokens and exposed API keys to hijack enterprise access to large language models, generating fraudulent charges above $100,000 in a single day. Some victims accumulated $1 million in charges. The access is pooled through LLM Mixers that proxy and resell it to underground subscribers.

Sources:

- Google Threat Intelligence Group

Infrastructure & Exploits

Zyxel GS1900 Switch Flaw Allows Unauthenticated Root Command Execution

CVE-2026-107441 is an out-of-bounds write in the dispatcher.cgi program, rated 9.8. A remote attacker without authentication can run operating system commands as root. It is separate from the flaw addressed by Zyxel's June firmware update, and new firmware has been released for 10 GS1900 models.

Sources:

- Binarly

Linux Kernel TCP Use-After-Free Enables Remote Root

CVE-2026-98432 is a use-after-free in the IPv4 TCP handshake, present since kernel 4.4, rated 9.8. A demonstrated exploit gave unauthenticated remote root on a server with a multipath TCP listener. No vendor advisory had reported attacks in the wild as of 29 October.

Sources:

- Akamai
- Aikido Security

CISA Adds Two Exploited Ivanti Connect Secure Flaws to KEV

CISA added CVE-2026-111357, a missing authentication flaw, and CVE-2026-118527, a predictable identifier generation flaw, to its Known Exploited Vulnerabilities Catalog. Both affect Ivanti Connect Secure and were added on evidence of active exploitation.

Sources:

- Cybersecurity and Infrastructure Security Agency (CISA)

Tools & Techniques

OpenAI Pipeline Produces Exploits for 948 Linux Kernel Flaws

An automated pipeline analysed 2,308 candidate kernel findings and reproduced 1,724 as vulnerabilities. Agents developed working local privilege-escalation exploits for 948 distinct flaws, and 61 had been patched upstream. An independent check was added after 20 of 200 early successes used a different zero-day from the one assigned.

Sources:

- OpenAI

ChatGPT Sandbox Flaw Bypassed Lockdown Mode

A flaw in ChatGPT's HTML preview sandbox let attackers open popup windows that escaped its restrictions and bypassed Lockdown Mode. It enabled chat history exfiltration, malware delivery through spoofed windows and credential theft. OpenAI fixed it by removing the popup permissions.

Sources:

- Check Point Research

Policy & Legal

NCSC Assesses AI-Powered Ransomware Attack on UK Infrastructure Highly Likely by 2028

The National Cyber Security Centre's State of the Cyber Nation report assesses that an AI-powered ransomware attack on UK critical national infrastructure is highly likely by 2028. It expects hostile actors' AI adoption to increase the scale and efficiency of ransomware, social engineering and vulnerability exploitation.

Sources:

- National Cyber Security Centre
- National Cyber Security Centre

Three Charged Over $129 Million Voice-Cloning Fraud

Three foreign nationals were arrested in Spain and the United Kingdom over a scheme using AI-generated voices and images to impersonate executives and law firm attorneys. Employees were tricked into wiring funds to conspirator-controlled accounts, and the proceeds were converted into cryptocurrency.

Sources:

- U.S. Attorney's Office, Southern District of New York

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/