Weekly Cybersecurity Briefing (16 November – 22 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 23, 2026

Editorial Analysis

The **VMware vCenter vulnerability** was the week's most significant infrastructure concern. CVE-2026-123890 was exploited as a zero-day, exposing organisations running the virtualisation management platform before a patch was available. Its significance stems from vCenter's central role in managing virtualised environments: compromising this layer can have consequences across multiple virtual machines and applications. The incident also highlighted the risk posed by vulnerabilities in centralised infrastructure, particularly when exploitation is already under way before disclosure.

AI-related attacks took several forms during the week. Reporting on the ARTEX and Cairn campaigns, including activity involving MONAI Label, showed attackers using AI-assisted tools to target systems and workflows associated with AI development and deployment. The incidents are distinct, but each highlights the security of the infrastructure supporting AI applications, including their interfaces and underlying services. As these tools become embedded in operational environments, vulnerabilities in the surrounding systems can provide access well beyond the AI applications themselves.

Software distribution and trusted web infrastructure were also prominent targets. The compromise of Optimizely's JavaScript delivery mechanism affected a channel used to serve code to websites, while malicious releases of Microsoft's Agent Framework on PyPI introduced harmful packages into a legitimate developer ecosystem. The two incidents involved different distribution mechanisms, but both exploited trust in established software channels. The Optimizely compromise was particularly significant because access to the service's delivery infrastructure allowed malicious code to reach downstream websites without attackers having to breach each site individually.

The **arrest of suspects linked to a Ukrainian ransomware operation** provided a further development in efforts to pursue cybercriminals beyond their online infrastructure. The arrests do not establish that the wider operation has been dismantled, but they represent a tangible step in the investigation of the individuals behind it.

Highlights of the Week

VMware vCenter Zero-Day Exploited to Create Rogue Administrator Accounts

A suspected China-linked group exploited CVE-2026-123890 in VMware vCenter from at least early September, obtaining administrator sessions without passwords through a flaw in VMware Directory Service. The attackers extracted directory data and SAML signing keys, created rogue SSO administrators and moved on to ESXi hosts and virtual machines. Broadcom released fixes on 19 November.

Sources:

- Google Threat Intelligence Group

Optimizely Compromise Injects Malicious JavaScript into 178 Organisations' Websites

Attackers exploited Optimizely's Experimentation platform to modify customer JavaScript snippets, affecting 188 snippets belonging to 178 organisations. The injected code skimmed payment card details, replaced cryptocurrency wallet addresses and delivered Coruna and DarkSword exploit kits to some Apple mobile devices. Optimizely says it fixed the vulnerability and removed the malicious code.

Sources:

- Report URI
- Optimizely

AI Agents Used in Credential Theft and Zero-Day Intrusions

The ARTEX framework was used in campaigns that stole AI provider keys, searched for further credentials and attempted cryptocurrency exchange identity checks using deepfakes. Separately, Cairn agents exploited a previously unknown MONAI Label vulnerability to gain root access to a university server. Amazon attributes the activity to Chinese-nexus groups with moderate confidence.

Sources:

- Amazon Threat Intelligence
- Amazon Web Services

Malicious Microsoft Agent Framework Releases Steal Cloud Credentials

Two compromised versions of Microsoft's Agent Framework package, 1.26.1 and 1.26.2, were published to PyPI on 17 November. A malicious `.pth` startup file executed whenever Python launched, collecting environment variables, credential files and cloud identity tokens without requiring the package to be imported. PyPI quarantined both releases.

Sources:

- StepSecurity
- SafeDep

Ukrainian Police and FBI Target Ransomware Operation Linked to Former Conti Members

Ukrainian police and the FBI conducted 66 searches across Kyiv and 11 regions, targeting a group of at least 27 people linked to ransomware operations. Investigators say the attackers used AI-generated voice calls impersonating IT support to obtain credentials or install remote-access software before stealing data and deploying ransomware. At least 28 companies reportedly paid more than $40.5 million.

Sources:

- National Police of Ukraine

Threats

Com-Affiliated Groups Use Vishing to Steal Cloud Data

Groups including Bling Libra, Vortex and Cinder impersonated IT support staff to direct employees to adversary-in-the-middle phishing pages that captured credentials and MFA approvals. Attackers enrolled their own MFA devices, accessed Microsoft 365 data and sent extortion messages through Teams or email. Healthcare became the most targeted sector in September.

Sources:

- Unit 42

Retadup USB Worm Returns with CoffeeLoader Malware

A revived Retadup worm is distributing a CoffeeLoader variant after its command servers were re-registered following a 2019 law-enforcement takedown. The malware uses a patched 7-Zip library and a packer that decrypts its payload on the GPU through OpenGL. Infection attempts were blocked on around 15,000 systems this year.

Sources:

- Kaspersky

Infrastructure & Exploits

Java4Shell Affects OpenJDK Base64 Decoder on AVX-512 Systems

CVE-2026-114938 is an integer underflow in OpenJDK's optimised Base64 decoder that can cause a 62-byte heap overflow on supported x86-64 processors with AVX-512 VBMI. Researchers

demonstrated unauthenticated code execution paths in products including Jenkins, Jira, Confluence and Keycloak. Oracle released fixes for affected OpenJDK versions, with no in-the-wild exploitation reported.

Sources:

- LevelBlue
- VulnCheck

Critical NGINX HTTP/2 Flaws Allow Unauthenticated Code Execution

F5 patched two critical HTTP/2 vulnerabilities, CVE-2026-123633 and CVE-2026-123007, which can allow unauthenticated remote code execution. The older integer-underflow flaw dates to NGINX 1.9.5, and public proof-of-concept code is available. Fixes were released in NGINX Open Source versions 1.30.7 and 1.31.9.

Sources:

- Wiz
- Orca Security

Danfoss AK-SM800A Controllers Affected by Six Remote Code Execution Flaws

Researchers disclosed six pre-authentication stack-based buffer overflows affecting Danfoss AK-SM800A controllers running firmware up to and including 5.3.2. The affected process runs as root, and one flaw bypasses strict session control. Firmware 5.4.0 does not fix the vulnerabilities, and CVE identifiers were pending at disclosure.

Sources:

- Midnight Blue

Tools & Techniques

Fake Banking Apps Use SMS Theft and Call Forwarding

A malware-as-a-service operation targeting Indian banks offered fraudsters a panel with 50 tenant accounts, 34 of them paid. Fake Google Play pages collected card details before delivering Android malware that disabled Play Protect, stole SMS messages containing one-time passwords and redirected bank verification calls to attackers.

Sources:

- CloudSEK

Typosquatters Register Lookalike Domains Within One Working Day

Attackers registered 25 one-letter lookalikes of newly registered campaign domains belonging to companies including CarMax, McDonald's and Sanofi. Most substituted the genuine domain by removing a plural "s". The lookalike sites hosted parked pages, deceptive search portals and brand-themed blogs that directed visitors to scam websites through advertising links.

Sources:

- Symantec

Policy & Legal

Project Zero Moves Towards 30-Day Vulnerability Disclosure

Google Project Zero replaced its "90 + 30" disclosure policy with "Full X", setting disclosure deadlines at 90 days initially, 60 days from 2 February 2027 and 30 days from 3 January 2028. The revised policy removes grace periods and variant handling, with exceptions for actively exploited, already public or vendor-requested disclosures.

Sources:

- Google Project Zero

Anthropic Shortens Coordinated Disclosure Timelines

Anthropic is reducing its coordinated vulnerability disclosure window from 90 to 30 days for high- and critical-severity flaws. Reports made from 1 January 2027 will have a 45-day window for other issues. Escalation to an external coordinator will begin after 15 days without a maintainer response, rather than 30.

Sources:

- Anthropic

US Agencies Establish Threat Disruption Centre to Target Scams

The US Department of Justice and Department of Homeland Security established a Threat Disruption Center to coordinate victim reports, criminal-network intelligence and cooperation with private companies. The centre has moved to take down more than 1,600 scam websites, added over 6,600 scam-linked cryptocurrency wallet addresses to a screening network and identified 101 fraud networks.

Sources:

- The Record

Alleged 8-Digits Team Member Extradited to the United States

Roman Khlynovskiy, a Kazakh national accused of involvement with the 8-Digits Team ransomware group, appeared in a Tennessee federal court after extradition from Italy. Prosecutors allege the group targeted hospitals and schools in at least six US states, claimed more than 70 million health records and issued ransom demands exceeding $500 million.

Sources:

- US Department of Justice

Weekly Topic Distribution

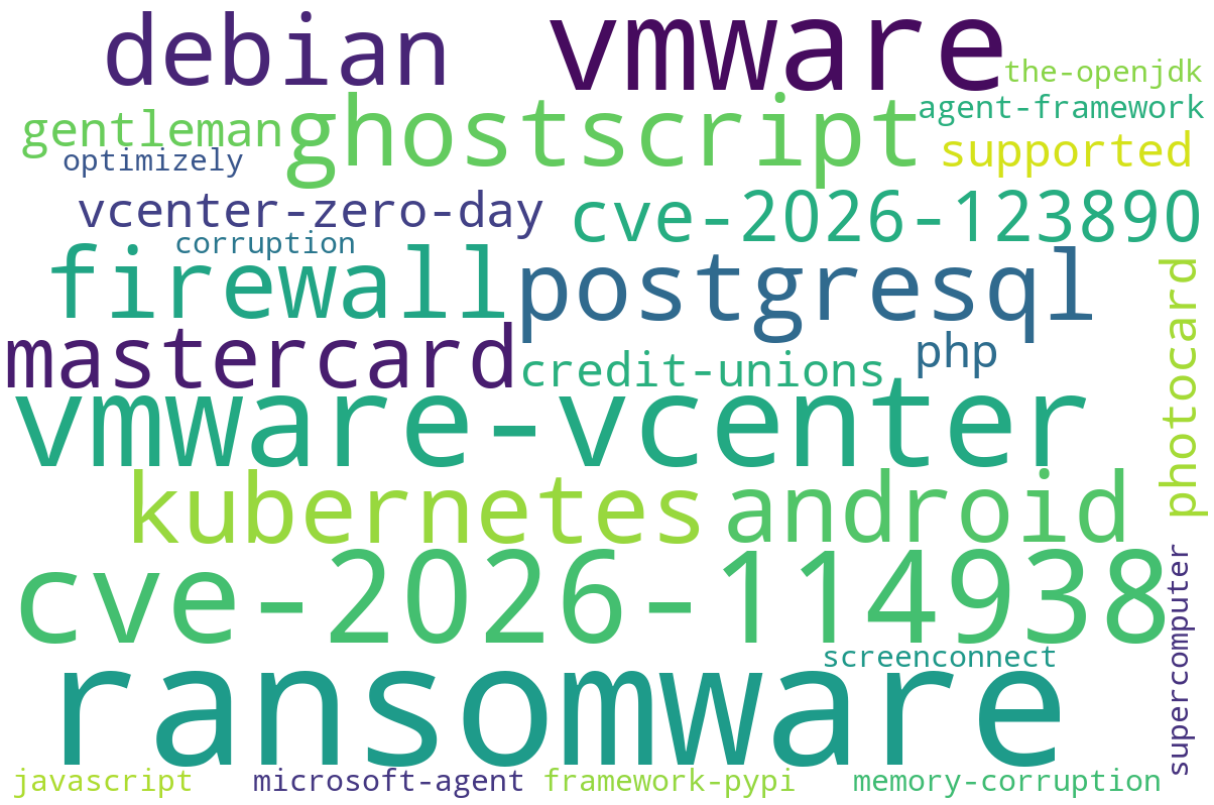

Copyright and Legal Notice

© 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/