
Weekly Cybersecurity Briefing (19 October – 25 October 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.



October 26, 2026

Editorial Analysis

PHP2Shell dominated the week, but the volume of reported exploitation attempts has yet to translate into evidence of widespread compromise. Amazon recorded an attempt against a honeypot on 17 October, three days before details leaked online and four days before PHP released fixes on 21 October. Cloudflare reported a rise from thousands of attempts to millions, while a UK minister told the House of Lords that 32 government services were unavailable, largely as a precaution. Confirmed compromises remain limited in the reporting available: Sophos identified one case involving the Potassium Mirai variant, while CrowdStrike found no successful exploitation in the wild or in its honeypots. WatchTowr found that much of the circulating proof-of-concept code crashes its target, and CrowdStrike says reliable code execution depends on matching the target's PHP, container and application versions. VulnCheck claims to have developed an exploit that bypasses ASLR on PHP 8.4.26, but has shared it only with customers, leaving the claim independently unverified. The disclosure process was contentious too. PHP says Chainguard published fixed images more than six hours before the coordinated release; Chainguard maintains that exploitation was already under way.

The **ConnectWise Cloud Backup** and **eMudhra** incidents exposed weaknesses in different forms of digital trust. In the ConnectWise case, attackers used the company's application identity to request Microsoft Graph tokens across customer tenants that had granted consent. Huntress reports that disabling the abused applications does not remove credentials added to other applications, and that, in one partner tenant, a ConnectWise Teams application granted itself Global Administrator privileges. Microsoft attributes the activity to Storm-0940, a China-aligned actor that it says used the Go-based ARTEX framework. The reported figure of more than 2,250 targeted tenants describes the scope of targeting, not confirmed compromise; Huntress has confirmed 162 affected tenants. At eMudhra, researchers recovered private keys for eight intermediate certificates from public certificates in under a second, with some recovered key values as small as 4, 56 and 115. Browsers blocked the intermediates within hours, but the company's explanation of a key wrap/unwrap flaw remains unconfirmed. eMudhra is withdrawing from the global TLS business, Chrome is expected to stop trusting the emSign ECC Root CA - G3 around 27 October, and approximately 30,800 unexpired certificates were still chained to that root as of 21 October.

Three separate incidents involving GitHub tokens demonstrated how stolen credentials can affect software distribution and cloud environments. **LiteLLM's** master key, committed to a public repository on 27 May, was never revoked after exposure was identified. It enabled the attacker to obtain a GitHub token belonging to an organisation owner and then exfiltrate tokens belonging to two maintainers. A temporarily relaxed approval rule allowed two stolen identities to approve the release of litellm 1.107.0, which remained available for about two hours. In

the **OVHcloud CLI** incident, the attacker did not need to alter workflows or tags: write access was sufficient to replace assets on five existing releases. A customer-pinned SHA-256 hash exposed the substitution within 94 minutes. The **Gheist** campaign extended token abuse across more than 13,500 planted branches in over 850 compromised accounts. GitHub's hold on suspicious workflow runs stopped the payloads in public repositories checked by Wiz, but audit logs and CloudTrail showed workflows completing in private environments. StepSecurity's follow-up analysis also narrowed the OVHcloud exposure: Homebrew rejected the replacement files because its cask pins hashes, while `install.sh` did not support arm64 Linux.

The **BlackX bootkit** was deployed on fully patched machines because Microsoft's Baton Drop mitigation does not revoke the vulnerable certificate automatically. Administrators must set a registry flag and create a scheduled task, and the planned automatic enforcement phase has not yet been implemented. Kaspersky attributes the campaign to APT41 with medium confidence; the targets were diplomatic agencies in Tajikistan, Uzbekistan and Turkmenistan. The week's **Brotli** disclosure raised a different problem. CVE-2026-111497 affects a vulnerable decoder present in the Python `Brotli` and `brotlicffi` packages, but those packages are not mapped to the CVE in advisory databases. As a result, tools such as `pip-audit` and Dependabot do not yet flag them. In this case, identifying affected software depends on recognising the vulnerable component even when automated advisory matching does not identify it.

Deep Dive: PHP2Shell, One Week On

PHP2Shell, the name given to CVE-2026-109618 and CVE-2026-109619, became one of the week's most prominent cybersecurity stories following their disclosure on 21 October. Both vulnerabilities received CVSS scores of 9.2 and affect PHP's core memory handling rather than a particular application. PHP powers a large proportion of websites, including installations running WordPress, Drupal, Magento and Joomla, giving vulnerabilities in the interpreter a potentially broad reach. A week into the response, however, the picture is more nuanced than the initial severity ratings suggested. The flaws were serious, but turning them into reliable remote code execution proved more difficult than the first wave of headlines implied.

The vulnerabilities

Both vulnerabilities involve PHP's handling of memory exhaustion and internal state. One relevant mechanism is `longjmp()`, which allows execution to jump out of a function without following the normal return path. When PHP encounters an out-of-memory condition, this abrupt unwind can leave internal data structures in an inconsistent state if an operation has not completed safely.

The vulnerable path involved `session.upload_progress`, a feature enabled by default since 2012. It processes session state during request handling, before the application has begun executing its own code. A crafted multipart request can trigger the vulnerable behaviour without

authentication or application-specific code, which helps explain the broad concern around the disclosure.

Reliable exploitation is more demanding than triggering the underlying error. An attacker must manipulate memory in a predictable way, including arranging a heap spray and accounting for build-specific offsets that vary with PHP versions, the C library and the container environment. At the time of disclosure, PHP's review of circulating proof-of-concept code found no working exploit, while CrowdStrike subsequently reported no successful exploitation in the wild or in its honeypots. These findings did not make the vulnerabilities harmless, but they did distinguish a remotely triggerable flaw from a reliable, broadly deployable exploit.

A long-standing problem, revisited

The underlying bug pattern had been identified as far back as 2004, when exploitation was considered impractical. An external audit by OSTIF and Quarkslab in 2024 raised related concerns again. A developer noticed the behaviour while addressing another issue and considered whether it might be exploitable, but no concrete exploitation path was established. Proposed fixes also caused compatibility problems, and the work stalled. The PHP Foundation later described this as a structural failure in how the issue had been handled.

The research that led to the 2026 disclosure built on that history. In August, a Calif researcher familiar with the earlier work used an AI model to help search PHP's source code for related patterns, identifying a vulnerable path through session upload progress. Wiz independently reached the same underlying flaws weeks later and identified additional vectors.

The significance is not that a model independently invented a new class of vulnerability. Previous researchers had already identified the underlying behaviour. AI-assisted analysis helped extend that work and find a practical exploitation path. The difficult remaining work, including chaining flaws and adapting an exploit to a target's memory layout, still required technical expertise. The episode nevertheless illustrates how research that once stalled for lack of a concrete exploit can become actionable when the cost of exploring related code paths falls.

The vulnerabilities' CVSS Attack Complexity rating of High is also relevant to how organisations interpret severity. It reflects the conditions required for successful exploitation, but does not mean the issue is unimportant or unlikely to be exploited eventually. At the same time, the history alone does not establish that the rating caused the earlier work to be deprioritised. The more defensible lesson is that a lack of demonstrated exploitation can leave technically serious problems unresolved when researchers cannot show a practical route from defect to impact.

Disclosure under pressure

The timeline helps explain why the response became contentious:

- 20 October, around 15:00 UTC: The two CVE identifiers appeared online alongside incorrect severity and release-date information. PHP brought the planned release forward by one day.
- 20 October, evening: A proof-of-concept demonstration video was published prematurely, intensifying efforts to reproduce the vulnerabilities.
- 21 October, 03:02 UTC: Cloudflare enabled Web Application Firewall rules to mitigate exploitation attempts.
- 21 October, 07:41 UTC: Chainguard published patched container images more than six hours before the coordinated release. PHP criticised the early availability, while Chainguard argued that exploitation was already under way.
- 17 October, reported retrospectively: Amazon disclosed that its MadPot honeypot had recorded an attempted exploit three days before the public leak. Amazon said it could not determine how the actor obtained the information during the embargo.

The pre-disclosure honeypot observation complicates any simple account of the leak as the starting point for attacker interest. It does not establish that the vulnerabilities were being successfully exploited, or explain how the actor obtained the relevant information. Likewise, the early availability of patched images created an opportunity for researchers and attackers to compare vulnerable and fixed code, but the evidence does not establish how much it accelerated exploitation.

The episode illustrates the difficulty of coordinating a disclosure once information begins to escape. Defenders need time to prepare mitigations and updates, while vendors and researchers must weigh the risks of releasing technical detail. An embargo can reduce that risk, but its effectiveness depends on information remaining contained and on participants coordinating their actions. PHP2Shell's timeline shows how quickly those conditions can change.

What the fix changes

PHP's fix, commit [01a40e89fb](#), changes how out-of-memory conditions are handled. Instead of immediately aborting execution through `longjmp()`, PHP records the memory-limit error and allows the current operation to return. A fatal error is then raised at a safe point, such as between virtual-machine opcodes, before `zend_call_function()`, after the main script or once request-body parsing has completed.

The change also retains a harder memory ceiling of approximately twice the configured `memory_limit`. If memory consumption continues to rise before execution reaches a safe point, the process can still be terminated. For the session-upload-progress path, PHP also prevents class resolution during the pre-request phase, blocking mechanisms such as autoloaders, `__wakeup()` and `__unserialize()` from running in that context.

The approach avoids removing `longjmp()` entirely, which could have introduced compatibility problems across PHP and its extensions. Instead, it changes when the error is acted upon and prevents potentially unsafe operations during the vulnerable request-processing stage. Having a concrete exploitation path gave developers a more specific basis for designing and assessing the fix than the 2024 investigation had provided.

The response and its costs

The initial response was shaped by two highly shareable facts: the vulnerabilities had CVSS scores of 9.2, and PHP runs a substantial proportion of the web. Headlines describing an internet-wide emergency spread quickly. Yet within days, CrowdStrike and WatchTowr were reporting that much of the circulating proof-of-concept material did not achieve reliable code execution.

That distinction matters for incident prioritisation. A high-severity vulnerability can warrant urgent patching even when reliable exploitation remains difficult, particularly when the affected component is widely deployed. But the urgency of remediation and the certainty of a successful attack are separate questions. Conflating them can make it harder for organisations to distinguish immediate evidence of compromise from precautionary work that is still necessary.

OX Security also reported a less visible consequence of the information vacuum: incomplete mitigations, AI-generated fixes containing errors and questionable detection scripts circulated on social media. This is a genuine risk during fast-moving disclosures, when defenders may be under pressure to act before reliable technical guidance is available. It does not mean early warnings were unjustified; it means the quality of the accompanying guidance matters as much as its speed.

Other developments during the same week offer useful context, though they are not direct equivalents. The PHP Foundation reserved CVE-2026-115560 for a separate remote-code-execution flaw that it described as requiring an unlikely combination of configuration and application code, and published mitigation advice rather than an emergency fix. GitLab issued a CVSS 10.0 fix covering a long history of releases, while the LiteLLM supply-chain compromise involved reported credential theft rather than a vulnerability awaiting exploitation. Their differing levels of coverage and response cannot be attributed to a single factor, but they underline how severity scores, practical exploitability, confirmed impact and public attention can diverge.

One week on

PHP2Shell is neither a story of mass exploitation proven by the available evidence nor a false alarm simply because reliable public exploits were initially scarce. The vulnerabilities exposed a long-standing weakness in core memory handling, and researchers eventually established practical paths that earlier investigations had not found. The response also showed how quickly

disclosure timelines can shift, and how incomplete technical information can complicate defenders' decisions.

The enduring question is how to handle that uncertainty. Organisations need to act on serious vulnerabilities before widespread exploitation is confirmed, but they also need reliable distinctions between a theoretical weakness, a demonstrated exploit, observed attempts and verified compromise. PHP2Shell is a useful case study in why those distinctions should inform the response without becoming an excuse to delay it.

Highlights of the Week

PHP2Shell Exploitation Attempts Reach Millions, With Few Confirmed Compromises

PHP released fixes on 21 October for CVE-2026-109618 and CVE-2026-109619, two unauthenticated remote code execution flaws rated CVSS 9.2. Amazon recorded an attempt on 17 October during the embargo, and Cloudflare saw attempts rise to millions. Sophos confirmed one compromise delivering the Potassium Mirai variant, while CrowdStrike found none in the wild. A UK minister said 32 government services were offline.

Sources:

- [SecurityWeek](#)
- [SecurityWeek](#)
- [The Register](#)
- [Sophos](#)
- [CrowdStrike](#)

ConnectWise Suspends Cloud Backup After Application Identity Abuse

ConnectWise suspended Cloud Backup, SaaS Security and Cloud Management after attackers used its Entra application identity to request Microsoft Graph tokens across consenting customer tenants. Huntress confirmed 162 affected tenants and reported that, in one partner tenant, a ConnectWise Teams application granted itself Global Administrator. Microsoft attributed the activity to Storm-0940, which it says used the Go-based ARTEX framework.

Sources:

- [The Register](#)

eMudhra Intermediate CA Keys Recovered in Under a Second

A researcher recovered private keys for eight eMudhra intermediate certificates from public certificates, including values as small as 4, 56 and 115. Mozilla blocked them through OneCRL,

Chrome used a CRLSet update, and eMudhra revoked them. eMudhra will exit the global TLS business, and Chrome is due to stop trusting the emSign ECC Root CA - G3 around 27 October.

Sources:

- [The Register](#)
- [Mozilla dev-security-policy](#)
- [The Hacker News](#)

Leaked Master Key Led to Malicious LiteLLM Release on PyPI

LiteLLM said a master key committed to a public repository on 27 May was used on 10 October to extract an organisation owner's GitHub token. The attacker stole two maintainer tokens, redirected the release pipeline to a lookalike fork and published litellm 1.107.0 on 21 October. The packages, live for about two hours, ran startup files that stole cloud credentials, SSH keys and AI tool tokens.

Sources:

- [LiteLLM](#)
- [SafeDep](#)
- [Socket](#)

OVHcloud CLI Releases Swapped With Credential Stealer

An attacker using a compromised contributor's GitHub account replaced binaries on five OVHcloud CLI releases on 19 October. The files were live for 94 minutes, downloaded about 250 times, and stole credentials for cloud services, SSH keys and AI-agent stores. GitHub release immutability was not enabled. StepSecurity found Homebrew installs were not exposed, and researchers linked the activity to GlassWorm.

Sources:

- [StepSecurity](#)
- [BleepingComputer](#)

Threats

Ransomware Groups Exploit PAN-OS GlobalProtect Flaw to Reach vCenter

Unit 42 observed Qilin and Inc exploiting CVE-2026-0257, a PAN-OS authentication bypass that allows GlobalProtect VPN sessions without authorisation. In one intrusion, attackers pivoted to a vCenter appliance, disabled ESXi security controls and dumped domain controller credentials from its virtual disk.

Sources:

- [Palo Alto Networks Unit 42](#)

Gheist Campaign Plants Secret-Dumping Workflows Across GitHub Accounts

Wiz recovered more than 13,500 planted branches from over 850 compromised accounts. Each carries a sync-check workflow that prints secrets into its log, and GitHub's approval hold blocks it in public repositories but not in private ones.

Sources:

- [Wiz](#)

Gentlemen Ransomware Uses Vulnerable Drivers to Kill Security Tools

Acronis reports the ransomware-as-a-service group has claimed more than 800 victims since mid-2025. A custom dropper loaded vulnerable Safetica and Lenovo drivers to terminate security processes before Go-based ransomware targeted Windows, Linux, NAS and BSD systems.

Sources:

- [Acronis Threat Research Unit](#)

Infrastructure & Exploits

Brotli Decoder Flaw Enables Code Execution in Node.js and Other Software

CVE-2026-111497 is an out-of-bounds write in Google's Brotli decoder, affecting Node.js, Bun, Envoy, Istio and Python packages Brotli and brotlicffi. Python packages remain unmapped in advisory databases, so pip-audit and Dependabot do not yet flag them.

Sources:

- [Wiz](#)
- [Aikido](#)

Plex Media Server Flaw Chain Exploited to Deploy XMRig

Researchers observed exploitation of four Plex Media Server flaws, CVE-2026-96652, CVE-2026-96654, CVE-2026-96655 and CVE-2026-96656, affecting versions before 1.43.3.10828. A honeypot attack on 17 October led to an XMRig download, and later activity deployed a proxy tool.

Sources:

- [BoredHackerBlog](#)
- [code-zm \(plex2shell\)](#)

Malware Distributed Through Rank Math PRO 3.0.123

Rank Math said PRO version 3.0.123 contained malware and was downloadable from 20:02 UTC on 15 October to 07:12 UTC on 16 October. The free plugin was unaffected, and the company released clean version 3.0.124.

Sources:

- [Rank Math](#)

Tools & Techniques

BlackX Bootkit Exploits Baton Drop on Fully Patched Systems

Kaspersky attributes BlackX to APT41 with medium confidence, with attacks on diplomatic agencies in Tajikistan, Uzbekistan and Turkmenistan. Machines were fully patched because Microsoft's fix requires manual certificate revocation through a registry flag and scheduled task.

Sources:

- [Kaspersky](#)

Planted AGENTS.md Files Hijack Coding Agents via Downloads Folder

Forever Security found a malicious AGENTS.md in a customer's Downloads folder that told Codex to install the npm package function-color. Codex ran the malware in over 90% of tests, and a crafted CLAUDE.md leaked Claude Code's OAuth token in about half of attempts.

Sources:

- [Forever Security](#)

AOMEI Backupper Promotional Channel Tampered to Deploy Backdoor

Attackers altered AOMEI Backupper remote configuration files for minutes at a time, causing the application to run a dropper through its promotional popup. It installed Radmin Server and an encrypted PocketPy implant, and Google and Mandiant track the activity as UNC7017.

Sources:

- [ManageEngine](#)

Policy & Legal

INTERPOL Operation Serengeti III Leads to 1,462 Arrests Across Africa

The operation across 21 countries disrupted 10,909 malicious networks and recovered over USD 105 million. Authorities identified 65,929 victims, and Uganda arrested 365 foreign nationals linked to two fraud rings.

Sources:

- [INTERPOL](#)


FBI and CISA Detail SafePay Ransomware Tactics Behind 600 Victims

A joint advisory with the US Secret Service and German and UK law enforcement says SafePay has affected at least 600 victims since August 2024. The group uses stolen passwords, VPN access and five known edge-device vulnerabilities, and encrypts files with ChaCha20 and RSA.

Sources:

- [CISA](#)

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/>