THE FORENSICS WAY

ISSUE 40 | FORENSICS AWARENESS THIS WEEK | TUESDAY, NOVEMBER 3, 2026

FRONT PAGE

EXPLOITATION & ANTI-FORENSICS

The Zero-Day That Scored 2.7: Warp Panda Weaponised a Logging Gap to Blind the Hypervisor Audit Trail

CrowdStrike disclosed on October 28 that the China-nexus actor it tracks as Warp Panda had used CVE-2026-41709 as a zero-day against US organisations, law firms among them. The flaw is an insufficient logging weakness (CWE-778) in VMware ESX: an actor already holding administrative privileges can carry out certain operations without those operations being written to the audit trail. Broadcom rates it 2.7 — Low — and fixed it on July 29 in VMSA-2026-0006 alongside four more serious bugs, after a report from CrowdStrike's Ian Barton. CISA added it to KEV on October 28 with a November 11 due date.

It grants no access and escalates nothing, which is exactly why it scored 2.7 and why most estates patched it in July without noticing. Its value to the operator is purely forensic. Warp Panda's pattern is initial access through an internet-facing edge device, a pivot to vCenter, then BRICKSTORM and SharpStorm implants, Microsoft 365 session-token theft and newly registered MFA devices — and the hypervisor log is the one record a guest-level compromise cannot reach. *Forensic priority:* CVSS measures what a bug does to the system, not what it does to your evidence. If a timeline across an early-2026 window rests on ESX audit data, absence of events is not evidence of absence — corroborate from vCenter, the SIEM and network telemetry the hypervisor does not control.

RESEARCH & MEDIA AUTHENTICATION

Ninety-One Per Cent Against Sixty-Nine, and Not One Wrong Answer: An AI Expert System Tested Against Twenty Examiners

Magnet Forensics launched the Magnet AI Research Lab on October 27, its stated purpose being to put the accuracy and reliability of AI-driven forensic output on a peer-reviewed footing. The first publication, in the *Journal of Forensic Sciences*, set fourteen video files — camera originals, edited clips, social media re-encodes, AI-generated media and files built to resist

THREAT BULLETIN

Active Exploit

CVE-2026-119977 — WordPress Core. Versions 4.7.0 through 7.1.3 stage a full core upgrade in predictable, web-reachable paths under *wp-content/upgrade*, and a time-of-check/time-of-use race there hands an unauthenticated attacker remote code execution during the seconds the package is unpacked. Fixed in 7.1.4 on October 28; Wordfence scores it 8.1 while WordPress's own advisory rates the race Low. CISA KEV November 3, due November 6. *Forensic Note:* The window is seconds long and no plugin is involved — hunt unexpected PHP files under *wp-content/upgrade* and requests to that path in web logs.

Active Exploit

CVE-2026-46300 — Linux kernel, tracked as Fragnesia. A local privilege escalation disclosed in May 2026 and added to KEV on October 28 with an October 31 due date. Dream reported a China-nexus operation chaining it behind last week's PHP2Shell flaws to take root on around twenty-five internet-facing Palo Alto Networks firewalls across the US, Japan, Taiwan and South Korea, planting webshells and pulling credentials from domain controllers; Unit 42 independently verified the tooling.

classification — against twenty trained examiners and the expert system in Magnet Verify, both measured against known ground truth. The system scored 91 per cent to the human average of 69, and made no incorrect classification at all.

Read the abstention rather than the headline gap. The system declined to reach a determination wherever the data did not support one, which is the behaviour a Rule 702 challenge actually probes, and the opposite of the confident wrong answer that makes AI output dangerous in a report. Two caveats belong in any citation of it: the study evaluates its sponsor's own product and the authors disclose Magnet funding, and fourteen files is a thin corpus to generalise from. *Forensic priority:* the finding worth carrying forward is not that a tool beat examiners but that published, falsifiable validation data now exists for one — which is what you will be asked for under cross-examination.

Forensic Note: This KEV entry carries a BOD 26-04 forensic triage requirement — triage the appliance before you patch it, because patching is a rebuild.

Malware Spotlight

Skyw4lker — ANY.RUN published analysis on October 27 of a device-code phishing kit aimed at Microsoft 365, with 618 sandbox submissions between June and October 21 and roughly 46 per cent of them from the US. It steals no password and drops no file: a lure document and CAPTCHA gate send the victim to Microsoft's genuine sign-in page to approve an attacker-initiated device-code request, while an obfuscated front end and PHP backend poll for completion and record IP and geolocation. ANY.RUN ties it to the Gabagool kit, likely the same developer. The only artefact is an Entra ID device-code sign-in from infrastructure the user never visited.

THE FORENSICS WAY

HISTORICAL ARCHIVES HISTORY OF DIGITAL INVESTIGATION TUESDAY, NOVEMBER 3, 2026

CASE STUDIES: HISTORICAL GRID

OLYMPIC DESTROYER (2018)

The Record That Lied: Forged Artefacts Built to Send

TJX (2005–2007)

“We May Never Be Able to Identify”: The Breach Whose Scope Was Never Established

Investigators After the Wrong Actor

During the opening ceremony of the Pyeongchang Winter Olympics on February 9, 2018, a wiper took down the Games' IT systems, the official ticketing site and Wi-Fi across venues. What made it a landmark is not the destruction but the deception: the malware carried forensic artefacts deliberately forged to implicate someone else, including a Rich header crafted to match known Lazarus tooling, so the first wave of attribution pointed at North Korea. Kaspersky's analysts unpicked the fake by establishing that the header was inconsistent with the rest of the binary — it had been grafted on. It is the case that forced the discipline to treat an artefact's provenance as a question in its own right rather than as a fact, and the reason attribution from tooling similarity alone no longer survives review.

Intruders reached TJX's network in July 2005 through a weakly protected wireless link at a store, and were not detected until December 2006. The company eventually reported data taken for at least 45.6 million payment cards; banks in later litigation asserted the true figure exceeded 94 million. The gap between those numbers is the forensic story. TJX told the SEC it believed it might never be able to identify much of the information taken, that the intruder's technology had made it impossible to determine the contents of most of the files stolen in 2006, and that it had itself deleted many of the affected files in the ordinary course of business. Eighteen months of intrusion produced a case in which the central question — what was actually taken — was never answerable, because the records that would have answered it were gone before anyone looked.

SHELLSHOCK (2014)

Twenty-Five Years in the Shell: The Pre-Auth Flaw That Made Web Logs the Only Witness

Disclosed on September 24, 2014, CVE-2014-6271 let bash execute commands appended to an environment variable's function definition, turning anything that passed attacker-controlled data into a shell's environment — above all CGI — into a pre-authentication execution path. The bug had been present for roughly twenty-five years; mass scanning began within hours and five further CVEs followed within a week. Execution happened inside the web server's own process and dropped nothing by default, so the first instruction in 2014, as now, was to preserve request logs before they rotated.

DIGINOTAR (2011)

Five Hundred Fraudulent Certificates: The Investigation That Ended a Certificate Authority

In mid-2011 an intruder reached the certificate infrastructure of the Dutch CA DigiNotar and issued at least 531 fraudulent certificates, including one for **.google.com* later used against an estimated several hundred thousand Iranian users. The Operation Black Tulip investigation found the systems poorly segmented and the issuance records unreliable, so investigators could state only a floor, never a total. Browsers removed DigiNotar within days and the company was bankrupt within weeks.

TOOLS OF THE TRADE

OPEN SOURCE

Edge DFIR Toolkit

EDGE-DFIR-TOOLKIT — NEW

OPEN SOURCE

ALEAPP

V2026.4.5 — BATTERY STATS DAILY

COMMERCIAL

Forensic OSINT

V3.8.1 — ASSISTED CAPTURE

UTILITY / SCRIPT

Unfurl

HINDSIGHT FOUNDRY — FBCLID

Triage material for compromised VPN concentrators, firewalls and gateways across ten platforms including FortiOS, Ivanti Connect Secure, PAN-OS and NetScaler: 92 catalogued artefacts, 142 Sigma and YARA detections and per-CVE exploitation timelines. Its playbooks open with explicit “do not” lists, which is the right instinct for a device you cannot image. An early release — check its per-CVE citations against the vendor advisories.

Parses *batterystats-daily.xml*, which research published in October shows retains about ten days of daily discharge and charge steps alongside package install and update records on Android 11 to 16 — plain XML through 13, ABX from 14. A correction in early October reports the F/O/D/Z screen-state fields as the raw AOSP bit value rather than the earlier prose label.

The October 27 release discovers an Instagram profile, then captures the posts you approve — expanding comment threads and cycling carousels within a date range, across Grid, Tagged, Reels and Stories, with TikTok and LinkedIn also covered. Each captured file is SHA-256 hashed and comment counts are checked against the figure the platform itself reports, so the capture carries its own completeness evidence.

Ryan Benson's URL breakdown tool now parses modern Facebook and Instagram click identifiers. Values issued since April 2024 run 100 to 190 characters and can carry a click timestamp and application identifiers, where the legacy 61-character form carried none — so a shared link in a chat thread may date the click as well as the message.

PRIORITY CVE ADVISORIES

ARCHIVE ALERT

CVE-2026-111497 (GOOGLE BROTLI)

Out-of-bounds write in the Brotli decoder across versions 1.0.3 to 1.2.0, reachable wherever a service decompresses attacker-supplied streams; CVSS v4 9.2, CISA KEV October 27. *Forensic Note:* Express 5 decompresses by default, so the exposure is often a dependency nobody chose.

ARCHIVE ALERT

CVE-2026-110447 (COLLABORA ONLINE)

OS command injection in the built-in CODE Server bundled with Nextcloud, reachable when the CODE server is not running; CISA KEV October 23. *Forensic Note:* The vulnerable path is the error handler, so a disabled feature is not a mitigation.

RECENT MALWARE WATCH

ARCHIVE: OCTOBER 27, 2026

ARTEX — THE AGENTIC FRAMEWORK BEHIND A SUPPLY CHAIN COMPROMISE

Microsoft tied the ConnectWise Cloud Backup compromise to Storm-0940 driving an open-source agentic AI pentest framework, reaching privileged cloud access in under three hours and authenticating through already-authorised multi-tenant applications. No interactive sign-in means no sign-in log.

ARCHIVE: OCTOBER 20, 2026

PEERSTITCH — THE IMPLANT WITH NO FILE, PORT OR CALLBACK

Unit 42's memory-resident ARM64 FortiGate implant, running inside the miglogd and iked daemons via ptrace with no persistent binary on disk, used to exfiltrate more than 500 firewall configurations. Capture process memory before any reboot.