MAG212CYBER

WORLD REPORT
WEEKEND EDITION · ISSUE 02



ANALYST EDITION · OCT 21–24, 2026

AI AGENTS STORM THE CLOUD

Eight threat briefings from the MAG212 analyst desk: supply-chain sabotage, OAuth identity abuse, exploited edge systems, and what to fix first.

INSIDE

- VeloCloud zero-day (KEV)
- Storm-0940 vs Microsoft 365
- LiteLLM supply-chain attack
- Nextcloud web shells (KEV)
- Phantom AI ransomware ops
- Session-token theft
- k0t edge credentials
- GitLab + Brotli

OCT 24, 2026 | MAG212 CYBER | EL MOSTAFA OUCHEN

FROM THE ANALYST DESK

Trust is the new attack surface

MAG212 CYBER

This week, AI agents went from conference talk to incident reports. Microsoft attributed a supply-chain campaign that abused ConnectWise's cloud apps to the China-aligned actor Storm-0940. The actor drove the ARTEX multi-agent AI framework and in some cases reached privileged cloud access in **under three hours** (Microsoft). Separately, SOCRadar exposed ransomware affiliates running an "AI-executed, human-approved" intrusion kit against about 62 organizations (The Hacker News).

Meanwhile, CISA gave federal agencies just **three days, until October 26**, to patch two exploited flaws: an Arista VeloCloud Orchestrator zero-day and a Collabora/Nextcloud command injection (The Hacker News). The popular AI gateway LiteLLM shipped an infostealer through its own release pipeline.

The common thread for MSPs and IT teams is identity. OAuth apps, session tokens, VPN credentials and leaked CI tokens did more damage this week than malware exploits did. Stories from our October 21 Special Edition (PHP2Shell, eMudhra, Plex2Shell, PAN-OS) appear as short follow-ups, because the news this week is somewhere else. One of those follow-ups has a hard date: Chrome stops trusting the eMudhra root on **October 27**.

The MAG212 CYBER Analyst Desk · Edited by El Mostafa Ouchen · Published by Dark Raven Labs

2,250+
Microsoft 365 tenants targeted via ConnectWise apps

79
organizations hit by session-token theft

70,000+
edge-device credentials exposed by k0t

49,000+
estimated downloads of poisoned LiteLLM

11.7M
CAPCOM ID email addresses exposed

3 days
to patch two KEV flaws (federal deadline)

IN THIS EDITION

- ▸ Priority matrix and deadlines
- 01 VeloCloud Orchestrator zero-day
- 02 Storm-0940 and the ConnectWise apps
- 03 LiteLLM 1.107.0
- 04 Nextcloud + Collabora CODE
- 05 Phantom framework
- 06 Microsoft 365 session-token theft
- 07 k0t firewall and VPN campaign
- 08 GitLab and Brotli
- \+ Breach roundup · Follow-ups · Quick hits · Glossary

THE COMMON THREAD

Non-human identities

Vendor apps, service principals, CI tokens, session tokens and AI-gateway keys did more damage this week than malware exploits. Our analysts' advice: protect them like administrator accounts.

PRIORITY MATRIX

What to fix first

MAG212 CYBER

Our analysts rank this week's threats by confirmed exploitation, exposure and how much access a successful attack gives. **P0** means act this weekend; **P1** means this week.

#	Threat	Rating	Exploited	Priority	First action
01	Arista VeloCloud Orchestrator zero-day CVE-2026-114625 (KEV)	Critical	Yes (zero-day)	P0 · patch by Monday	Upgrade on-prem VCO to 5.2.3.18 / 6.1.3.10 / 6.4.2.12 and restrict UI access
02	Storm-0940 uses ARTEX AI agents against M365 via ConnectWise apps	Critical	Yes (active campaign)	P0 · review this weekend	Keep ConnectWise app connections disabled; audit multi-tenant app IDs, consents and role grants
03	LiteLLM 1.107.0 supply-chain infostealer (plus Railway secrets incident)	Critical	Yes (stolen keys abused within hours)	P0 for anyone who installed it on Oct 21	Find any 1.107.0 install, rebuild on 1.106.2, rotate every secret
04	Collabora CODE / Nextcloud command injection CVE-2026-110447 (KEV)	Critical	Yes (since Sep 10)	P0 · KEV due Oct 26	Update richdocumentscode to 26.4.303+ or remove it; hunt for `shx*.php` web shells
05	"Phantom" AI intrusion framework used by INC / The Gentlemen affiliates	High	Yes (about 62 organizations targeted)	P1 · this week	Patch FortiClient EMS, SimpleHelp and cPanel/WHM; eliminate password reuse
06	Microsoft 365 session-token "spraying" hits 79 organizations	High	Yes (79 organizations)	P1 · this week	Enforce Conditional Access; audit registered devices, MFA methods and inbox rules
07	k0t firewall/VPN campaign: 70,000+ credentials, ransomware links	High	Yes (70,000+ credentials)	P1 · this week	Patch CVE-2025-64446 / CVE-2026-0257; reset VPN and LDAP credentials on exposed devices
08	GitLab max-severity RCE CVE-2026-110264 and Brotli "Node2Shell" CVE-2026-111497	Critical / High	No (not yet)	P1 · patch before exploitation	Upgrade GitLab; update Brotli in Node.js, Python, Envoy and Istio, then rebuild images

HARD DEADLINES

Oct 26
CISA KEV federal due date: Arista VeloCloud CVE-2026-114625 and Collabora CVE-2026-110447

Oct 27
Chrome stops trusting the eMudhra emSign ECC Root CA - G3

Nov 3
KEV due date for Plex CVE-2026-96652 and CVE-2026-96655

Jan 12, 2027
Windows Server 2016 reaches end of support

KEV due dates bind US federal agencies; for everyone else they are a strong signal of real-world exploitation. Ratings and priorities are MAG212 analyst judgments based on the sources cited in each feature.

FEATURE 01 · EDGE · SD-WAN

Third VeloCloud Orchestrator zero-day in four months hits on-prem SD-WAN

Arista has patched another actively exploited SQL injection in on-prem VeloCloud Orchestrator, and CISA wants it fixed by Monday.

What happened. Arista found the flaw itself and confirmed it was exploited before a fix existed. Patches came on October 22 (BleepingComputer; SecurityWeek). CISA added it to the Known Exploited Vulnerabilities catalog on October 23 with a three-day federal deadline. This is the **third VCO flaw exploited in the wild since July**. The earlier two were CVE-2026-16812 (CVSS 10.0, July) and CVE-2026-93952 (CVSS 10.0, September) (The Hacker News). Arista has not named an actor or given victim counts.

Why it matters / client impact. VCO is the management brain of VeloCloud SD-WAN. Whoever controls it can see and reshape branch connectivity for every tenant it manages. That makes it a high-value target for any MSP or multi-site business that self-hosts it. Arista says VCO is exposed by default and no configuration prevents that exposure (SecurityWeek). Three exploited zero-days in four months means attackers are actively researching this product, so expect more.

How the attack works. The attacker needs network access to the VCO web interface and *any* valid session, even the lowest tenant role ("Enterprise Read Only"). The SQL injection then lets them gain operator privileges, impersonate the local "superuser operator" and reach internal-only functions (SecurityWeek). Arista has not said how attackers got their initial sessions. Stolen or reused tenant credentials are a plausible route, but that is our inference, not a confirmed fact.

EXPLAINER

What is an SD-WAN orchestrator?

SD-WAN connects a company's branch offices over ordinary internet links. The orchestrator is the central console that pushes configuration and policy to every site. For an MSP it can manage many customers at once, so it deserves the same protection as a domain controller.

What to do now

1. Upgrade on-prem VCO to 5.2.3.18, 6.1.3.10 or 6.4.2.12 today. Until you can, restrict the VCO UI to management networks and block 74.82.197[.]86 and 38.54.50[.]151 (SecurityWeek).
2. Preserve web, backend, database and system logs *before* you remediate. Then hunt for high request rates, URL-like or encoded path components, references to internal services, backdoor daemons and web shells, and block unneeded outbound ports.
3. Review operator and admin activity, and audit low-privilege tenant accounts. Any valid session is enough to start the attack chain.

ANALYST ASSESSMENT

P0 · patch by Monday

EXPLOITED
Yes (zero-day)

CONFIDENCE
High

WHO SHOULD ACT
MSPs and multi-site businesses that self-host VeloCloud Orchestrator

Three exploited VCO zero-days since July shows sustained attacker interest in this product. Because any valid low-privilege account is enough to start, weak or shared tenant credentials raise the risk. Patch first, then review accounts and logs.

KEY FACTS

CVE / ACTOR
CVE-2026-114625, CVSS 9.9, SQL injection. No actor named. Attack IPs geolocate to Japan.

AFFECTED
On-prem VeloCloud Orchestrator (VCO) releases before 5.2.3.18, 6.1.3.10 and 6.4.2.12. Hosted VCO is already patched.

FIXED IN / STATUS
5.2.3.18, 6.1.3.10 and 6.4.2.12. Exploited as a zero-day. Added to CISA KEV Oct 23, federal due date **Oct 26, 2026**.

DATE
Patches Oct 22. KEV addition Oct 23.

VERIFICATION CHECKLIST

Tick each item before closing the ticket

- [ ] Confirm the on-prem VCO version is 5.2.3.18, 6.1.3.10 or 6.4.2.12 or later.
- [ ] Confirm the VCO web interface is reachable only from management networks or a VPN.
- [ ] Export and keep web, database and system logs before any rebuild.
- [ ] List every tenant and operator account; remove unused ones and reset shared passwords.
- [ ] Check logs for the two IPs Arista published and for unusual request volumes.

Sources · SecurityWeek, Oct 23 · BleepingComputer, Oct 22 · The Hacker News, Oct 24

FEATURE 02 · CLOUD IDENTITY · NATION-STATE

China-aligned Storm-0940 turns ConnectWise cloud apps and AI agents on Microsoft 365

Microsoft says a state-aligned actor used stolen ConnectWise app credentials and the ARTEX multi-agent AI framework to try to get into more than 2,250 tenants. Some intrusions reached privileged access in under three hours.

What happened. Microsoft attributes the ConnectWise multi-tenant app compromise to Storm-0940. The actor used credentials stolen from ConnectWise cloud apps to try to sign in to **more than 2,250 tenants**. Because it was app access, no interactive user sign-in was needed. Several intrusions reached privileged cloud access in **under three hours**. That included an application assigning itself Global Administrator, new credentials added to OAuth apps, and bulk reads of Key Vault secrets (Microsoft). The Record reports targets in government, aerospace, defense, finance, manufacturing, technology and energy (The Record). The ARTEX developer, whom the WSJ identified as Li Puhua, says the tool was misused and made it closed-source on October 8. China's Foreign Ministry denied knowledge (The Record).

Source discrepancy: Microsoft cites a password spray between October 10 and 14 that hit 327 accounts at 8 organizations with more than 151,000 failures and no successes (Microsoft). The Record cites 2,900 sign-in attempts per minute and more than 111,000 failed logons on October 11, and describes "hundreds" of organizations targeted (The Record). These figures likely describe different measurement windows.

Why it matters / client impact. This is the attack MSPs worry about most: a trusted vendor app already holding consent in hundreds of customer tenants becomes the way in. MFA on user accounts does nothing when the attacker authenticates as an application. AI-driven automation also compresses the time from first access to Global Admin into an afternoon, faster than most manual SOC escalation paths.

How the attack works. The attacker takes stolen app credentials and authenticates as the ConnectWise service principal in each tenant where it is consented. Then ARTEX agents enumerate permissions and escalate: they add credentials to OAuth apps, assign roles and read secrets. Alongside this ran a high-volume password spray against user accounts (Microsoft).

EXPLAINER

Why user MFA did not stop this

When you connect a SaaS product to Microsoft 365, you grant permissions to an enterprise application. That app signs in with its own secret or certificate, not with a user's password and MFA. If the vendor's app credentials are stolen, the attacker can act as the app in every tenant that granted it access, with whatever permissions it was given. App governance, not user MFA, is the control that matters.

What to do now

1. Inventory enterprise apps by application ID. Keep the affected ConnectWise connections disabled until ConnectWise issues guidance, and preserve evidence *before* deleting any enterprise app.
2. Review service-principal sign-ins from October 1 onward. Audit credential, consent and role changes, especially self-assigned Global Administrator and bulk Key Vault reads.

ANALYST ASSESSMENT

P0 · review this weekend

EXPLOITED

Yes (active campaign)

CONFIDENCE

High (Microsoft primary report)

WHO SHOULD ACT

Tenants that connected ConnectWise cloud apps; all MSP-managed tenants

The lesson for every IT team: third-party apps with broad consent are privileged identities. Inventory them, scope them to the minimum, alert on any change to their credentials or roles, and make sure you can disable a vendor app within minutes.

KEY FACTS

ACTOR

Storm-0940, a China-state-aligned group active since 2021. Tool: ARTEX, an open-source multi-agent LLM pentest framework.

AFFECTED

Organizations that granted ConnectWise multi-tenant cloud apps (Cloud Backup, MS Teams, MS Graph) access to Microsoft 365 or Google Workspace.

STATUS

ConnectWise shut down several services. Microsoft says to keep the app connections disabled until ConnectWise issues restoration guidance. Microsoft says its own products were not compromised.

DATE

Microsoft report Oct 21, updated Oct 23. Activity since as early as September 2026.

3. Block the published indicators: 144.172.64[.]128, 144.172.64[.]99, 144.172.64[.]60 and VPN IPs 187.15.88[.]186, 187.15.90[.]217, 187.15.89[.]186. Alert on the user agents Python-urllib/3.12, python-requests 2.31.0/2.34.2, curl/7.88.1 and azure-toolkit/1.0. Apply least privilege to every multi-tenant app (Microsoft).

VERIFICATION CHECKLIST

Tick each item before closing the ticket

- [ ] Keep the affected ConnectWise app connections disabled until ConnectWise publishes restoration guidance.
- [ ] Export the list of enterprise apps with their permissions and owners before changing anything.
- [ ] Review Entra audit logs since Oct 1 for new app credentials, consents and role assignments.
- [ ] Check who holds Global Administrator and remove any application or account that should not.
- [ ] Review Key Vault access logs for unusual bulk reads; rotate secrets if in doubt.
- [ ] Block the indicators Microsoft published and apply least privilege to every multi-tenant app.

KQL · Microsoft Sentinel (verify table and field names in your workspace)

```
// Changes to app credentials, consents and role membership since Oct 1
AuditLogs
| where TimeGenerated > datetime(2026-10-01)
| where OperationName has_any ("Add service principal credentials",
        "Consent to application", "Add member to role",
        "Add app role assignment to service principal")
| project TimeGenerated, OperationName, InitiatedBy, TargetResources
| order by TimeGenerated desc
```

Sources · Microsoft Security Blog, Oct 21/23 · The Record, Oct 22 · SecurityWeek · CyberScoop

FEATURE 03 · SUPPLY CHAIN · AI INFRASTRUCTURE

LiteLLM's own release pipeline ships an infostealer, and Railway tells developers to rotate secrets

A leaked, never-revoked key let attackers publish a malicious LiteLLM 1.107.0 to PyPI and Docker. The package harvested the cloud and AI-provider keys that AI gateways hold.

What happened. On October 21 (UTC), attackers overrode branch protection and pointed LiteLLM's release pipeline at a lookalike repository between 10:08 and 11:08. They published 1.107.0 to PyPI at about 11:39 and pushed Docker images at 11:46. Upwind researcher Lidor Machluf reported it at 12:06, and PyPI quarantined the release at 13:23 (The Hacker News). The Docker Hub tags pointed at the malicious image until 14:20 (BleepingComputer). The exfiltration domain was registered on October 13, a week before the attack (The Hacker News). Wiz saw stolen AWS credentials used to query Amazon Bedrock models about **4.5 hours after exfiltration** (SecurityWeek). This is LiteLLM's second compromise of 2026, after the March TeamPCP/Trivy incident (Help Net Security).

Root cause, and a discrepancy. The Hacker News (Oct 24) reports that a LiteLLM master key was committed to a public repository on May 27. The commit was force-pushed over, but the key was never revoked. That key led to an unscoped GitHub CLI OAuth token tied to an organization owner. LiteLLM had also relaxed release controls so "any two engineers" could approve a release, and the attackers used two stolen tokens (The Hacker News). SecurityWeek's earlier account says the tokens came from an internet-exposed internal deployment, where LiteLLM suspects an AI coding assistant embedded them in a logged response body (SecurityWeek). The two accounts appear to reflect different stages of LiteLLM's post-mortem. The later THN version is the more complete one.

Railway, a separate incident. Cloud platform Railway (about 4.8M users, 16M+ hosted services) warned that an unauthorized party may have accessed encrypted customer environment variables. For a subset of users, the attacker obtained decrypted contents, including API keys and database and registry credentials. Railway paused all deployments and says the incident is part of a series of attacks on platform companies aimed at AI inference providers (BleepingComputer). **No outlet links Railway to LiteLLM.**

Why it matters / client impact. AI gateways concentrate the most valuable keys a company has: OpenAI, Anthropic and Bedrock keys alongside cloud credentials. Attackers now cash these in quickly through "LLMjacking", running AI workloads on the victim's bill. Any client with a data-science or AI pilot that ran `pip install litellm` or pulled `:latest` on October 21 should be treated as compromised.

How the attack works. The malicious package dropped startup hooks (`cook.pth`, `cook.py`, `sitecustomize.py`) into site-packages, so the stealer runs **every time Python starts**, including when an admin runs pip to investigate. It collects cloud credentials (AWS, Azure, GCP, Kubernetes, Docker), LLM keys, Git, npm, PyPI and CI tokens, SSH and GPG keys, `.env` files, browser data and shell history. It encrypts the haul and sends it to litellm[.]dev. A loader then polls GitHub hourly for new commands (BleepingComputer).

EXPLAINER

Deleting a leaked secret is not the same as revoking it

LiteLLM's root cause, according to the later account, was a key that had been removed from a public repository but never revoked. Once a secret has been public, assume it was copied. The only real fix is to revoke it at the provider and issue a new one.

What to do now

1. Search lockfiles, container images, CI caches and developer machines for litellm 1.107.0 and the affected Docker tags. Inspect *without launching Python*, look for the three startup files, and check egress logs for litellm[.]dev.

ANALYST ASSESSMENT

P0 for anyone who installed it or

EXPLOITED

Yes (stolen keys abused within hours)

CONFIDENCE

High on impact; medium on root cause

WHO SHOULD ACT

Data-science teams, AI pilots, CI pipelines and containers that pull :latest

AI gateways hold some of the most valuable keys a company has, and stolen ones were used within hours. Pin dependency versions, verify image digests, and keep AI-provider keys in a secrets manager with spending alerts.

KEY FACTS

PACKAGE / ACTOR

litellm==1.107.0 on PyPI. Docker Hub tags v1.107.0, latest, main-latest and main-stable, plus GHCR images. No group named. Attacker IP 69.48.229[.]143.

AFFECTED

Anyone who installed or pulled 1.107.0 on Oct 21. Estimated **49,000+ downloads** (SecurityWeek). The enterprise edition and private repository were not affected.

FIXED / STATUS

Malicious version quarantined. Last clean release: **1.106.2**.

DATE

Compromise Oct 21. Root-cause details published Oct 24.

2. Isolate and rebuild affected hosts and images on 1.106.2. Removing only the LiteLLM directory can leave the startup hooks behind (The Hacker News).
3. Rotate every reachable secret: cloud keys, LLM keys, GitHub, npm and PyPI tokens, SSH keys. Watch Bedrock and LLM billing for abuse. Railway customers should rotate environment variables and registry credentials and review service logs.

VERIFICATION CHECKLIST

Tick each item before closing the ticket

- [ ] Search lockfiles, requirement files, CI caches and container images for litellm 1.107.0 and the affected Docker tags.
- [ ] Rebuild affected hosts and images from clean sources on 1.106.2; do not just uninstall the package.
- [ ] Rotate cloud, AI-provider, GitHub, npm, PyPI and SSH credentials that the environment could reach.
- [ ] Turn on billing and usage alerts for AI providers and Amazon Bedrock.
- [ ] Pin versions and image digests in production; avoid pulling :latest.
- [ ] Railway customers: rotate environment variables and registry credentials and review service logs.

Sources · BleepingComputer, Oct 22 · SecurityWeek, Oct 23 · The Hacker News, Oct 24 · Help Net Security, Oct 22 · BleepingComputer (Railway), Oct 22

FEATURE 04 · COLLABORATION SERVERS · KEV

Nextcloud's built-in Collabora server is being backdoored, and KEV gives agencies three days

CISA added a command-injection flaw in Collabora's built-in CODE server to KEV. Attackers have exploited it since September to plant web shells on Nextcloud servers.

What happened. This is not a new disclosure, but the KEV listing makes it urgent. Exploitation was first reported on October 14, and the earliest logged attack dates to September 10 (The Hacker News, Oct 14). CISA added the flaw on October 23 alongside the Arista bug (The Hacker News, Oct 24). There is a catch for older installs: Nextcloud 29 and earlier are end-of-life and **cannot install the fix**. Version 26.4.402 also reportedly failed on some setups (The Hacker News, Oct 14).

Why it matters / client impact. Self-hosted Nextcloud is popular with SMBs, schools and privacy-conscious firms as a file-sharing service. A web shell on that server exposes every file and credential the web-server account can read. The attackers specifically looked for credential stores and crypto wallets.

How the attack works. The vulnerable script inserted part of the request path into a shell command. Attacker-supplied shell syntax therefore executed, but only while the local CODE server was down, for example during a restart or crash. Observed attacks wrote a one-line PHP web shell (`shx*.php`), sent system information to webhook.site, then searched for secrets. No root access has been reported. Attacker IPs: 139.28.18[.]122 and 89.38.98[.]50 (The Hacker News, Oct 14).

EXPLAINER

Patching does not remove an existing backdoor

This flaw was exploited for about six weeks before the KEV listing. A server patched today may already have a web shell (a small script that gives the attacker remote control) left from September. Patch, then check for files the attackers left behind.

What to do now

1. Update `richdocumentscode` to 26.4.303 or later, or **remove** the app. Disabling it is not enough. On Nextcloud 29 or older, upgrade Nextcloud first or remove the app.
2. Delete any `shx*.php` files. Search logs back to September 10 for `proxy.php?status` requests containing `$(`, `${IFS}`, `%2e%2e` or `base64 -d | sh`.
3. Treat all data readable by the web-server account as exposed. Rotate the credentials stored there and check for wallet or key theft.

ANALYST ASSESSMENT

P0 · KEV due Oct 26

EXPLOITED
Yes (since Sep 10)

CONFIDENCE
High

WHO SHOULD ACT
Self-hosted Nextcloud with the built-in CODE app: SMBs, schools, privacy-focused firms

Nextcloud 29 and older cannot install the fix, so end-of-life installs need an upgrade or app removal now. Treat anything the web-server account could read as potentially exposed and rotate stored credentials.

KEY FACTS

CVE / ACTOR
CVE-2026-110447, CVSS 9.2, OS command injection via `proxy.php`. No actor named.

AFFECTED
Nextcloud servers running the Collabora Online built-in CODE server app (`richdocumentscode`). Standalone Collabora Online servers are **not** affected.

FIXED IN / STATUS
26.4.303 (September 2026), which requires Nextcloud 30 or later. Exploited. In KEV since Oct 23, federal due date **Oct 26, 2026**.

DATE
First logged attack Sep 10. Exploitation reported Oct 14. KEV addition Oct 23.

VERIFICATION CHECKLIST

Tick each item before closing the ticket

- [ ] Update richdocumentscode to 26.4.303 or later, or remove the app (disabling is not enough).
- [ ] Upgrade Nextcloud 29 or older before installing the fix.
- [ ] Search the web root for unexpected shx*.php files and remove them after taking a copy for investigation.
- [ ] Check web logs since Sep 10 for proxy.php requests from the published IPs.
- [ ] Rotate credentials stored on the server and review it for unexpected outbound connections.

Sources · The Hacker News, Oct 24 · The Hacker News, Oct 14

FEATURE 05 · RANSOMWARE · AI-ASSISTED OPERATIONS

"Phantom" shows ransomware affiliates letting AI agents run the kill chain

An exposed server revealed INC and The Gentlemen affiliates using an AI-agent framework against about 62 organizations. The AI executed the steps; a human approved them.

What happened. SOCRadar found an exposed directory at 45.227.254[.]43 holding the operator's "Autonomous Red Team Kill Chain Framework". It runs through the OpenCode agent with commercial models via OpenRouter and has **18 skill modules**, including EDR fingerprinting. It also used "WormGPT-MCP" (**32 MCP servers**), Ghost C2, RevSocks and a bring-your-own-vulnerable-driver (BYOVD) EDR killer. The logs hold **8,257 messages across 133 sessions** (The Hacker News). SOCRadar did not name the specific models used, according to THN.

Why it matters / client impact. Every entry point here is ordinary SMB and hosting tooling: remote support, endpoint management, web hosting panels and VPNs. The biggest force multiplier was not the AI. It was **credential reuse**: one reused password unlocked VPN, servers, RMM, backup and security platforms (The Hacker News). AI lowers the skill needed to run a full intrusion, so expect more affiliates to work at this speed.

How the attack works. Affiliates exploit an internet-facing management product or reuse a password to get in. The agent framework then enumerates the network, fingerprints EDR, and kills it with a vulnerable signed driver (Zemana). It moves on to the backup and security consoles before deploying ransomware. When a model refused a task, the operator reframed the prompt as an authorized red-team engagement (The Hacker News).

EXPLAINER

The real force multiplier was password reuse

SOCRadar's findings show AI speeding up attacks, but the weakness that let one foothold become a full compromise was a single password reused across VPN, servers, remote-management, backup and security consoles. Unique credentials and phishing-resistant MFA on each admin console break that chain whatever tools the attacker uses.

What to do now

1. Patch and restrict exposure of FortiClient EMS, SimpleHelp and cPanel/WHM. None of them should have an open management interface on the internet.
2. Eliminate shared or reused passwords across VPN, RMM, backup and security consoles, and put phishing-resistant MFA on each. Keep backup admin credentials separate from domain credentials.
3. Turn on EDR tamper protection and vulnerable-driver blocking, and alert on new kernel driver loads on servers.

ANALYST ASSESSMENT

P1 · this week

EXPLOITED
Yes (about 62 organizations targeted)

CONFIDENCE
Medium (single research source via press)

WHO SHOULD ACT
Organizations exposing remote-support, endpoint-management, hosting-panel or VPN interfaces

Every entry point named is ordinary SMB tooling with patches available. Close internet exposure of management consoles, separate backup credentials from domain credentials, and make sure EDR cannot be switched off without an alert.

KEY FACTS

ACTOR
Russian-speaking operator working with INC and The Gentlemen ransomware affiliates (SOCRadar research).

AFFECTED
About 62 organizations targeted, 4 later listed on leak sites. More than half in the US. Mainly industrial, manufacturing and energy, then healthcare and pharma.

ENTRY POINTS
FortiClient EMS (CVE-2026-35616), SimpleHelp (CVE-2024-57727), cPanel/WHM (CVE-2026-41940), VPNs, and reused passwords.

DATE
Reported Oct 23.

VERIFICATION CHECKLIST

Tick each item before closing the ticket

- [ ] Patch FortiClient EMS, SimpleHelp and cPanel/WHM, and remove their admin interfaces from the internet.
- [ ] Find and eliminate shared or reused admin passwords across VPN, RMM, backup and security consoles.
- [ ] Enforce phishing-resistant MFA on every administrative console.
- [ ] Keep backup administrator accounts separate from Active Directory and test offline restores.
- [ ] Enable EDR tamper protection and Microsoft's vulnerable driver blocklist; alert on new driver installs on servers.

Sources · The Hacker News, Oct 23

FEATURE 06 · CLOUD ACCOUNT TAKEOVER · BEC

Session-token "spraying" compromises Microsoft 365 at 79 organizations

Huntress tracked a campaign that replayed stolen session tokens within minutes of login. It added rogue devices and swapped victims' MFA methods to set up payroll fraud.

What happened. Huntress counted **497 successful sign-ins** and **more than 200,000 failed attempts**. At least **48 rogue devices** were registered in victims' Entra directories, and at least **19 organizations** had their MFA methods wiped and replaced. One attacker IP was blocked 1,052 times, every time by one of four organizations that had Conditional Access policies (Help Net Security).

Why it matters / client impact. The end goal was business email compromise. Attackers set inbox rules to hide payroll and bank-detail emails, and in one case sent a payroll email to an external accountant. Password resets do not fix this, because registered devices and attacker-added authenticators **survive a reset**. Blocking IPs had limited value, while Conditional Access stopped the attacks outright (Help Net Security).

How the attack works. Attackers obtained session tokens through adversary-in-the-middle phishing, device-code phishing and social engineering, then replayed them in bulk within minutes of the victim's login. Using the hijacked session, they registered a device and planted their own MFA method, which gave them persistence independent of the user's password.

EXPLAINER

Why a password reset was not enough

A session token proves you already signed in, so an attacker who steals it skips the password and MFA. Here attackers also registered their own devices and MFA methods. Those survive a password reset, so recovery means removing them and revoking sessions, not just changing the password.

What to do now

1. Enforce Conditional Access: require MFA and a compliant or registered device, block risky sign-ins, and restrict device-code flow to users who need it.
2. Audit every user for unknown registered devices and authentication methods, remove them, and revoke sessions. Do not rely on password resets alone.
3. Review mailbox inbox rules, especially rules that move or delete mail mentioning payroll, banking or invoices. Require out-of-band verification for payroll and bank-detail changes.

ANALYST ASSESSMENT

P1 · this week

EXPLOITED
Yes (79 organizations)

CONFIDENCE
High (Huntress research via press)

WHO SHOULD ACT
Any Microsoft 365 tenant, especially finance and payroll users

Conditional Access was the control that stopped this campaign outright. Tenants without it should treat this as the reason to deploy it. Pair it with a payroll-change verification process done by phone, since the end goal was payment fraud.

KEY FACTS

ACTOR
Not named (Huntress research).

AFFECTED
Microsoft 365 / Entra ID tenants. **79 organizations** compromised.

STATUS
Campaign ran July 14 to Sept 22, 2026. Not a software vulnerability, so there is nothing to patch.

DATE
Reported Oct 23.

VERIFICATION CHECKLIST

Tick each item before closing the ticket

- [ ] Require MFA plus a compliant or registered device through Conditional Access; block risky sign-ins.
- [ ] Limit device-code sign-in to the users who genuinely need it (MAG212 recommendation).
- [ ] Review every user's registered devices and MFA methods; remove unknown ones and revoke sessions.
- [ ] Audit inbox rules that move or delete mail about payroll, banking or invoices.
- [ ] Verify any payroll or bank-detail change through a known phone number before acting.

KQL · Microsoft Sentinel (verify table and field names in your workspace)

```
// New devices and MFA methods registered in the last 90 days
AuditLogs
| where TimeGenerated > ago(90d)
| where OperationName in ("Register device", "User registered security info",
                          "User changed default security info")
| project TimeGenerated, OperationName, InitiatedBy, TargetResources
| order by TimeGenerated desc
```

Sources · Help Net Security, Oct 23

FEATURE 07 · EDGE CREDENTIALS · INITIAL ACCESS

k0t harvests 70,000+ credentials from firewalls and VPNs, with links to Qilin and The Gentlemen

Group-IB and NyxLab tie a mass edge-device campaign against FortiGate and GlobalProtect to two ransomware operations. k0t may be supplying them with initial access.

What happened. Researchers identified **more than 70,000 exposed credentials** tied to network devices and LDAP dumps (The Hacker News). Many victims later appeared on Qilin and The Gentlemen leak sites, which points to k0t acting as an initial access broker or affiliate (The Hacker News).

Why it matters / client impact. This was opportunistic mass scanning, not targeted research. Any exposed, unpatched or weakly protected appliance was in scope. If a client's firewall was unpatched at any point between February and June, its VPN and LDAP credentials should be presumed stolen, even if the device is patched now.

How the attack works. The actor found devices with the FOFA and ZoomEye search engines, then exploited the flaws or brute-forced logins. They connected over OpenConnect VPN as a legitimate user and enumerated Active Directory and SMB. For persistence they used RMM tools (MeshCentral, Atera, AnyDesk) and the AdaptixC2 framework, with Mimikatz and Impacket for credentials. Data left via Rclone to Wasabi S3 and MEGA (The Hacker News).

EXPLAINER

Patched today does not mean safe since February

If a firewall or VPN was unpatched at any point during the campaign, its stored VPN and directory credentials may already have been copied. Patching stops new break-ins, but stolen credentials keep working until they are changed.

What to do now

1. Confirm CVE-2025-64446 and CVE-2026-0257 are patched, and take VPN and management interfaces off the open internet or restrict them.
2. Reset VPN, local-admin and LDAP bind credentials on any device that was exposed and unpatched this year. Enforce unique credentials and monitor for brute force.
3. Hunt for unsanctioned MeshCentral, Atera or AnyDesk installs and for Rclone traffic to Wasabi or MEGA.

ANALYST ASSESSMENT

P1 · this week

EXPLOITED
Yes (70,000+ credentials)

CONFIDENCE
Medium (Group-IB/NyxLab via press)

WHO SHOULD ACT
Organizations running FortiGate/FortiWeb or Palo Alto GlobalProtect at the edge

This was mass, opportunistic scanning. Any exposed appliance was in scope. The links to Qilin and The Gentlemen suggest stolen access is being resold to ransomware groups, so credential resets on affected devices should not wait.

KEY FACTS

ACTOR
k0t (also CryptoFeel, Member19, QazWsxShop). Assessed as linked to Qilin and The Gentlemen ransomware. The exact relationship is unclear.

AFFECTED
Fortinet FortiGate/FortiWeb and Palo Alto GlobalProtect/PAN-OS devices worldwide. Victims are about 38% APAC, 31.9% LATAM, 15.8% Europe and 12.7% MENA.

CVES
CVE-2025-64446 (FortiWeb path traversal) and CVE-2026-0257 (PAN-OS/Prisma Access authentication bypass), plus brute force. Patches exist for both.

DATE
Most compromises Feb to June 2026. Reported Oct 23.

VERIFICATION CHECKLIST

Tick each item before closing the ticket

- [ ] Confirm patches for CVE-2025-64446 and CVE-2026-0257 are installed.
- [ ] Restrict VPN and management interfaces to known addresses where possible.
- [ ] Reset VPN, local-admin and LDAP bind credentials on any device exposed and unpatched this year.
- [ ] Look for remote-management tools (MeshCentral, Atera, AnyDesk) that IT did not install.
- [ ] Alert on large uploads to cloud storage services such as Wasabi or MEGA.

Sources · The Hacker News, Oct 23

FEATURE 08 · DEVOPS · DEPENDENCIES

GitLab ships a max-severity RCE fix, and a Brotli bug opens Node.js, Python and service meshes

Neither flaw is exploited yet. GitLab's is unauthenticated, and Brotli's reaches deep into the software supply chain. Patch both before weaponization.

What happened. GitLab patched an unauthenticated RCE that an attacker could trigger by getting the server to compile a crafted regular expression. The same researcher, "joaxcar", reported a CVSS 10 GitLab flaw (CVE-2026-97229) about ten days earlier (BleepingComputer). On October 23, Ruby fixed three regex-engine bugs behind four critical GitLab RCEs (The Hacker News). For Brotli, Google rates the flaw **CVSS 4.0 9.2**, Red Hat rates it **8.1**, and Node.js calls it "high", so severity ratings vary by vendor. CISA's October 23 assessment was exploitation "none" and automatable "no" (The Hacker News). Imperva demonstrated RCE on Node.js 24.22.0 on Linux ARM64. It took several attempts, and the exploit has not been released (The Hacker News).

Product	Brotli-fixed versions
Brotli	1.0.10, 1.1.1, 1.2.1
Node.js	22.23.4, 24.22.1, 26.12.1 (20.x and older are end-of-life)
Bun	1.4.4
Envoy	1.40.0, 1.39.3, 1.38.6, 1.37.8, 1.36.12
Istio	1.31.2, 1.30.6, 1.29.9
Python Brotli / brotlicffi	1.2.1 / 1.2.1.0

Source: The Hacker News, Oct 24

Why it matters / client impact. A self-managed GitLab server holds source code, CI secrets and deploy keys, which makes it a launchpad for the kind of supply-chain attack LiteLLM just suffered. Brotli is the week's most widely deployed library flaw. Fixing it means restarting services and **rebuilding container images**, not just running package updates.

How the attack works. GitLab: crafted input makes the server's regex compiler overflow and run attacker code without authentication. Brotli: a malicious compressed response or request body corrupts memory in the decoder. That affects Node.js and Bun `fetch()` decoding `br` responses, Express 5 and body-parser 2 decoding Brotli bodies, Envoy and Istio only when Brotli decompression is enabled, and Python's requests, urllib3, httpx and aiohttp when brotli or brotlicffi is installed (The Hacker News).

EXPLAINER

Why library fixes need rebuilds

Brotli is a compression library compiled into Node.js, Python packages, Envoy and Istio. Updating the operating system's package is not enough: each runtime and every container image that bundles its own copy has to be updated and rebuilt, and running services restarted.

What to do now

1. Upgrade self-managed GitLab to a fixed release now, and keep its admin interface off the open internet where possible.
2. Update Brotli across Node.js, Bun, Python, Envoy and Istio, then restart services and rebuild images. Scan your container registry for the old library versions.
3. If you cannot patch yet, reject `Content-Encoding: br` at the edge, disable Brotli decoding in clients, or uninstall Python brotli/brotlicffi. Removing `br` from Accept-Encoding does **not** help, and Node.js has no workaround (The Hacker News).

ANALYST ASSESSMENT

P1 · patch before exploitation

EXPLOITED
No (not yet)

CONFIDENCE
High

WHO SHOULD ACT
Self-managed GitLab owners; teams running Node.js, Python, Envoy or Istio

Neither flaw is exploited yet, which makes this the cheapest week to fix them. A self-managed GitLab server holds source code and CI secrets, so it is a prime target for the kind of supply-chain attack LiteLLM just suffered.

KEY FACTS

CVE
GitLab CVE-2026-110264. Integer overflow in the regex compiler leading to unauthenticated RCE. Rated "maximum severity", no numeric CVSS published. **Brotli CVE-2026-111497 ("Node2Shell").** Memory corruption in the C decoder.

AFFECTED
GitLab CE/EE (Shadowserver sees **2,700+** internet-exposed CE instances). The Brotli decoder in Node.js, Bun, Python, Envoy and Istio.

FIXED IN
GitLab 19.5.1, 19.4.3, 19.3.6, 19.2.9, 19.1.9, 19.0.10 and 18.11.13. Brotli fixes are in the table below.

STATUS / DATE
GitLab: Oct 22, not exploited, GitLab.com already patched. Brotli: fixes from Oct 21, follow-ups Oct 23-24. No exploitation reported.

BROTLI
1.0.10, 1.1.1, 1.2.1

NODE.JS
22.23.4, 24.22.1, 26.12.1 (20.x and older are end-of-life)

BUN
1.4.4

ENVOY
1.40.0, 1.39.3, 1.38.6, 1.37.8, 1.36.12

ISTIO
1.31.2, 1.30.6, 1.29.9

PYTHON BROTLI / BROTLICFFI
1.2.1 / 1.2.1.0

VERIFICATION CHECKLIST

Tick each item before closing the ticket

- [ ] Upgrade self-managed GitLab to a fixed release and keep its admin pages off the open internet.
- [ ] Update Node.js, Bun, Python brotli/brotlicffi, Envoy and Istio to the fixed versions in the table.
- [ ] Scan your container registry for old Brotli versions and rebuild affected images.
- [ ] Restart services after updating so the fixed library is actually loaded.
- [ ] If patching must wait, reject Brotli-encoded request bodies at the edge where feasible.

Sources · BleepingComputer, Oct 22 · The Hacker News, Oct 23 (Ruby) · The Hacker News, Oct 24 (Brotli) · The Hacker News, Oct 23 (Node2Shell)

BREACH WATCH

Breach roundup

None of these incidents has a confirmed ransomware attribution, named actor or disclosed ransom as of October 24.

Capcom: 11.7 million CAPCOM ID accounts exposed. Capcom detected suspicious activity on October 19, confirmed customer data was compromised on October 20, reported to police and Japan's privacy regulator on October 21, and disclosed publicly on October 22. Exposed data covers **more than 11.7 million email addresses** with account creation dates, last-login timestamps and IPs, and login counts. For 43 accounts, birth date, gender, nickname and country were also exposed. Attackers changed the passwords on 196 accounts to lock owners out. Capcom says other passwords and card data were not exposed (BleepingComputer). SecurityWeek's headline rounds the figure to "11 Million" (SecurityWeek). *Client action:* expect phishing aimed at those addresses, and warn staff who use CAPCOM ID with work email.

Quest Diagnostics: intrusion into Quanum Lab Services Manager. An unauthorized party accessed the system customers use to order lab tests and look up results. Quest believes "some patient data" was accessed but has not said how many people or which data types. It paused that system and several others, moved customers to paper requisitions, and aimed to restore service by October 24. Quest says the incident **did not involve PHP2Shell** (The Record). SecurityWeek reports the disclosure came via a regulatory filing citing "certain patient records" (SecurityWeek; BleepingComputer). *Client action:* healthcare clients that order through Quanum should expect follow-up notices and watch for lab-themed phishing.

City of Columbia, South Carolina. Suspicious activity disrupted MyColaSC (utility payments and non-emergency reports), GIS and animal services. Police, fire, 911 and water kept running. The city would not say whether ransomware was involved (The Record).

Toronto Crime Stoppers / Navigate360 (P3 Global Intel). Toronto Crime Stoppers paused operations after CBC reviewed about 1.5M tips stolen from the vendor. For Toronto, that included 176,859 tips containing 14,973 names. The hacktivist group "Internet Yiff Machine" claimed the March breach, and the vendor's account of the impact shifted between March and June (The Record). *Lesson:* write breach-notification timelines and data-retention limits into vendor contracts.

Ongoing: Ahold Delhaize is still restoring online ordering after an attack that began October 12. The company reports no evidence of data theft (SecurityWeek).

SINCE SPECIAL EDITION 01

Follow-ups

PHP2Shell (CVE-2026-109618/109619), plus a new unpatched sibling, CVE-2026-115560. Attack attempts have grown sharply, but no successful compromise had been confirmed as of October 23. Amazon saw the first attempts on **October 17, four days before the patches** (BleepingComputer). Cloudflare saw attempts climb from tens of thousands within 45 minutes of disclosure to **millions by October 22** (The Hacker News). Sophos ties some attempts to "Potassium", a Mirai variant, and corrected an earlier claim that an attempt had succeeded. Hunt.io saw 10 rounds against 123 servers with no success. VulnCheck built a private ASLR-bypass exploit, which shows reliable RCE is achievable. Qualys counted about 199,000 instances patched in 12 hours, but about 71% of PHP instances run end-of-life branches, and 32 UK government services went offline while patching. Neither CVE was in KEV as of the morning of October 23 (SecurityWeek). The UK NCSC told organizations to patch or take servers offline (BleepingComputer).

New: On October 23 PHP disclosed CVE-2026-115560, an RCE with no emergency patch and no CVSS. It affects only non-default setups that change the session serializer at runtime with upload progress enabled, and no exploitation is reported. **The October 21 releases do not fix it.** The mitigation, `session.upload_progress.enabled=0`, also hardens against PHP2Shell (The Hacker News; BleepingComputer). *Action:* make sure every server runs 8.5.12, 8.4.27, 8.3.36 or 8.2.35, set the upload-progress mitigation, and restart workers (BleepingComputer).

eMudhra: Chrome removes the root on October 27. Chrome and Mozilla committed to removing the `emSign ECC Root CA - G3` root, and Chrome's removal takes effect **October 27**. eMudhra estimates about **30,000 certificates** need reissuing. On October 21 eMudhra disputed Mozilla's account, saying seven of the eight intermediates were unused. Trust-store fixes remain uneven. As of October 22, Amazon Corretto, Amazon Linux, RHEL/UBI and Bottlerocket had no fixed packages. Lambda Python 3.13 and Java 21 images still trusted the chain, Debian 11 had no fix, and Apple's block does not cover the built-in macOS command-line tools. Fixed runtimes include Go 1.26.10/1.27.3, Node.js 22.23.4/24.22.1/26.12.1 and certifi 2026.10.19 (Konfirmity advisory, updated Oct 23). Note that Konfirmity is a compliance vendor's page, not a primary source, and some of its timestamps are inconsistent. OpenJDK published CVE-2026-91403, and Oracle issued a matching Java alert (SecurityWeek). *Action:* replace any eMudhra-chained certificates before Monday.

Plex2Shell. CISA added four Plex Media Server CVEs on October 20. Federal due dates for CVE-2026-96654 and CVE-2026-96656 were October 23, and are November 3 for CVE-2026-96652 and CVE-2026-96655. Chained, they give unauthenticated RCE that attackers are using for cryptominers and proxy malware. The fix is PMS 1.43.3.10861 (The Hacker News). We found no new developments after October 21.

PAN-OS GlobalProtect CVE-2026-0257. We found no new patch, KEV change or vendor update dated October 21-24. However, the flaw is central to the k0t campaign (Briefing 7), and Unit 42's October 19-20 update describes Qilin and INC exploitation since early June (Start With Identity; Gridinsoft). *Action:* patch, or disable authentication override / use a dedicated override-cookie certificate. Terminate active sessions after patching, and hunt for VPN sessions from hosts named "kali" (The Arabian Post).

EVERYTHING ELSE THAT MATTERED

Quick hits

- **Oracle October CPU:** 1,295 fixes covering 1,209 unique CVEs. Fusion Middleware has 206 fixes, 134 of them remotely exploitable without authentication. E-Business Suite has 194 and Java SE 55. None are flagged as exploited. Next CPU is January 19, 2027 (SecurityWeek).
- **Chrome 156:** 353 security fixes (6 critical, 61 high), with no reported in-the-wild exploitation. Make sure the browser has actually relaunched (SecurityWeek).
- **Cisco:** 34 flaws fixed, including a critical Expressway root RCE (CVE-2026-76466, fixed in X15.5.20), three critical Unified Communications bugs, and Snort 3 criticals (fixed in 3.12.4.0) that affect ASA, FTD, FMC and Meraki MX. None exploited (SecurityWeek).
- **HPE Aruba AOS-8/AOS-10:** 72 flaws, 8 of them max-severity unauthenticated RCE or auth-bypass bugs (CVE-2026-79844 to -79851) in controllers, gateways and APs. No public PoC or exploitation. Put management on a dedicated VLAN (BleepingComputer).
- **Lazarus "FlashHole":** Kaspersky reports South Korean firms targeted through GitHub-hosted lures and a chain that abuses INISAFE CrossWeb EX. The LPEClient loader now hides payloads with an ML-KEM-1024 post-quantum key exchange (BleepingComputer).
- **Sandworm and the Yggdrasil mesh:** ESET documents UAC-0099 handing access to Sandworm. Sandworm used the MythicRatatoskr backdoor, with command-and-control over the Yggdrasil mesh network, before wiping a Ukrainian logistics firm via Group Policy in September 2025. The attack is from 2025; the disclosure is new (The Hacker News).
- **BlackX UEFI bootkit:** Kaspersky reports a Secure Boot bypass (via CVE-2022-21894) against Central Asian diplomats, with APT41 attribution only suspected. Windows Update alone does not stop it: admins must manually apply Microsoft's DBX revocations, including "Windows Production PCA 2011" (The Hacker News).
- **Talos callback phishing:** Attackers abused more than 25 legitimate platforms to send lures that pass SPF, DKIM and DMARC. One example: about 1,400 fake Dell/Staples order confirmations, such as a $736.09 "Microsoft Project" charge with a "Helpline" phone number. Train users to distrust unexpected receipts that list a support number (Help Net Security).
- **Symantec LLMjacking agent:** A honeypot caught an attack agent that hijacks exposed OpenAI-compatible endpoints, AI gateways and MCP servers. It uses the victim's own model to plan attacks and deploy XMRig, and can run up to 100,000 steps without human approval. The honeypot logged more than 127,000 requests (Help Net Security). Never expose model servers without authentication.
- **Interpol Serengeti III:** 1,462 arrests, more than $105M recovered and about 66,000 victims with losses above $385M. Sources differ slightly on scope: BleepingComputer cites 10,909 networks dismantled across 21 African countries, while The Record says 20 African countries plus the US (BleepingComputer; The Record).
- **UK undersea-cable law:** A white paper proposes up to 14 years in prison for deliberate damage to undersea cables and pipelines, replacing the 1885 Act's £1,000 maximum fine. Consultation closes January 14, 2027 (The Record).
- **FCC covered-components ban:** First enforcement under the rules that took effect October 13. The FCC blocked a wireless video transmitter and a dashcam containing Huawei HiSilicon chips (The Record).
- **Windows Server 2016 end of support:** January 12, 2027. The January 2027 update will be the last, and Microsoft recommends Windows Server 2025. Start migration planning now (BleepingComputer).
- **Teams AI phishing detection:** Context-aware impersonation detection for external first-contact chats. Targeted release runs early to mid-November, with GA expected in December. No admin action is needed (BleepingComputer).
- **More patches:** Pillow CVE-2026-112900, critical, fixed in 13.0.0 (The Hacker News). Atlassian fixed 75 third-party flaws (SecurityWeek). Splunk and SolarWinds also released updates (SecurityWeek).
- **Incidents in brief:** 333 Tor exit relays (about 10% of exits) were hijacked to intercept SSH, traced to a reused password. PingCAP lost TiDB Cloud customer and billing data to stolen service-account credentials. Bullhorn disabled system access after an incident, and API users should rotate API passwords (SecurityWeek roundup).
- **ShinyHunters allegations:** KrebsOnSecurity sources, relayed by SecurityWeek, say investigators believe a US consulting-firm executive helped run the "Shinywire" extortion site. These are allegations, not charges (SecurityWeek roundup).
- **Also noted:** Iran's "Prince of Persia" APT hosted C2 in the US for the first time since 2016 (SecurityWeek). A critical x402 SDK flaw with no default spending limit could let malicious APIs drain AI agents' balances; August releases added a limit (The Hacker News).

ANALYST TAKEAWAY

Treat every non-human identity like an admin account

This week's incidents had less to do with clever exploits than with misplaced trust. A vendor app with standing consent, a session token that outlived its login, a VPN password reused on the backup console, and a key that was hidden from view but never revoked: each gave attackers legitimate-looking access that perimeter tools and user MFA cannot see. AI agents are now driving that access at machine speed. Storm-0940 and the Phantom operators reached privileged positions in hours, so control reviews that happen weekly will be too slow.

For MSPs, the practical shift is to treat **non-human identities** (OAuth apps, service principals, CI tokens and AI-gateway keys) as tightly as admin accounts: inventory them, scope them, rotate them and alert on them. The two KEV deadlines (October 26) and Chrome's eMudhra cutoff (October 27) make Monday the next deadline. Patch VeloCloud and Collabora this weekend.

REFERENCE

Glossary

CISA KEV
The US Known Exploited Vulnerabilities catalog. A listing means confirmed exploitation, and federal agencies must patch by a set date.

CVSS
A 0-10 severity score. Different scorers can rate the same flaw differently.

Conditional Access
Microsoft Entra policies that decide who can sign in, from where and from which device.

Web shell
A small script left on a web server that lets an attacker run commands remotely.

LLMjacking
Using stolen AI-provider keys to run AI workloads at the victim's expense.

SD-WAN
Software-defined wide-area networking that links branch offices under central control.

Zero-day
A flaw exploited before the vendor has released a fix.

Service principal
The identity an application uses inside your Microsoft Entra tenant.

Session token
Proof that a user already signed in. If stolen, it can bypass password and MFA until revoked.

Supply-chain attack
Compromising software or a vendor that many organizations trust, to reach them all at once.

Initial access broker
A criminal who breaks in and sells that access, often to ransomware groups.

Phishing-resistant MFA
MFA such as passkeys or security keys that cannot be relayed by a fake login page.

How we sourced this issue: Status is as of the morning of October 24, 2026, from trade press (BleepingComputer, The Hacker News, SecurityWeek, The Record, Help Net Security, CyberScoop) and Microsoft's primary report. KEV details come from trade-press coverage, not the CISA catalog directly. Most vendor research (Huntress, SOCRadar, Group-IB, Kaspersky, ESET, Talos, Symantec) is cited through news coverage. No exploit code is published in this magazine.

Analyst assessments, priorities, explainers and verification checklists are MAG212 CYBER analysis. Queries are starting points for authorized defensive use and should be validated in your own environment. No exploit code is published in this magazine.

MAG212CYBER

CYBERSECURITY WORLD REPORT

Independent cybersecurity intelligence for IT administrators, MSPs and security teams. Written by the MAG212 CYBER analyst desk.

PUBLISHED BY

MAG212 CYBER

CYBER.MAG212.COM