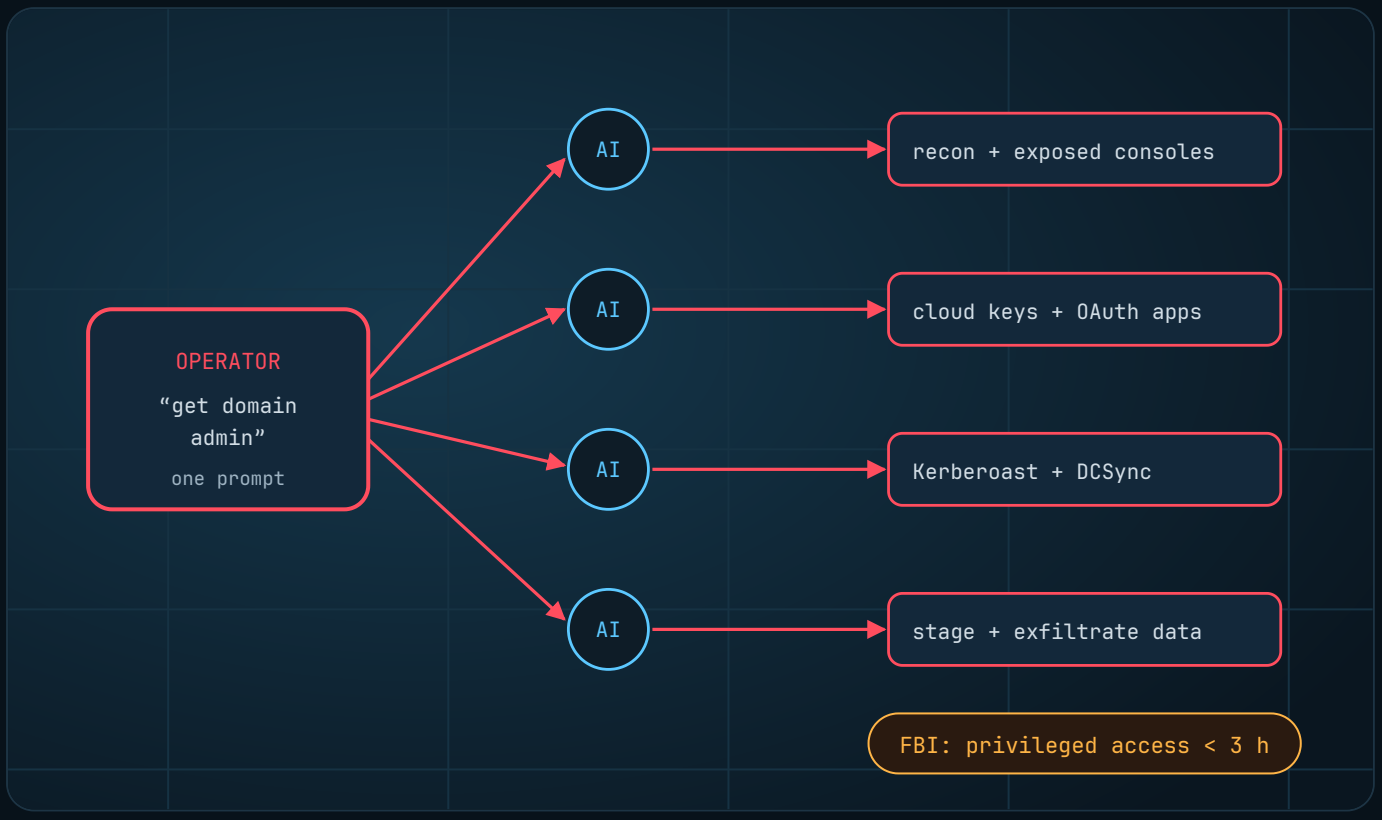
MAG212
CYBER

DAILY

Sunday, 15 November 2026

DAILY CYBER INTELLIGENCE · CUTOFF 7:45 AM PST




THE SPEED GAP / NOV 13–15

# Machine speed, human gaps

The FBI says AI agents now take some intruders from first scan to privileged access in under three hours. The weekend's other stories show where they find the gaps: an RMM fix that is still not out, suppliers that hold other companies' data, and outage emails nobody verifies.

03 FBI FLASH on ARTEX: AI agents that run the whole intrusion

06 N-central: miners on patched servers, HF3 still not out

07 Expedia confirms a supplier breach as n0n claims 1B records

08 Maileva: 22 million letters' recipients exposed

09 Outage-themed phishing after the Microsoft 365 disruption

10 KEV deadline tomorrow: UnrealIRCd and Check Point

# What matters today

This edition covers 13–15 November developments that change what defenders should do. Stories from our 14 November Weekly return only where something new happened.

### AI-run intrusions are now an FBI warning

FLASH-20261113-001 says threat actors use ARTEX, an open-source multi-agent pentest framework, against ten sectors, and that some intrusions reached privileged access in under three hours. It brings no new exploits: speed and parallelism are the threat. [1][2]

### N-central: the patch race is not over

Hotfix 3 missed its 14 November target. Huntback counted 21 servers mining Monero via 4EverAble2Shell on 12 November, 19 of them on Hotfix 2. Keep the Legacy API off. [7][8]

### Suppliers keep exposing other people's customers

Expedia confirmed attackers reached it through a third-party data platform provider; n0n claims nearly a billion records (unverified). Maileva, a La Poste print-and-mail service, confirmed 22 million recipients' names and addresses were stolen. [13][18]

### Deadlines and lures

CISA's KEV due date for the UnrealIRCd and Check Point flaws is tomorrow, 16 November. Microsoft warned of phishing that impersonates its own outage notices. [26][21]

## THE THREE PRIORITY ACTIONS

1. **Hunt ARTEX-style activity in identity and cloud logs today.**
   Review Entra audit records for new credentials on applications and service principals, unfamiliar service-principal sign-ins, and outbound calls to LLM APIs from servers. Lock Kudu/SCM and cloud admin endpoints to admin networks. [1]
2. **Close the two KEV items due tomorrow.**
   Upgrade UnrealIRCd to 6.2.7.1 (or apply the hotfix or spamfilter workaround) and finish the Check Point Jumbo Hotfix and log hunt from our Weekly. [25][26]
3. **Keep N-central's Legacy API off and verify.**
   Confirm it is disabled and unreachable from the internet, apply HF2, upgrade Agents and Probes, hunt the miner indicators, and install HF3 when N-able ships it. [7][8][9]

## CONTENTS


03 ARTEX: the FBI's warning on AI-run intrusions
04 ARTEX: what to hunt, with commands
05 ARTEX response + Opinion: is AI making us worse?
06 N-central update: HF3 delayed, miners on HF2
07 Expedia and the third-party data platform
08 Maileva: 22 million recipients exposed
09 Outage-themed phishing after Exchange Online
10 Due tomorrow: UnrealIRCd and Check Point
11 Briefs: FAA, PostgreSQL, Apple, Turla, iOS
12 Defender's Action Desk
13 Awareness: outage emails and 'press these keys' pages
14 Sources

FBI FLASH-20261113-001 | ISSUED 13 NOV | TEN SECTORS | NO NEW EXPLOITS

# ARTEX: the FBI warns AI agents now run whole intrusions

An open-source autonomous pentest framework has become a shared attack platform for criminals and state actors. The FBI's alert is the clearest official statement yet that agentic AI is speeding up real intrusions.

> STATUS **Confirmed by the FBI.** The FLASH describes ARTEX use since at least August 2026, publishes historical IOCs, and recommends specific controls. Campaign attribution comes from the vendors it cites (CrowdStrike, Microsoft, Zscaler and others). [1]

## WHAT HAPPENED

On 13 November the FBI published FLASH-20261113-001 on ARTEX, a Go-based, multi-agent framework that turns a plain-language objective into commands run by AI agents. It has no model of its own; operators plug in commercial or open models, and on attacker-run instances Chinese-origin DeepSeek Flash-family models dominate. [1]

The developer published ARTEX on GitHub in July and pulled it on 8 October after reports of abuse, but forks, mirrors and a commercial ARTEX PRO keep it in circulation. Researchers counted more than 1,600 exposed ARTEX consoles in early November. [1][2]

## WHAT IS STILL UNKNOWN

Actor names differ across sources: the FLASH ties the ConnectWise activity to Storm-3358, while Microsoft's 21 October report is filed under Storm-0940. MAG212 reports both and does not merge them. [1][5]

## WHO IS AFFECTED

- Ten named sectors: financial services, government, technology, defence, semiconductors, energy, healthcare, manufacturing, telecoms and legal. [1]
- Microsoft 365 tenants with over-privileged multi-tenant apps or vendor integrations (ConnectWise apps named).
- Anyone whose cloud admin, Kudu/SCM, hypervisor or backup consoles are reachable from the internet.

## WHY IT MATTERS

The FLASH cites a South Korean financial campaign (about 67,000 people's data taken) and a ConnectWise Cloud Backup compromise reaching 2,250+ downstream tenants. One small team can now run several intrusions at once. [1][4]

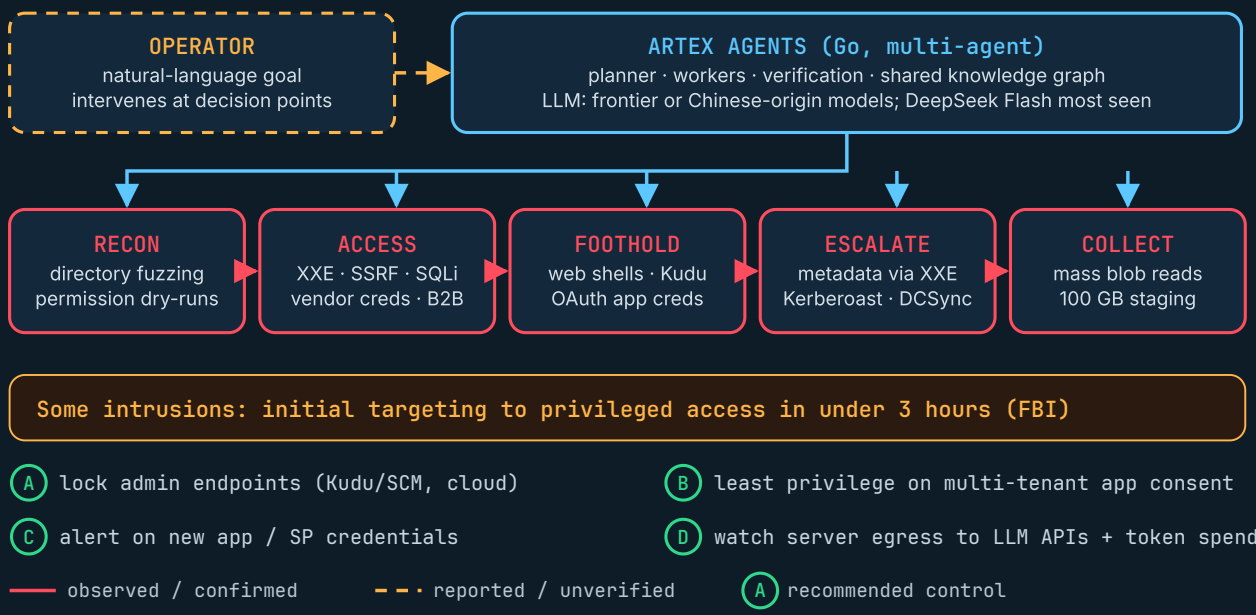



Figure 1 · ARTEX intrusion flow as described in FBI FLASH-20261113-001. Techniques are examples from the FLASH, not one single intrusion. Diagram: MAG212 CYBER.

# Hunting machine-speed intruders

ARTEX chains familiar techniques. The defence is the same hygiene, applied fast: shrink what is exposed, watch identity changes, and treat LLM traffic from servers as a signal.

## WHAT DEFENDERS SHOULD LOOK FOR

- Outbound HTTPS from servers to `api.deepseek.com`, `open.bigmodel.cn`, `api.x.ai` or `api.anthropic.com` where no business use exists. [1]
- Self-hosted LLM gateways (CLIProxyAPI, New API) or unknown resellers in proxy logs; unusual token spend. [1]
- Bursts of sequential lookups against portals, login and search endpoints; 500+ directory-fuzzing requests.
- New .aspx/.jsp/.php files in web roots; Kudu/SCM API calls you did not schedule.
- PowerShell sweeps of domain controllers, LAPS reads, DCSync; 100 GB staging volumes; output sent to webhook.site.

**False positives:** developers and data teams legitimately call LLM APIs; vulnerability scanners also fuzz directories. Correlate with identity and host activity before escalating.

### FBI IOCs · HISTORICAL · VET BEFORE BLOCKING

| | |
|---|---|
| ARTEX host | 38.244.50[.]120 (open directory) |
| WISP C2 | 199.19.95[.]66, .67, .68 (active 6 Nov) |
| ARTEX srv | 91.92.240[.]238, 85.11.168[.]45, 178.16.54[.]129, 31.77.229[.]35 |
| C2 | 85.11.167[.]52 |
| Graph abuse | 144.172.64[.]60, .99, .128 · 187.15.88[.]186, 187.15.90[.]217, 187.15.89[.]186 |
| Reseller | xcai[.]pro (LLM API) |

Cloud IPs get reassigned; the FBI says to vet these against current telemetry before blocking. Full list: [1]; also reported by [3].

PowerShell · Microsoft Graph SDK | Application.Read.All | READ-ONLY | NOT TESTED BY MAG212

ENTRA · SP CREDENTIAL SWEEP

```
Connect-MgGraph -Scopes 'Application.Read.All'
Get-MgServicePrincipal -All |
  Where-Object { $_.PasswordCredentials.Count -gt 0 -or $_.KeyCredentials.Count -gt 0 } |
  Select-Object DisplayName, AppId,
    @{N='Secrets';E={$_.PasswordCredentials.Count}},
    @{N='Certs';E={$_.KeyCredentials.Count}}
```

Lists service principals that hold their own secrets or certificates. Credentials added directly to a service principal do not show in the portal's App registrations blade. Any you cannot explain go to the incident lead. [42]

KQL · Sentinel / Log Analytics | Log reader | READ-ONLY | NOT TESTED BY MAG212

ENTRA · CREDENTIAL ADDS

```
AuditLogs
| where TimeGenerated > ago(45d)
| where Category == "ApplicationManagement"
| where OperationName has_any ("Add service principal credentials",
    "Certificates and secrets management")
| project TimeGenerated, OperationName, InitiatedBy, TargetResources
```

Finds credential additions to apps and service principals. Requires Entra audit logs streamed to Log Analytics; the portal keeps about 30 days. Expect noise from legitimate certificate rotation. [44]

# ARTEX: investigate, contain, verify

## HOW TO INVESTIGATE

- Pull Entra sign-in and audit logs, service-principal sign-ins (successful and failed, reviewed separately), Graph activity and Purview audit. [1]
- Check AWS IAM (AssumeRole, PassRole, instance profiles), Azure ARM dry-runs, blob write/delete timing and Kudu/SCM logs.
- If ConnectWise apps were ever authorised, search tenants by application ID, not display name, from 1 October onward. Sophos advises hunting back to at least September 2026. [1][5][6]

## CONTAIN AND REMEDIATE (IN ORDER)

1. Preserve Entra, Purview, SIEM and Graph evidence, but do not wait for exports before urgent containment.
2. Remove unauthorised credentials, consent grants, owners, role assignments and identities.
3. Rotate exposed keys and secrets; remove web shells and rogue SSH keys.
4. Keep affected ConnectWise app connections disabled until ConnectWise issues restoration steps; do not delete enterprise apps first.
5. Lock Kudu/SCM, cloud admin, hypervisor and backup consoles to admin paths.

Source: FBI recommendations [1].

## VERIFY RECOVERY

- Re-run the credential sweep: only expected secrets remain.
- No credential adds from unexpected actors for 14 days.
- Admin endpoints answer only from admin networks.
- Server LLM egress matches an approved list.

## PREVENTION

Rate-limit high-impact APIs and portal lookups; scan code, images and pipelines for embedded secrets.

OPINION · MAG212 ANALYSIS

### Is AI making us worse at cybersecurity?

SecurityWeek asked it in a 14 November headline. ARTEX answers half: attackers are getting faster. The other half is about defenders.

In ISC2 research, 65% of professionals who use AI say they spend time deciding when to trust its output; 62% say it has not reduced the need for foundational skills (26% say it has); 32% say it made stress worse. [35][36]

Our view: AI makes teams worse when it replaces understanding and better when it removes toil. ARTEX chains ordinary techniques at machine speed; an analyst who can read an Entra audit log will catch a rogue service-principal secret faster than a tool they cannot explain.

- Use AI to triage; keep human sign-off on containment.
- Train juniors on raw logs, not only AI summaries; measure how often AI output is wrong.

SecurityWeek's column text could not be retrieved for this issue; we cite only its headline. Survey figures: ISC2, as reported by ITPro and SecureWorld [35] [36].

FOLLOW-UP TO WEEKLY P.06 | HF3 NOT YET RELEASED | EXPLOITED IN THE WILD

# N-central: miners on patched servers, and the next fix slips again

**CORRECTION What changed.** Our Weekly called exploitation *disputed* and Huntback's figures unverified; we have now read Huntback's report. Exploitation of internet-exposed Legacy API servers is evidenced, and N-able has published IoCs for servers with the Legacy API on. [8][9]

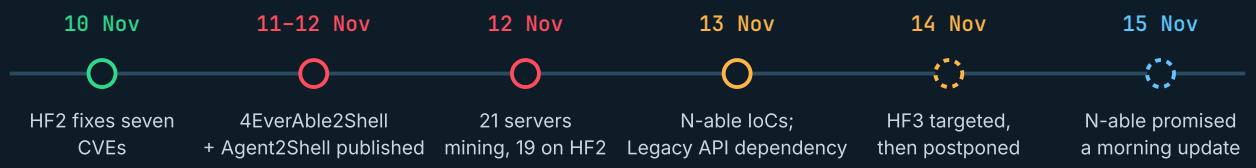



Figure 2 · N-central timeline, 10–15 Nov. Dashed = scheduled or promised. Sources: [7] [8] [12].

## WHAT HAPPENED

Huntback's decoys received the 4EverAble2Shell exploit on 12 November. Public mining-pool statistics for the attacker's wallet showed 21 rigs named `nc_<victim-ip>`; 19 reported build 2026.4.1.26 (HF2) and one 2025.3.1.9. Huntback treats 21 as a lower bound. [8]

The chain: a crafted SOAP fault to `/dms/services/ServerUI`, code execution as `nable`, root via `nable_wrapper.pl` and a flag-file symlink, then XMRig. [8]

Late on 14 November (US time) N-able said HF3 'will not be available today' and promised an update on the morning of 15 November ET. No HF3 had shipped at our cutoff. [7][41]

**Unknown:** whether the 19 HF2 servers were hit before or after patching; victim addresses are not published.

## WHO IS AFFECTED

- On-premises N-central with the **Legacy API enabled and internet-facing**, including HF2. [9][10]
- Agent2Shell needs the Legacy API too. Hosted N-central is not affected by the server chain. [9]

### CONTAIN NOW

1. Confirm the Legacy API is disabled (N-able KB); block it at the edge.
2. Keep HF2; upgrade every Agent and Probe.
3. Block 157.20.241[.]16 and pools 139.99.69[.]109, 79.76.51[.]156.
4. Any hit: isolate, rebuild, rotate RMM and customer credentials.
5. Install HF3 when released; re-scan.

Sources: [8], [9].

### VERIFY RECOVERY

- Legacy API path unreachable from outside; no ServerUI SOAP faults with exceptionName in logs.
- No pool traffic; CPU at baseline. N-able confirms patch state (version strings may not change).

Linux shell on the N-central server | root | READ-ONLY | NOT TESTED BY MAG212

N-CENTRAL · MINER CHECK

```
ls -la /var/opt/n-central/flag/flag_default.json /tmp/.xvd /var/tmp/.xvd 2>/dev/null
ls -la /var/opt/n-central/tmp/ | grep serverchild
systemctl status xvd.service --no-pager; grep -rs 'xvkeep' /etc/cron.d /var/spool/cron
ps aux | grep -E -- '--rig-id nc_|--donate-level 1' | grep -v grep
```

Looks for Huntback's artefacts: symlinked flag file, `.xvd` miner, serverchild scripts, systemd/cron persistence, XMRig flags. Any hit = full compromise. Single-vendor indicators. [8]

EXPEDIA: CONFIRMED | SCALE: UNVERIFIED | STATEMENT 12 NOV

# Expedia confirms a supplier breach as n0n claims nearly a billion records

STATUS **Confirmed:** attackers used a compromised third-party data platform provider to reach Expedia and other companies. **Not confirmed:** data taken, the provider's name, n0n's figures. [13][14]

## WHAT HAPPENED

n0n listed Expedia Group in early November, claiming more than a terabyte of data. On Thursday 12 November Expedia acknowledged a supply-chain attack: a compromised data platform provider's trusted access reached its systems. Expedia says it contained the issue, is investigating with outside experts, and that its platforms are safe. [13][14][15]

Skift notes that Qubole, whose AWS account was abused against customers on 14–26 October (Weekly p.13), was a former Expedia vendor. Expedia has not said whether the incidents are linked, and researchers have not named Expedia as a Qubole victim. [13]

## HOW THE ATTACK WORKS

No entry vector has been disclosed. SOCRadar, which found an exposed n0n server in September, says every documented n0n intrusion used weak or compromised credentials, with no exploit payloads in the operator's files. MAG212 reads that as a hypothesis for this case, not a finding. [17]

## WHAT TO LOOK FOR

- Credential stuffing against loyalty and travel accounts; points drained or redeemed oddly. [16]
- Phishing quoting real bookings, airports or loyalty IDs (Expedia, Hotels.com, Vrbo, Orbitz).
- Vendor-held tokens, roles or keys to your data platforms still active after the contract ended.

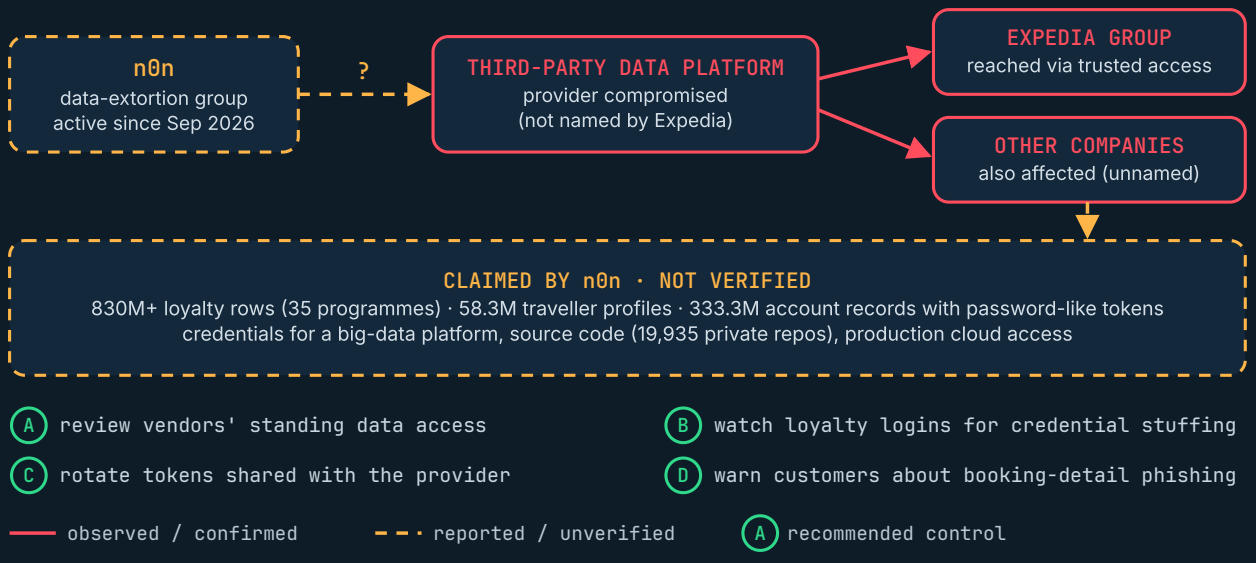



Figure 3 · Solid = Expedia statement; dashed = claims or unproven links. [13] [14] [15]. Diagram: MAG212.

AWS CLI v2 | iam:ListRoles | READ-ONLY | NOT TESTED BY MAG212

AWS · WHO CAN ASSUME YOUR ROLES

```
aws iam list-roles \
  --query "Roles[].{Name:RoleName,Trust:AssumeRolePolicyDocument.Statement[].Principal}" \
  --output json
```

Lists each role and who may assume it. Flag external account IDs you cannot match to a current vendor contract. Run per account.

CONFIRMED BY LA POSTE | 22 MILLION PEOPLE | CNIL NOTIFIED

# Maileva: 22 million people who never chose the vendor

A print-and-mail service inside La Poste's group lost the names and addresses of the people its business clients write to: tenants, patients, employees and debtors.

> STATUS **Confirmed.** La Poste told AFP that names and postal addresses of 22 million people were stolen in a mid-September attack on Maileva; bank details were not compromised. Maileva notified the CNIL, filed a criminal complaint and informed all affected client companies. [18][20]

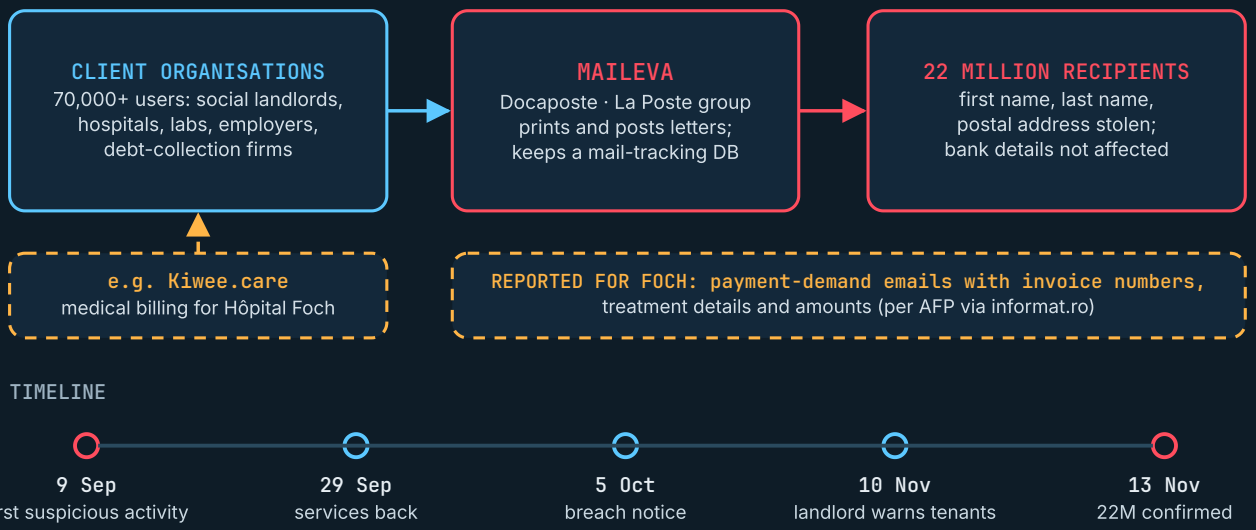



Figure 4 · Who holds whose data. Solid = confirmed by Maileva/La Poste; dashed = reported for individual clients. Timeline: Docaposte notice and client letters via [19]; 22M figure [18] [19]. Diagram: MAG212 CYBER.

## WHAT HAPPENED

Maileva sends payslips, registered letters, invoices and notices on behalf of companies. A Docaposte notice dated 5 October says an attacker reached a mail-tracking database covering about 22 million people, kept for up to ten years depending on the contract. On 13 November Maileva confirmed the 22 million figure. [18][19]

## WHY IT MATTERS

Most of the 22 million have never heard of Maileva. A name, an address and the knowledge that a hospital or landlord writes to you is enough for a convincing fake payment demand. [18][20]

## HOW THE ATTACK WORKS

Not disclosed. Neither the entry point nor the attacker has been named.

### IF YOU OUTSOURCE PRINT, MAIL OR BILLING

1. Ask the provider in writing whether your recipients' data was in the affected database, and for how long it was kept.
2. Warn recipients: you will never ask for payment by a new link or new bank details.
3. Set retention limits in the contract; tracking data rarely needs to live ten years.
4. Map sub-processors: Hôpital Foch's exposure came via Kiwee, a billing subcontractor.

### IF YOU RECEIVED A LETTER

- Treat unexpected payment demands as suspect; call the sender on a number you already have.
- Expect phishing that names your landlord, lab or hospital.

MICROSOFT-CONFIRMED CAMPAIGN | CREDENTIAL THEFT NOT CONFIRMED | US SECTORS

# The outage is over. The outage-themed phishing is not

**STATUS Confirmed by Microsoft Threat Intelligence.** The campaign targeted organisations across several US sectors. Microsoft calls it consistent with credential phishing but has not confirmed that credentials were collected or malware delivered. [21][22][23]

ILLUSTRATIVE RECREATION

From: **ACME Outage Notification** <alerts@moments-notice[.]com>

Subject: **Outage – K7Q2ZP**

Due to incident **EX1550543**, (4) messages were held. They will be **deleted within 24 hours** unless you restore them.

Restore messages

Figure 5 · Illustrative interface — not an actual incident screenshot. Built from the lure details Microsoft published [22] [23].

## ANATOMY OF THE LURE

| | |
|---|---|
| Display name | organisation prefix + 'Outage Notification' or 'Service Health' |
| Sender domain | moments-notice[.]com · planetmaison[.]com · newbooker[.]com |
| Subject | 'Outage – ' + six-character code |
| Real incident ID | EX1550543 is genuine, which makes the lure believable |
| Urgency | held mail deleted in 24 hours |

## WHAT HAPPENED

The disruption began 12 November, hitting Exchange Online, Outlook and Teams; Microsoft blamed servers becoming unavailable and extended mail timeouts to 48 hours. Service was restored by 14 November. On 13 November Microsoft detected emails claiming messages had been held. [22][21]

## WHAT TO DO

- Confirm service status only in the Microsoft 365 admin center Service health page; Microsoft does not email end users about outages. [22][24]
- Check Safe Links policies cover all users; hunt with the queries below or Threat Explorer. [22]
- Block the three domains after checking your telemetry.

KQL · Defender XDR advanced hunting | Security reader | READ-ONLY

NOT TESTED BY MAG212

M365 · OUTAGE LURES AND CLICKS

```
let lures = EmailEvents
| where Timestamp > datetime(2026-11-12)
| where SenderFromDomain in~ ("moments-notice.com", "planetmaison.com", "newbooker.com")
    or (Subject has "Outage" and SenderDisplayName has_any ("Outage Notification",
"Service Health"));
lures | project Timestamp, RecipientEmailAddress, SenderFromAddress, Subject,
DeliveryAction;
// who clicked through Safe Links:
UrlClickEvents | where Timestamp > datetime(2026-11-12)
| join kind=inner (lures | project NetworkMessageId) on NetworkMessageId
| project Timestamp, AccountUpn, Url, ActionType, IsClickedThrough
```

Run the two statements separately if your console returns only the last result. Needs Defender for Office 365; internal 'outage' mail may match the subject filter. For each click-through: reset password, revoke sessions, check MFA methods and inbox rules. [43][22]

# Due tomorrow: UnrealIRCd and Check Point

CISA added both on Friday 13 November with a three-day remediation window for US federal agencies. Everyone else should treat the date as a warning, not a formality.

| CVE | PRODUCT | KEV ADDED | DUE | MAG212 COVERAGE |
|---|---|---|---|---|
| CVE-2026-114713 | UnrealIRCd 6.0.0 – before 6.2.7.1 | 13 Nov | 16 Nov | This page |
| CVE-2026-123039 | Check Point Gaia (DynamicCLI) | 13 Nov | 16 Nov | Weekly p.04–05 |
| CVE-2026-124194 | Microsoft Outlook (NTLMv2 leak) | 10 Nov | 24 Nov | Weekly p.08 |
| CVE-2026-125350 | Google Chrome (use-after-free) | 10 Nov | 24 Nov | Weekly p.08 |

KEV dates from CISA's 13 November alert [26] and our Weekly's KEV board. Due dates bind US federal civilian agencies; CISA urges all organisations to prioritise them.

### UNREALIRCD · WHAT HAPPENED

CVE-2026-114713 (CVSS 8.6) is an improper neutralisation of delimiters in UnrealIRCd, the open-source IRC server. The maintainer, Bram Matthys (Syzop), said attackers were exploiting it on the internet; attempts reportedly began the day after the 7 November hotfix. [25][26]

### WHO IS AFFECTED

UnrealIRCd 6.0.0 up to (not including) 6.2.7.1, on networks with two or more linked servers. Any connected user can trigger it; no IRCOp rights needed. [25]

### CONTAIN AND VERIFY

1. Upgrade every linked server to 6.2.7.1, or apply the hotfix; use the spamfilter workaround only as a stopgap. [25]
2. From an IRC client, run /VERSION against each server and confirm 6.2.7.1 or later.
3. Review IRCOp actions, new links, K-lines and config changes since 7 November.
4. Check Point: finish the sk1000362 log hunt and fixed Jumbo Hotfix take (Weekly p.04–05). [27]

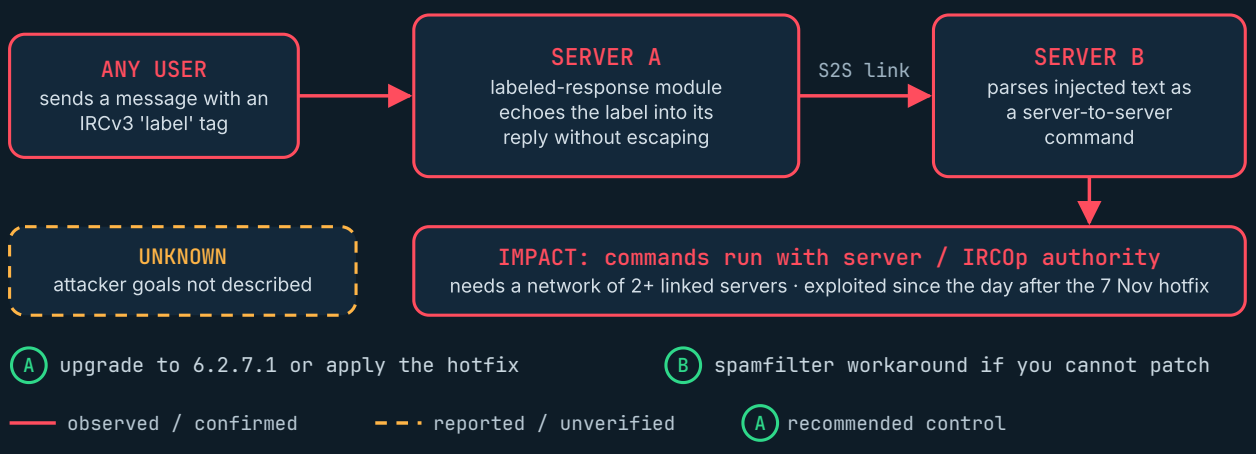



Figure 6 · How CVE-2026-114713 works, per the maintainer's description as reported by The Hacker News [25]. MAG212 did not test this. Diagram: MAG212 CYBER.

# Also on our radar

Stories from your reading list and the wider week, with what changed and what to do. Items older than 72 hours are labelled.

UPDATE · NOT A HACK

### FAA ground stop: FBI reportedly looks at a software update

Officials found no evidence of a cyberattack behind Friday's four-hour nationwide ground stop. Bloomberg's source says the FBI is examining a software update; CBS sources say an update to the private network linking air-traffic facilities made it restart and crash. Secretary Duffy blamed overloaded legacy switches and routers. No official root cause yet. [28][29][30][31]

PATCH

### PostgreSQL fixes 52 CVEs, including a 22-year-old RCE

CVE-2026-6460 (CVSS 8.8), a type confusion in composite types, lets a user who can create objects run code as the database's OS user. Fixed in 18.7, 17.12, 16.16, 15.20 and 14.25 (14's final release); rebuild C-language extensions too. CVE-2026-116322 lets an unauthenticated client restart PostgreSQL 18. No exploitation reported. [32][33]

PRIVACY DESIGN

### Apple explains how Secure Exclaves guard camera and mic data

Apple detailed isolated runtimes that process camera and microphone data outside the main OS and only release full sensor access once the on-screen indicator has been shown. Apple says formal verification found serious kernel flaws early, but does not claim complete coverage. [34]

ESPIONAGE OLDER THAN 72 H

### Turla's 'Cereal' turns Exchange servers into their own C2

ESET found a .NET toolset on government Exchange servers in Kazakhstan and Romania that listens on normal-looking OWA/Autodiscover paths and injects into EdgeTransport.exe without registering a transport agent. A clean agent list does not prove a clean server. [37][38]

MOBILE OLDER THAN 72 H

### iOS 'worse_query' sandbox escape: no confirmed fix

Binance warned that a public sandbox escape could let a malicious app read other apps' files, including wallet seed phrases. Apple's 27.2.1 and 26.7.3 notes do not mention it. Remove sideloaded apps, update, and move funds to a new wallet if you suspect exposure. [39][40]

ALSO NOTED

- Vanderbilt University is restoring systems after a BYOD ransomware attack disclosed on 9 November; the group claims over a million SSNs (unverified). [45]
- OpenAI's GPT-6 Cyber (5 November) reported 81 high-severity Firefox flaws and 948 working Linux kernel exploits; access is restricted to vetted customers. [46]

# Fourteen actions, in order

Owner tags show who usually does the work. Every action points back to its story.

## Today

5 ACTIONS

| | |
|---|---|
| SOC | Run the Entra credential-add query and service-principal sweep (p.04); triage anything unexplained. [1] |
| NETWORK | Upgrade UnrealIRCd to 6.2.7.1 on every linked server; confirm with /VERSION. [25] |
| MSP / RMM | Confirm N-central Legacy API is off and unreachable; run the miner check (p.06). [7][8] |
| EMAIL | Hunt the outage lures and Safe Links clicks (p.09); reset anyone who clicked through. [22] |
| FIREWALL | Check Point: confirm the fixed Jumbo Hotfix take and the sk1000362 log hunt (Weekly p.04–05). |

## Within 48 hours

5 ACTIONS

| | |
|---|---|
| CLOUD | Restrict Kudu/SCM, cloud admin, hypervisor and backup consoles to admin networks. [1] |
| IDENTITY | Review multi-tenant app consent grants and high-privilege apps; remove what nobody owns. [1] |
| PROXY | Alert on server egress to LLM APIs and self-hosted LLM gateways. [1] |
| CLOUD | List IAM roles trusted by external accounts (p.07); remove vendors you no longer use. |
| DBA | Patch PostgreSQL to 18.7/17.12/16.16/15.20/14.25 and rebuild C extensions; plan off 14. [32] |

## This week

4 ACTIONS

| | |
|---|---|
| VENDOR MGMT | Ask print, mail, billing and data-platform vendors which of your customers' data they hold and for how long (p.07–08). |
| LEADERSHIP | Ask your MSP whether it runs N-central and whether the Legacy API was ever enabled. |
| AWARENESS | Brief staff on outage-themed emails and 'press these keys' pages (p.13). |
| CUSTOMER | If you run a loyalty programme, tune credential-stuffing detection and step-up checks on redemptions. |

# Four messages to ignore this week

Share this page with staff. Each trick comes from this issue's stories; the advice is MAG212 guidance.

### “Your messages were held during the outage”

Real outages make fake notices believable, and attackers quote genuine incident numbers.

DO Check service status in the admin portal or with IT. Never restore mail through a link in an email.

### “Press Win+R, then Ctrl+V, then Enter”

Fake CAPTCHA or 'fix' pages ask you to paste a command the page already copied to your clipboard. That runs malware with your rights.

DO No real website asks you to press key combinations. Close the tab and tell IT.

### “Pay your hospital invoice here”

Breaches such as Maileva leak who writes to you. A letter or email can name your real landlord or lab.

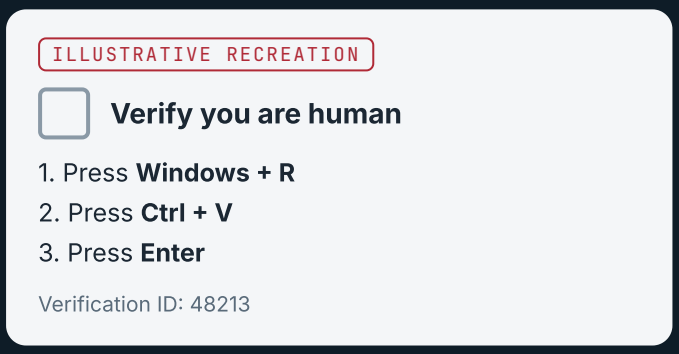
DO Pay only through a portal or number you already use. New bank details are a red flag.

### “Your loyalty points expire today”

Travel and loyalty data claimed in the Expedia case is ideal for convincing lures.

DO Open the app or type the address yourself. Turn on MFA for travel accounts.

ILLUSTRATIVE RECREATION




### Spot the fake check

Step 2 pastes a command the page secretly copied to your clipboard; step 3 runs it. Real CAPTCHAs never ask you to open the Run box or a terminal.

Illustrative interface — not an actual incident screenshot. Pattern: MAG212 guidance.

FOR LEADERS · THREE QUESTIONS FOR MONDAY

1. Which suppliers hold our customers' or patients' data, and how long do they keep it?
2. If an AI-driven intruder reached admin in three hours, how quickly would we see it?
3. Who outside IT can confirm a service outage, and how do staff reach them?

IF YOU ALREADY CLICKED

- Tell IT immediately; speed matters more than embarrassment.
- Change the password from a different, trusted device, and sign out other sessions.
- If you pasted a command, disconnect from the network and leave the computer on for IT.

### Get MAG212 CYBER free

Daily intelligence, weekly deep dives and a defender action list in your inbox.

SUBSCRIBE FREE

cyber.mag212.com/subscribe-free/

# Sources


[1] FBI / IC3 · 13 Nov · FLASH-20261113-001: Threat Overview of ARTEX Autonomous Penetration Testing Framework

[2] The Record · 14 Nov · FBI flash alert says Artex tool automating cyberattacks at 'machine speed'

[3] Cyberpress · 14 Nov · FBI warns hackers are using ARTEX AI agents to breach networks in under 3 hours

[4] SafeBreach / Security Boulevard · 14 Nov · ARTEX AI attack framework: coverage for FBI FLASH-20261113-001

[5] Microsoft Security · 21 Oct · AI-driven, large-scale nation-state campaign using ARTEX enables supply chain compromise

[6] Sophos · 11 Nov · ARTEX: Attackers apply agentic AI

[7] N-able · 14–15 Nov · Status dashboard, event 202613 (HF3 updates)

[8] Huntback · 12 Nov · 4EverAble2Shell is already mining Monero on live N-central servers

[9] Huntress · 10–13 Nov · N-able disables cloud services following disclosure of Able2Shell exploit chains (updates to 13 Nov)

[10] SOCRadar · 13 Nov · 4EverAble2Shell: N-central zero-day exploited

[11] Field Effect · 12 Nov · N-able takes N-central offline as alleged zero-day exploits surface

[12] BleepingComputer · 12 Nov · N-able releases security hotfixes for zero-day vulnerabilities

[13] Skift · 12 Nov · Expedia confirms cyberattack through third-party vendor

[14] Security Affairs · 14 Nov · Extortion group claims theft of nearly one billion Expedia traveler records

[15] Cybernews · Nov · Nearly a billion traveler records at risk as hackers target Expedia

[16] Security Boulevard · 14 Nov · Top 10 breaches of the week

[17] Cyber Daily · 12 Nov · Cyber extortion group claims breach of Expedia

[18] Security Affairs · 15 Nov · Data breach at La Poste subsidiary Maileva affects 22 million people in France

[19] Anas Sahel (tracker) · Nov · Maileva: mail-tracking data of about 22 million people accessed (timeline, Docaposte notice)

[20] Informat.ro (AFP) · 15 Nov · Branch of the French postal service confirms theft of data of 22 million people

[21] CybersecurityNews · 14 Nov · Microsoft warns of active password stealing attack exploiting the Microsoft 365 service disruption

[22] The Tribune (ANI) · 14 Nov · Microsoft restores 365 services globally; warns of phishing campaign exploiting outage

[23] TVC News · 14 Nov · Microsoft warns of phishing campaign exploiting Microsoft 365 service disruption

[24] HawkEye · 14 Nov · Microsoft 365 outage advisory: the phishing started before the service came back

[25] The Hacker News · 14 Nov · Attackers exploit UnrealIRCd flaw that enables server-level command injection

[26] CISA · 13 Nov · CISA adds two known exploited vulnerabilities to catalog

[27] Security Affairs · 14 Nov · U.S. CISA adds Check Point and UnrealIRCd flaws to its KEV catalog

[28] SecurityWeek · 14 Nov · No signs of hacking in FAA ground stop as FBI reportedly probes software update

[29] NBC News · 14 Nov · Thousands of flights canceled or delayed Saturday after FAA's nationwide ground stop

[30] The Next Web · 14 Nov · FBI reportedly probes software update behind US flight ground stop

[31] SecurityWeek · 13 Nov · FAA ground stop issued due to 'extraordinary circumstances'

[32] The Hacker News · 13 Nov · PostgreSQL patches 52 flaws, including 22-year-old bug enabling authenticated RCE

[33] PostgreSQL · 12 Nov · Security information (18.7, 17.12, 16.16, 15.20, 14.25)

[34] SecurityWeek · 14 Nov · Apple reveals how Secure Exclaves shield camera and microphone data

[35] ITPro · 2026 · Cyber professionals are flocking to AI tools, but they're getting tired of fixing mistakes (ISC2)

[36] SecureWorld · 2026 · Why AI is actually increasing the cognitive load on cyber teams (ISC2)

[37] Security Affairs · 10 Nov · Russia-linked APT Turla turned Exchange servers into their own control room

[38] CyberScoop · 10 Nov · Russian Turla group used new toolset to pilfer emails from Kazakhstan, Romanian governments

[39] BleepingComputer · 10 Nov · Binance warns of iOS flaw that could leak crypto wallet secrets

[40] Glitchwire · 11 Nov · iOS sandbox escape puts crypto wallets at risk

[41] DEV Community (Varaxon) · 15 Nov · N-central vulnerabilities (Able2Shell): what MSPs should do now (HF3 status)

[42] Microsoft Learn · docs · Get-MgServicePrincipal (Microsoft Graph PowerShell)

[43] Microsoft Learn · docs · Advanced hunting schema tables (EmailEvents, UrlClickEvents)

[44] Microsoft Learn · docs · Understand service principal creation with new audit log properties

[45] Malwarebytes · 11 Nov · Vanderbilt restores systems after cyberattack as hackers claim data theft

[46] BleepingComputer · 6 Nov · OpenAI's new GPT-6 Cyber AI finds 81 Firefox flaws, 948 Linux bugs


## HOW WE VERIFIED THIS ISSUE

- Primary sources checked where reachable (FBI FLASH, CISA, PostgreSQL, Huntback, vendor statements). Criminal claims (n0n, BYOD) are labelled unverified.
- No command was executed by MAG212. N-central correction: see p.06.
- Not retrieved: SecurityWeek's 'Is AI Making Us Worse at Cybersecurity?' text (headline only) and Microsoft's original post (reported via [21]–[23]).

# Read.
# Verify.
# Defend.

Free daily intelligence, CVE explainers and practical defensive guidance from MAG212 CYBER.

SUBSCRIBE FREE

**Get every issue in your inbox**

cyber.mag212.com/subscribe-free

MAG212 ATLAS · LIVE DEMO

**Map every network before you test it**

demo.atlas.mag212.com

WORK WITH MAG212 CYBER

**Penetration testing**
Find the forgotten edge before someone else does.

**Network & security consulting**
Inventory, segmentation and remote-access design.

**Training**
For staff, IT teams, leaders and schools.

cyber.mag212.com/services

Watch the Atlas walkthrough: youtu.be/j27B5g3viKw

MAG212 CYBER

Daily · 15 November 2026
Understand the threat. Build the defense.