MAG212 CYBER

DAILY INTELLIGENCE

DAILY EDITION

NOVEMBER 2, 2026



TODAY'S THREATS. PRACTICAL DEFENSES.

VPN EXPLOITS / AI TOOLS / IDENTITY / SUPPLY CHAIN

EL MOSTAFA OUCHEN
cyber.mag212.com

CONTENTS | 9 INVESTIGATIONS | 27 TECHNICAL PAGES

Inside the daily edition

Nine investigations, each in three parts: a **briefing** with a threat card and the first commands to run, a **field guide** with indicators, a GUI reconstruction and paired CLI panels, and an **atlas** with an attack-flow diagram, MITRE ATT&CK mapping, a detection query and the ordered response.

PAGE	INVESTIGATION	KEY ARTIFACT	TOOLING INSIDE
03	**Ivanti VPN: close the entire attack chain** EDGE SECURITY	CVE-2026-111357 / -118527 / -104935	PowerShell KQL GUI
06	**A trusted extension becomes a foothold** DEVELOPER SUPPLY CHAIN	anthropic.claude-code 2.1.312 / 2.1.313	zsh PowerShell KQL GUI
09	**Fake IT support reaches the certificate tier** TEAMS + ACTIVE DIRECTORY	AD CS ESC1 / ESC8, events 4886/4887	PowerShell KQL GUI
12	**SysGoRAT: persistence on Linux infrastructure** LINUX THREAT HUNTING	auto_prepend_file, 104.194.158[.]65:443	Bash Sigma GUI
15	**Polymarket: the browser session is the target** WEB SUPPLY CHAIN	Third-party JS, SRI, CSP, registrar lock	Bash HTTP/CSP GUI
18	**Chainalysis: identity exposure after intrusion** AI + CYBER / CUSTOMER DATA	Entra sign-in + audit log review	PowerShell KQL GUI
21	**Binance theft: attribution and asset tracing** CRYPTO + INCIDENT RESPONSE	FBI address list, TraderTraitor	Python Bash GUI
24	**GitLab: patch the runtime, not just the label** PATCH VALIDATION	CVE-2026-97229, Ruby 3.3.13 / 3.4.12 / 4.0.9	Bash GUI
27	**GMO: security warnings need accountable owners** GOVERNANCE	Risk register, owners, verified closure	PowerShell JQL GUI
30	**The first 24 hours: operations checklist** OPERATIONS	Windows + Linux triage collectors	PowerShell Bash
31	**Sources and verification** REFERENCES	10 primary sources + standards	

HOW TO READ THE PANELS

Blue windows are Windows PowerShell; **dark windows with coloured dots** are macOS / Linux shells; **navy code panels** hold KQL, Sigma or config. A **SYNTHETIC OUTPUT** tag means the output lines are invented teaching data, never a real capture.

Legend: PowerShell | Bash / zsh | KQL | Sigma | GUI

EDITION IDENTITY

DAILY: cyan edition badge + a specific calendar date. **WEEKLY:** green visual identity + WEEKLY INTELLIGENCE.

Editorial pictures are conceptual AI illustrations. GUI panels are teaching reconstructions; commands are read-only examples, not recorded executions. Run nothing against systems you are not authorized to examine.

01 / EDGE SECURITY | EXPLOITED IN THE WILD | INTERNET-FACING

Ivanti VPN: close the entire attack chain

STAGE 01	STAGE 02	STAGE 03	STAGE 04
Internet-facing VPN	Chained weaknesses	Root on appliance	Identity exposure

PRODUCT	STATUS	CHAINED CVES	FIXED BUILDS
Ivanti Connect Secure	Exploited in the wild	CVE-2026-111357 CVE-2026-118527 + CVE-2026-104935	25.1.1.3 22.8R2.6
IMPACT	**DOWNSTREAM**	**CAVEAT**	**NOT ENOUGH**
Unauthenticated root on the appliance	LAPS passwords, BitLocker recovery keys queried	Integrity Checker Tool can be evaded	October single-flaw fix

VERIFIED REPORTING

Ivanti disclosed two additional Connect Secure flaws, CVE-2026-111357 and CVE-2026-118527, exploited with CVE-2026-104935. Rapid7 reports fixes in 25.1.1.3 and 22.8R2.6 for the affected branches; the earlier October fixes do not cover both additional flaws. [1]

TECHNICAL EXPLANATION

The chain combines authentication bypass, controlled file placement, and local privilege escalation. Individually modest severity can conceal the combined impact: unauthenticated root access. Research cited by Rapid7 describes credential theft and queries for Windows LAPS passwords and BitLocker recovery keys. This is more than a VPN availability issue. [1]

WHAT TO INVESTIGATE

Compare appliance model, release branch and exact build with the current vendor advisory. Review administrative changes, unusual outbound sessions and downstream directory activity. Rapid7 warns that observed activity can evade the Integrity Checker Tool: a clean result is insufficient assurance. [1]

RESPONSE AND PREVENTION

MAG212 guidance: prioritize exposed gateways outside the routine patch cycle. Preserve volatile evidence where feasible before resets; investigate the identity tier and rotate exposed secrets from a trusted host. If compromise is suspected, follow vendor recovery procedures rather than assuming an in-place patch removes persistence.

Windows PowerShell - Domain Controller (Security log reader) — SYNTHETIC OUTPUT — READ-ONLY

```
# 1. Scope the VPN directory bind account (name is an example)
PS C:\> $svc = "svc-ics-ldap"
PS C:\> Get-ADUser $svc -Properties LastLogonDate,PasswordLastSet,MemberOf |
    Select-Object SamAccountName,LastLogonDate,PasswordLastSet,@{n="Groups";e={$_.MemberOf.Count}}
SamAccountName LastLogonDate         PasswordLastSet       Groups
svc-ics-ldap   11/01/2026 23:41:07   03/14/2024 09:02:55        4
# 2. Directory object reads by that account (needs DS Access auditing + SACLs)
PS C:\> Get-WinEvent -LogName Security -MaxEvents 500 -FilterXPath "*[System[EventID=4662] and
EventData[Data[@Name='SubjectUserName']='$svc']]" |
    Group-Object {$_.Properties[6].Value} | Sort-Object Count -Descending | Select-Object -First 5
Count,Name
   37 %{ms-LAPS-Password read on CN=WS-0142,OU=Workstations,...}
   12 %{msFVE-RecoveryInformation under CN=WS-0088,...}
```

Sources: Rapid7 technical advisory [1]; Volexity investigation (via [1]). Full references on page 31. Commands are read-only examples, not recorded executions.

01 / FIELD GUIDE | GUI + CLI

Ivanti VPN: close the entire attack chain



AI-generated conceptual editorial illustration. Not incident photography or a captured product screen.

INDICATORS, ARTIFACTS & KEY FACTS

TYPE	VALUE	NOTE
CVE	CVE-2026-111357	additional flaw
CVE	CVE-2026-118527	additional flaw
CVE	CVE-2026-104935	chained with both
Fixed build	25.1.1.3 / 22.8R2.6	per branch; check exceptions
HTTP	welcome.cgi + unexpected X-DSCGI-PATH	Volexity lead
File	/tmp/ml.py /tmp/rq5.py	may be deleted
File	/data/.icscap/h.py	forensic image only
Log gap	/home/runtime/dlogs/debuglog	and audit.log

GUI / INVESTIGATION WORKSPACE

MAG212 Edge Inventory - Connect Secure gateways — ILLUSTRATIVE UI / SYNTHETIC DATA

Filter: internet-facing = yes | Compare to advisory | Export CSV

- Gateways (3)
- Advisories
- Admin changes
- Egress sessions
- Identity tier

GATEWAY	BRANCH	BUILD	ADVISORY MIN.	EXPOSURE	ICT RESULT	VERDICT
vpn-gw-01.example.net	22.8	22.8R2.5	22.8R2.6	Public 443	clean	BELOW FIX
vpn-gw-02.example.net	25.1	25.1.1.3	25.1.1.3	Public 443	clean	AT FIX
vpn-gw-03.example.net	25.1	25.1.1.2	25.1.1.3	Public 443	clean	BELOW FIX
Note	ICT "clean"	is not	assurance	per Rapid7		INFO

3 gateways / 2 below fixed build — Rows are synthetic examples

CLI / READ-ONLY INVESTIGATION

Sentinel KQL - VPN bind account directory reads — KQL

```
// Requires 4662 collection from DCs
SecurityEvent
| where EventID == 4662
| where SubjectUserName =~ "svc-ics-ldap"
| summarize Reads = count(),
            Objects = dcount(ObjectName)
    by bin(TimeGenerated, 1h), Computer
| order by TimeGenerated desc
```

bash - forensic image review (mounted read... — OFFLINE

```
# Never run collectors on the live appliance
analyst@srv:~$ sudo mount -o ro,noexec,loop ics.img
/mnt/ics
analyst@srv:~$ ls -la /mnt/ics/tmp/ml.py
/mnt/ics/tmp/rq5.py \
     /mnt/ics/data/.icscap/h.py 2>&1
analyst@srv:~$ sha256sum
/mnt/ics/home/runtime/dlogs/debuglog
# Absent files are leads, not proof of safety
```

REMEDIATE / EXECUTION CHECKLIST

01 Preserve volatile evidence

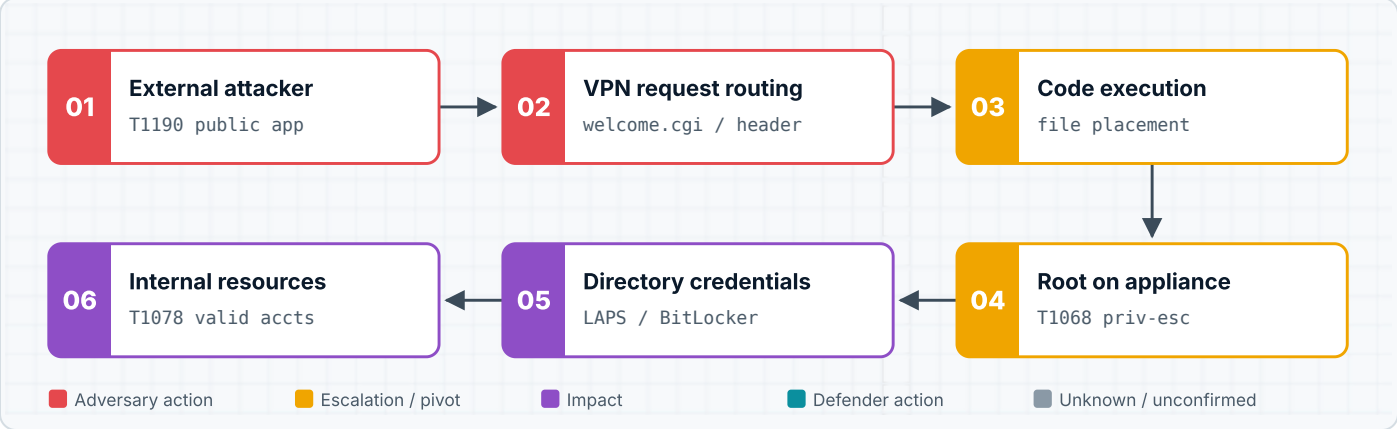
02 Apply complete supported update

03 Recover appliance; rotate secrets

GUI panels are original teaching reconstructions populated with synthetic data (example hostnames, RFC 5737 documentation IP ranges, CONTOSO domain). They are not findings about any environment. Published indicators are reproduced from the cited sources and defanged where they are network destinations. Validate every command in a lab before production use.

01 / ATLAS | MITRE ATT&CK MAPPING | DETECTION

Ivanti VPN: close the entire attack chain



VPN boundary crossed: expand the investigation into Active Directory. Log deletion (T1070) was observed by Volexity.

EXPOSURE
A vulnerable build plus reachability establishes exposure, not compromise. Record model, build and dates before changes.

CORRELATED EVIDENCE
Request-path/header mismatches, log deletion and downstream credential use. Normal login traffic alone is not malicious.

HIGH-PRIORITY LEADS
Gaps in debuglog and audit.log; the three Python filenames in collected images. Filenames may have been deleted.

ATT&CK MAPPING (MAG212 ANALYSIS)

ID	TECHNIQUE → WHAT TO LOOK FOR
T1190	**Exploit Public-Facing Application** Requests to welcome.cgi with unexpected X-DSCGI-PATH
T1068	**Exploitation for Privilege Escalation** Root-owned new files, process anomalies in image
T1070	**Indicator Removal** Gaps in debuglog / audit.log; compare remote syslog
T1078	**Valid Accounts** VPN bind account used from new hosts
T1552	**Unsecured Credentials** LAPS / BitLocker attribute reads (4662)

Splunk SPL - reverse proxy / WAF header ... SPL

```
// Only if your proxy or WAF logs request headers
index=edge sourcetype=waf uri_path="*welcome.cgi*"
| where isnotnull(http_x_dscgi_path)
| eval
mismatch=if(like(http_x_dscgi_path,"%welcome%"),0,
1)
| where mismatch=1
| stats count min(_time) as first max(_time) as
last
        values(http_x_dscgi_path) by src_ip, dest
| convert ctime(first) ctime(last)
```

REMEDIATE / ORDERED RESPONSE

1. Restrict access to the exposed service under an emergency change; preserve memory, disk and remote logs with vendor or forensic support. Do not install arbitrary collectors on the appliance.
2. Apply the full supported update for the actual branch (25.1.1.3 / 22.8R2.6). The earlier single-flaw update is not an all-clear; review model exceptions.
3. Investigate the VPN directory service account, secret access and unusual LDAP activity. Rotate compromised credentials; identify exposed LAPS passwords and recovery material.
4. For suspected compromise, follow vendor recovery steps and validate from trusted infrastructure. Recheck downstream systems, identity alerts and log collection.

CLOSURE CRITERIA Fixed build verified, appliance trust restored, exposed secrets rotated and downstream investigation documented.

Evidence: Rapid7 technical advisory [1]; Volexity investigation (via [1]). Incident details are source-derived; ATT&CK mappings, queries and checklists are MAG212 defensive analysis. No environment was assessed.

02 / DEVELOPER SUPPLY CHAIN | SUPPLY CHAIN | DEVELOPER ENDPOINTS

A trusted extension becomes a foothold

STAGE 01	STAGE 02	STAGE 03	STAGE 04
Marketplace package	Extension loads	SSH key + persistence	Persistent access

CHANNEL	PACKAGE	BAD VERSIONS	DATE SERVED
Cursor extension marketplace	anthropic.claude-code	2.1.312, 2.1.313	October 29
UNAFFECTED	**KEY COMMENT**	**NETWORK**	**CLEANUP**
Official VS Marketplace and npm	svc-telemetry	134.209.82[.]225 : 8099	Uninstall alone may be insufficient

VERIFIED REPORTING

Cursor disclosed tampered anthropic.claude-code versions 2.1.312 and 2.1.313 served through its marketplace on October 29. Microsoft says the official Visual Studio Marketplace and npm distributions were unaffected, and Anthropic did not develop these malicious packages. [2,3]

TECHNICAL EXPLANATION

The injected JavaScript attempted host reconnaissance, SSH-key installation, reverse SSH tunneling and persistence. Privilege and configuration failures prevented some remote-access steps, but keys and persistence artifacts could remain. Cursor says uninstalling the extension alone may be insufficient. [2,3]

WHAT TO INVESTIGATE

Correlate extension source, version and installation time with process and network records. Published indicators include key comment svc-telemetry, destination 134.209.82[.]225, and HTTP port 8099. A matching comment requires key verification; an absent marker does not prove a clean host. [2,3]

RESPONSE AND PREVENTION

MAG212 guidance: preserve the package and logs, isolate hosts with execution evidence, and follow the vendor cleanup procedure. Review every relevant user profile, scheduled task and LaunchAgent. Assess rebuilding where trust cannot be restored. Revoke accessible sessions and rotate exposed developer secrets from a clean device.

zsh - macOS developer workstation — SYNTHETIC OUTPUT — READ-ONLY

```
# Which versions came from the Cursor marketplace?
analyst@mac ~ % ls -1d ~/.cursor/extensions/anthropic.claude-code-* 2>/dev/null
/Users/dev/.cursor/extensions/anthropic.claude-code-2.1.313
# Key comment + fingerprints for verification against the notice
analyst@mac ~ % grep -n "svc-telemetry" ~/.ssh/authorized_keys; ssh-keygen -lf ~/.ssh/authorized_keys
3:ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAA...redacted svc-telemetry
256 SHA256:q2Xh...redacted svc-telemetry (ED25519)
analyst@mac ~ % launchctl list | grep -v com.apple; ls -la ~/Library/LaunchAgents
analyst@mac ~ % lsof -nP -iTCP | grep -E ":8099|134\.209\.82\.225"
# No match is inconclusive: the destination may be offline today
```

Sources: Cursor incident notice [2]; Microsoft package analysis [3]. Full references on page 31. Commands are read-only examples, not recorded executions.

02 / FIELD GUIDE | GUI + CLI

A trusted extension becomes a foothold



AI-generated conceptual editorial illustration. Not incident photography or a captured product screen.

INDICATORS, ARTIFACTS & KEY FACTS

TYPE	VALUE	NOTE
Package	`anthropic.claude-code 2.1.312`	Cursor marketplace
Package	`anthropic.claude-code 2.1.313`	Cursor marketplace
Date	`October 29`	install / update window
SSH key	`comment: svc-telemetry`	verify fingerprint
IP	`134.209.82[.]225`	defanged
Port	`8099 / HTTP`	historic traffic
Technique	`reverse SSH tunnel`	may not have succeeded
Clean	`VS Marketplace, npm`	per Microsoft

GUI / INVESTIGATION WORKSPACE

Extensions - Installed (editor reconstruction) — ILLUSTRATIVE UI / SYNTHETIC DATA

Search: @installed claude | Show source | Copy SHA-256

- Installed
- Recommended
- Disabled
- Outdated
- Provenance log

EXTENSION	VERSION	SOURCE	INSTALLED (UTC)	SIGNATURE	VERDICT
`anthropic.claude-code`	`2.1.313`	Cursor marketplace	2026-10-29 14:02	n/a	TAMPERED VERSION
`ms-python.python`	`2026.10.0`	Open VSX	2026-10-02 08:15	verified	OK
`esbenp.prettier-vscode`	`11.0.0`	Open VSX	2026-09-18 11:40	verified	OK
`redhat.vscode-yaml`	`1.18.0`	Open VSX	2026-08-30 16:21	verified	OK
`local.internal-tools`	`0.4.2`	VSIX file	2026-07-11 10:05	unsigned	REVIEW

1 extension matches the published versions — Synthetic workstation

CLI / READ-ONLY INVESTIGATION

PowerShell - Windows developer workstation — READ-ONLY

```
PS C:\> Get-ChildItem
"$env:USERPROFILE\.cursor\extensions" -Directory |
    Where-Object Name -like "anthropic.claude-code-*"
PS C:\> Select-String
"$env:USERPROFILE\.ssh\authorized_keys" -Pattern svc-
telemetry
PS C:\> Get-ScheduledTask | Where-Object {
    $_.Actions.Execute -match "ssh|node" } |
    Select-Object TaskName,TaskPath,State
PS C:\> Get-NetTCPConnection -RemotePort 8099 -EA 0
```

bash - Linux developer host — READ-ONLY

```
analyst@srv:~$ ls -1d
~/.cursor/extensions/anthropic.claude-code-*
analyst@srv:~$ grep -n svc-telemetry
~/.ssh/authorized_keys
analyst@srv:~$ systemctl --user list-unit-files --
state=enabled
analyst@srv:~$ crontab -l; ls -la ~/.config/autostart
analyst@srv:~$ ss -tnp | grep -E
":8099|134\.209\.82\.225"
# Repeat for every local user profile
```

REMEDIATE / EXECUTION CHECKLIST

01 Isolate confirmed execution

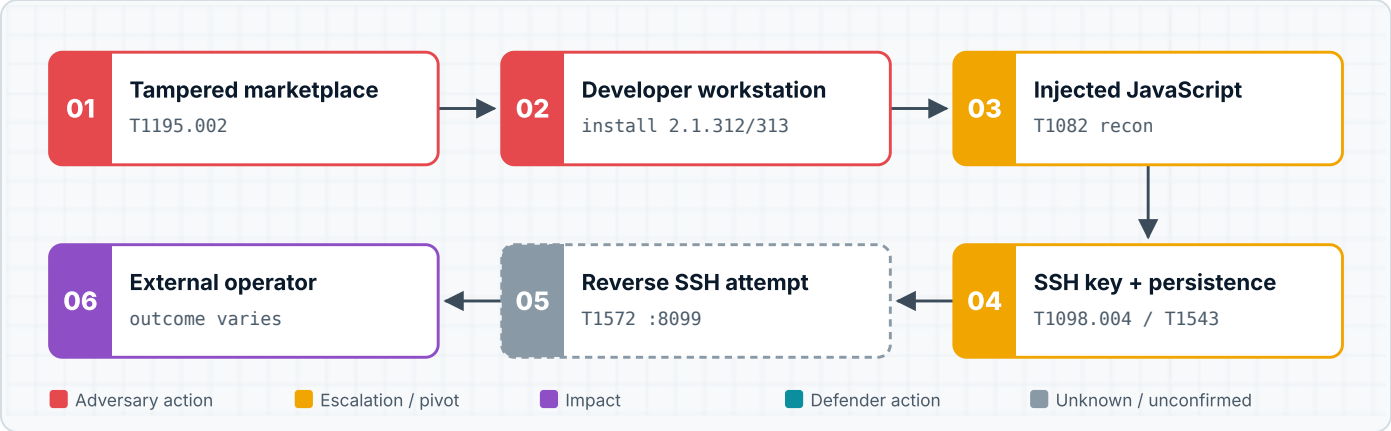
02 Remove persistence, not just plugin

03 Revoke exposed developer access

GUI panels are original teaching reconstructions populated with synthetic data (example hostnames, RFC 5737 documentation IP ranges, CONTOSO domain). They are not findings about any environment. Published indicators are reproduced from the cited sources and defanged where they are network destinations. Validate every command in a lab before production use.

02 / ATLAS | MITRE ATT&CK MAPPING | DETECTION

A trusted extension becomes a foothold



Attempted capability does not mean every tunnel succeeded; persistence artifacts can remain after uninstall.

EXPOSURE
Cursor marketplace origin, version 2.1.312 or 2.1.313 and October 29 install records. A version number without provenance is insufficient.

ENDPOINT EVIDENCE
svc-telemetry plus the published fingerprint; all users' authorized_keys, LaunchAgents, scheduled tasks and startup entries.

NETWORK EVIDENCE
Historic traffic to 134.209.82[.]225 and port 8099; Cursor or child shells spawning ssh.

ATT&CK MAPPING (MAG212 ANALYSIS)

ID	TECHNIQUE → WHAT TO LOOK FOR
T1195.002	**Compromise Software Supply Chain** Package origin = Cursor marketplace, Oct 29
T1082	**System Information Discovery** Node child processes running host recon
T1098.004	**SSH Authorized Keys** New key with comment svc-telemetry
T1543 / T1053	**Launch Agent / Scheduled Task** New LaunchAgent plist or task calling ssh
T1572	**Protocol Tunneling** ssh -R to external host, port 8099 traffic

Defender XDR KQL - editor spawning ss... KQL

```
DeviceNetworkEvents
| where Timestamp > datetime(2026-10-29)
| where RemoteIP == "134.209.82.225"
    or RemotePort == 8099
    or (InitiatingProcessFileName =~ "ssh"
        and InitiatingProcessParentFileName
            has_any ("Cursor", "node"))
| project Timestamp, DeviceName,
    InitiatingProcessCommandLine, RemoteIP,
RemotePort
```

REMEDIATE / ORDERED RESPONSE

1. Preserve extension bytes, timestamps, process trees and persistence artifacts. Isolate endpoints showing payload execution while keeping an out-of-band IR channel.
2. Follow the platform-specific remediation linked in the Cursor notice. Remove the confirmed malicious key, persistence and running tunnel, not just the extension.
3. Review developer credentials, cloud tokens, repository access and browser sessions on the host. Revoke or rotate exposed secrets from a clean workstation.
4. Rebuild if integrity remains uncertain. Re-scan all profiles and startup locations, reinstall from a verified channel and monitor for repeat execution.

CLOSURE CRITERIA Artifacts removed, persistence cannot restart, scope assessed and exposed developer access revoked.

Evidence: Cursor incident notice [2]; Microsoft package analysis [3]. Incident details are source-derived; ATT&CK mappings, queries and checklists are MAG212 defensive analysis. No environment was assessed.

03 / TEAMS + ACTIVE DIRECTORY | SOCIAL ENGINEERING | IDENTITY TIER

Fake IT support reaches the certificate tier

STAGE 01	STAGE 02	STAGE 03	STAGE 04
Fake support call	Remote session	ESC1 or ESC8	Privileged identity

INITIAL ACCESS Teams help-desk impersonation	REMOTE TOOL Quick Assist (user-approved)	PKI PATH A ESC1 template abuse + DCSync	PKI PATH B Coerced DC auth + ESC8 relay
NOT A Teams software exploit	KEY EVENTS 4886 / 4887 / 4768 / 4662	MITIGATION REF Microsoft KB5005413	TRAP Password reset does not revoke certs

VERIFIED REPORTING

Microsoft investigated two intrusions starting with Teams help-desk impersonation and user-approved remote access, primarily Quick Assist. Attackers staged portable runtimes and explored AD CS. One case involved ESC1 and DCSync; another involved coerced domain-controller authentication and ESC8 relay. These were not Teams software exploits. [4]

TECHNICAL EXPLANATION

ESC1 abuses permissive certificate templates to request an authentication certificate for another identity. ESC8 relays NTLM authentication to inadequately protected web enrollment. Password resets do not revoke issued certificates or invalidate every existing ticket. In one case, an identity excluded from containment preserved an exfiltration path. [4]

WHAT TO INVESTIGATE

MAG212 guidance: correlate Teams external-contact records, remote-support sessions, process trees and CA issuance records. Where auditing is enabled, examine Security events 4886/4887 for certificate requests and issuance, 4768 for Kerberos authentication, and 4662 for relevant directory access. Absence may reflect audit configuration or retention.

RESPONSE AND PREVENTION

Restrict template enrollment and requester-supplied identities; review authentication-capable templates with the PKI owner. Disable unused enrollment services and implement supported relay protections. During response, identify and revoke unauthorized certificates, publish revocation information, contain affected identities and endpoints, and assess directory-wide credential exposure.

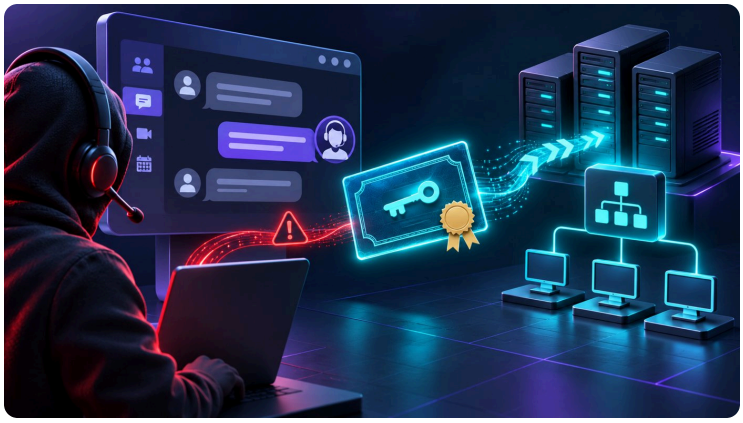
PowerShell - templates where the enrollee supplies the subject (ESC1 precondition) SYNTHETIC OUTPUT READ-ONLY

```
PS C:\> $cfg = (Get-ADRootDSE).configurationNamingContext
PS C:\> Get-ADObject -SearchBase "CN=Certificate Templates,CN=Public Key Services,CN=Services,$cfg" `
  -LDAPFilter "(objectClass=pKICertificateTemplate)" -Properties msPKI-Certificate-Name-
Flag,pKIExtendedKeyUsage |
  Where-Object { ($_."msPKI-Certificate-Name-Flag" -band 1) -and
                 ($_.pKIExtendedKeyUsage -contains "1.3.6.1.5.5.7.3.2") } | Select-Object Name
Name : UserAuth-Legacy        # ENROLLEE_SUPPLIES_SUBJECT + Client Authentication EKU
# Then review who can enroll with the PKI owner (ACL on the template object)
PS C:\> Get-WinEvent -FilterHashtable @{LogName="Security"; Id=4886,4887} -MaxEvents 100 | Select
TimeCreated,Id
```

Sources: Microsoft intrusion research [4]; Microsoft KB5005413. Full references on page 31. Commands are read-only examples, not recorded executions.

03 / FIELD GUIDE | GUI + CLI

Fake IT support reaches the certificate tier



AI-generated conceptual editorial illustration. Not incident photography or a captured product screen.

INDICATORS, ARTIFACTS & KEY FACTS

TYPE	VALUE	NOTE
Vector	`Teams external "IT support"`	social engineering
Tool	`QuickAssist.exe`	user-approved session
Staging	`portable Python / Bun / Java`	suspicious in context
Event	`4886 / 4887`	cert request / issued
Event	`4768`	Kerberos TGT (PKINIT)
Event	`4662`	directory access / DCSync
Path	`ESC1: requester-supplied SAN`	template flag 0x1
Path	`ESC8: HTTP web enrollment relay`	KB5005413

GUI / INVESTIGATION WORKSPACE

certsrv - Certification Authority (CONTOSO-CA) - Issued Certificates — ILLUSTRATIVE UI / SYNTHETIC DATA

View: Issued Certificates | Filter: last 7 days | Revoke... | Publish CRL

- CONTOSO-CA
  - Revoked Certificates
  - **Issued Certificates**
  - Pending Requests
  - Failed Requests
  - Certificate Templates

REQ ID	REQUESTER	TEMPLATE	ISSUED TO (SAN)	SERIAL	EFFECTIVE
4412	CONTOSO\jlee	UserAuth-Legacy	administrator@contoso.com	61 00 00 11 3c	10/27 15:44
4411	CONTOSO\WS-0142$	Machine	WS-0142.contoso.com	61 00 00 11 3b	10/27 15:12
4410	CONTOSO\DC01$	DomainController	DC01.contoso.com	61 00 00 11 3a	10/27 15:09
4407	CONTOSO\mrivera	User	mrivera@contoso.com	61 00 00 11 37	10/26 09:30

Requester jlee obtained a certificate for administrator — Synthetic CA

CLI / READ-ONLY INVESTIGATION

PowerShell - on the CA — READ-ONLY

```
# Issued (Disposition 20) since Oct 1
PS C:\> certutil -view -restrict
"Disposition=20,NotBefore>=10/1/2026" `
  -out "RequestID,RequesterName,CertificateTemplate,
  CommonName,SerialNumber,NotBefore" csv > issued.csv
# ESC8 surface: web enrollment role present?
PS C:\> Get-WindowsFeature ADCS-Web-Enrollment
```

PowerShell - containment (change ticket) — CHANGE

```
# Reason 1 = Key Compromise; then publish CRL
PS C:\> certutil -revoke 610000113c 1
PS C:\> certutil -crl
# Verify relying parties fetch the new CRL
PS C:\> certutil -URL C:\case\4412.cer
# Coordinate with PKI owner before running
```

REMEDIATE / EXECUTION CHECKLIST

01 **Isolate remote-access endpoint**

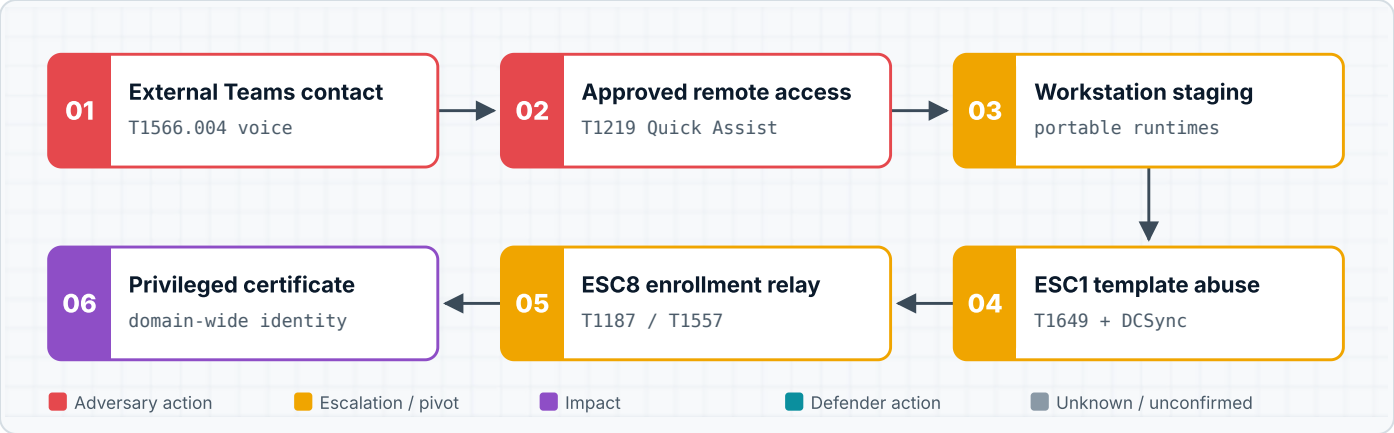
02 **Fix template permissions / relay**

03 **Revoke unauthorized certificates**

GUI panels are original teaching reconstructions populated with synthetic data (example hostnames, RFC 5737 documentation IP ranges, CONTOSO domain). They are not findings about any environment. Published indicators are reproduced from the cited sources and defanged where they are network destinations. Validate every command in a lab before production use.

03 / ATLAS | MITRE ATT&CK MAPPING | DETECTION

Fake IT support reaches the certificate tier



Steps 04 and 05 are two alternative certificate paths from the two intrusions; they are not mandatory consecutive steps.

INITIAL CONTACT
External-tenant Teams contact, Quick Assist session time and Downloads-folder execution. Portable runtimes are suspicious only in context.

CERTIFICATE EVIDENCE
Request ID, requester, template, requested identity and serial on the CA; 4886/4887 where enabled; KDC authentication on the same timeline.

CONTAINMENT GAPS
Privileged identities excluded from automatic response. A certificate can outlive the first compromised endpoint.

ATT&CK MAPPING (MAG212 ANALYSIS)

ID	TECHNIQUE → WHAT TO LOOK FOR
T1566.004	**Phishing: Spearphishing Voice** External Teams call posing as help desk
T1219	**Remote Access Tools** QuickAssist.exe session, then downloads
T1649	**Steal or Forge Authentication Certificates** 4887 where requester ≠ subject
T1003.006	**OS Credential Dumping: DCSync** 4662 replication rights from non-DC
T1187 / T1557	**Forced Authentication / AitM relay** DC NTLM to /certsrv over HTTP

Defender XDR KQL - Quick Assist then n... KQL

```
let qa = DeviceProcessEvents
  | where FileName =~ "QuickAssist.exe"
  | project DeviceId, QAStart = Timestamp;
DeviceProcessEvents
| join kind=inner qa on DeviceId
| where Timestamp between (QAStart .. QAStart +
2h)
| where FolderPath has @"\Downloads\"
| project Timestamp, DeviceName, FileName,
    ProcessCommandLine, AccountName
```

REMEDIATE / ORDERED RESPONSE

1. End the unauthorized remote session and isolate its endpoint. Preserve Teams, EDR, CA, IIS and domain-controller logs; scope every identity used.
2. ESC1: restrict enrollment and requester-supplied subjects on authentication-capable templates. Review effective permissions with the PKI owner; do not disable production templates blindly.
3. ESC8: apply KB5005413 - require HTTPS and Extended Protection for Authentication on enrollment services; disable HTTP and unneeded roles. Test clients and TLS termination.
4. Revoke unauthorized certificates and publish revocation data; verify relying systems obtain it. Address tickets and replicated secrets through coordinated identity recovery.

CLOSURE CRITERIA Unsafe templates and relay paths corrected, certificates revoked, identity scope assessed and containment exclusions reviewed.

Evidence: Microsoft intrusion research [4]; Microsoft KB5005413. Incident details are source-derived; ATT&CK mappings, queries and checklists are MAG212 defensive analysis. No environment was assessed.

04 / LINUX THREAT HUNTING | GOVERNMENT TARGETS | LINUX SERVERS

SysGoRAT: persistence on Linux infrastructure

STAGE 01	STAGE 02	STAGE 03	STAGE 04
Web-shell access	Host discovery	RAT + persistence	Data collection

CLUSTER CL-UNK-1211 (Unit 42)	ACTIVITY WINDOW Oct 2025 - Apr 2026	MALWARE SysGoRAT (Go)	C2 TRANSPORT WebSocket
DEFAULT C2 104.194.158[.]65:443	PERSISTENCE PHP auto_prepend_file	PATHS /etc/php/ , /opt/plesk/php/	AI NOTE AI-assisted scripts, not autonomous op

VERIFIED REPORTING

Unit 42 describes CL-UNK-1211 activity targeting government and state-affiliated infrastructure from October 2025 through April 2026. Its findings include web shells, credential harvesting, database and mailbox targeting, and defense evasion. Script annotations suggest AI-assisted generation; they do not establish that an autonomous AI conducted the entire operation. [5]

TECHNICAL EXPLANATION

The Go-based SysGoRAT supports remote command execution, interactive access and proxying, using WebSocket command-and-control. Unit 42 also observed PHP auto_prepend_file modifications under /etc/php/ and /opt/plesk/php/, enabling code to execute before other PHP scripts. Reported default C2: 104.194.158[.]65:443; operators can change it. [5]

WHAT TO INVESTIGATE

MAG212 guidance: hunt web-server processes spawning shells, unexpected executables in writable directories, new service units, cron changes, unauthorized SSH keys and altered PHP configuration. Correlate outbound connections with process lineage. Encrypted traffic on 443 is not evidence of benign HTTPS, and an IP match alone does not establish infection.

RESPONSE AND PREVENTION

Capture volatile evidence and isolate confirmed hosts while preserving business-critical evidence. Remove the original access path and investigate credentials, mailboxes and databases, not only the RAT binary. Restore from trusted images when integrity is uncertain. Centralize audit and network logs so local deletion cannot erase the only record.

bash - web server (Debian/Ubuntu + Plesk paths) SYNTHETIC OUTPUT READ-ONLY

```
# Effective PHP setting, per SAPI, and where it is set
analyst@srv:~$ sudo grep -rn "auto_prepend_file" /etc/php /opt/plesk/php 2>/dev/null
/etc/php/8.2/fpm/php.ini:2014:;auto_prepend_file =
/etc/php/8.2/fpm/conf.d/99-perf.ini:1:auto_prepend_file = /var/lib/php/sessions/.sess_cache.php
# Inspect, never execute, the referenced file
analyst@srv:~$ sudo stat /var/lib/php/sessions/.sess_cache.php; sudo sha256sum
/var/lib/php/sessions/.sess_cache.php
# Web workers with shell children, and their sockets
analyst@srv:~$ ps -eo pid,ppid,user,lstart,cmd --forest | grep -A3 -E "php-fpm|apache2|nginx"
analyst@srv:~$ sudo ss -tpn state established | grep -v -E "127\.0\.0\.1|::1"
```

Sources: Unit 42 campaign analysis [5]. Full references on page 31. Commands are read-only examples, not recorded executions.

04 / FIELD GUIDE | GUI + CLI

SysGoRAT: persistence on Linux infrastructure



AI-generated conceptual editorial illustration. Not incident photography or a captured product screen.

INDICATORS, ARTIFACTS & KEY FACTS

TYPE	VALUE	NOTE
Cluster	`CL-UNK-1211`	Unit 42 naming
Malware	`SysGoRAT (Go)`	cmd / shell / proxy
C2	`104.194.158[.]65:443`	default; changeable
Transport	`WebSocket over 443`	TLS not = benign
Config	`auto_prepend_file`	php.ini / conf.d
Path	`/etc/php/ /opt/plesk/php/`	compare to baseline
Behaviour	`web worker -> sh/bash`	process lineage
Targets	`databases, mailboxes`	scope beyond host

GUI / INVESTIGATION WORKSPACE

htop-style process tree - web01.example.org — ILLUSTRATIVE UI / SYNTHETIC DATA

Tree view F5 | Filter: www-data | Sockets

PID	PPID	USER	COMMAND	NET	FLAG
812	1	root	`php-fpm: master process (/etc/php/8.2/fpm/php-fpm.conf)`	-	baseline
1931	812	www-data	`└ php-fpm: pool www`	unix	baseline
24410	1931	www-data	`└ sh -c cd /tmp; ./.x`	-	WEB->SHELL
24412	24410	www-data	`└ /tmp/.x (unknown Go ELF)`	203.0.113.50:443 ws	EGRESS
2201	1	mysql	`mysqld --defaults-file=/etc/mysql/my.cnf`	127.0.0.1:3306	baseline

Web worker spawned a shell which runs an unknown binary — Synthetic host / RFC 5737 IP

CLI / READ-ONLY INVESTIGATION

bash - persistence sweep — READ-ONLY

```
analyst@srv:~$ systemctl list-unit-files --
state=enabled
analyst@srv:~$ systemctl list-timers --all
analyst@srv:~$ ls -la /etc/cron.*
/var/spool/cron/crontabs
analyst@srv:~$ sudo awk '{print FILENAME": "$NF}' \
    /root/.ssh/authorized_keys
/home/*/.ssh/authorized_keys
analyst@srv:~$ sudo find /tmp /var/tmp /dev/shm -
type f -perm -u+x -ls
```

bash - auditd (monitor, then search) — AUDIT

```
# /etc/audit/rules.d/50-mag212.rules
-w /etc/php -p wa -k php_conf
-w /opt/plesk/php -p wa -k php_conf
analyst@srv:~$ sudo augenrules --load
analyst@srv:~$ sudo ausearch -k php_conf -i | tail
-20
analyst@srv:~$ sudo journalctl --since "24 hours
ago" -p warning
```

REMEDIATE / EXECUTION CHECKLIST

01 Preserve evidence; contain host

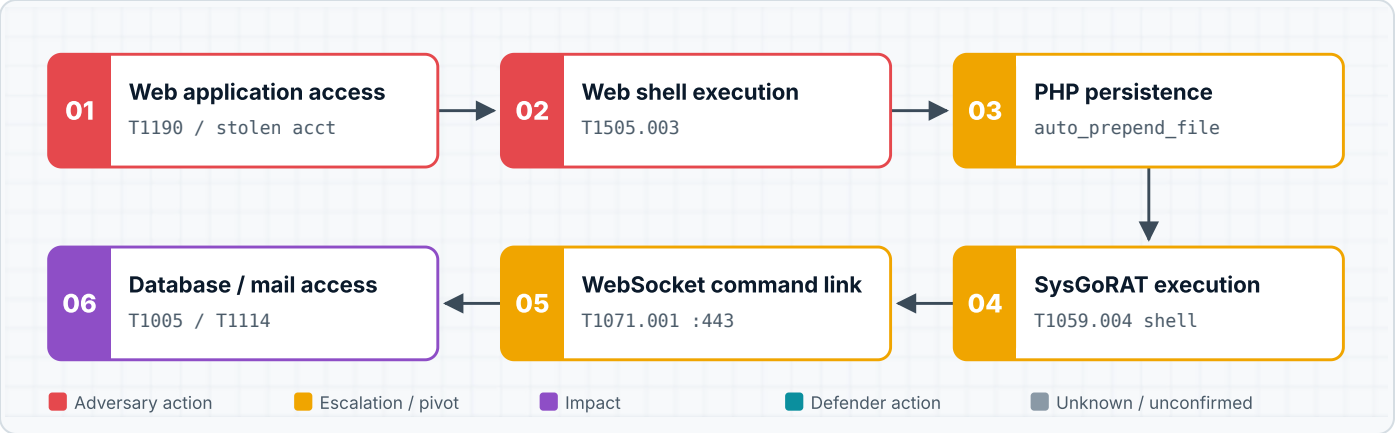
02 Fix entry point; rebuild if needed

03 Rotate application and mail secrets

GUI panels are original teaching reconstructions populated with synthetic data (example hostnames, RFC 5737 documentation IP ranges, CONTOSO domain). They are not findings about any environment. Published indicators are reproduced from the cited sources and defanged where they are network destinations. Validate every command in a lab before production use.

04 / ATLAS | MITRE ATT&CK MAPPING | DETECTION

SysGoRAT: persistence on Linux infrastructure



Conceptual synthesis of observed techniques; not a universal event sequence.

WEB-PROCESS BEHAVIOUR

php-fpm or web-server child shells, new executables and unusual egress. Go binaries or WebSocket traffic alone do not identify SysGoRAT.

PERSISTENCE CHECKS

Compare auto_prepend_file against approved config. Inspect referenced files without executing them. Review units, timers, cron and SSH keys.

SCOPE AND EVASION

Agent shutdown, audit gaps, new mailbox access and database exports. The default C2 is only one lead.

ATT&CK MAPPING (MAG212 ANALYSIS)

ID	TECHNIQUE → WHAT TO LOOK FOR
T1505.003	**Server Software Component: Web Shell** Shell children of php-fpm / httpd
T1059.004	**Unix Shell** sh -c chains from www-data
T1071.001	**Web Protocols (WebSocket C2)** Long-lived 443 sessions from odd binaries
T1543.002 / T1053.003	**Systemd Service / Cron** New units, timers, crontab lines
T1562.001	**Impair Defenses** EDR / auditd stopped, log gaps

Sigma rule - web worker spawns a shell SIGMA

```
title: Web Server Worker Spawns Shell
status: experimental
logsource: {product: linux, category:
process_creation}
detection:
  parent:
    ParentImage|contains: [/php-fpm, /apache2,
/httpd, /nginx]
  child:
    Image|endswith: [/sh, /bash, /dash]
  condition: parent and child
falsepositives: [legitimate CMS cron wrappers]
level: high
```

REMEDIATE / ORDERED RESPONSE

1. Use an authorized forensic collection path for memory, process details and remote logs. Contain confirmed hosts and preserve web access logs before rotation.
2. Investigate the exploited application or stolen account that enabled entry. Correct that path and other exposed software before restoring production.
3. Rebuild from a trusted image when root-level or broad tampering cannot be excluded. Restore approved PHP settings; validate code against signed releases.
4. Rotate application, SSH, database and mail credentials from clean hosts. Check mailbox rules and export activity. Restore centralized audit and egress visibility.

CLOSURE CRITERIA Initial entry corrected, host rebuilt or verified, persistence absent and data-access scope documented.

Evidence: Unit 42 campaign analysis [5]. Incident details are source-derived; ATT&CK mappings, queries and checklists are MAG212 defensive analysis. No environment was assessed.

05 / WEB SUPPLY CHAIN | SESSION RIDING | DNS / REGISTRAR

Polymarket: the browser session is the target

STAGE 01	STAGE 02	STAGE 03	STAGE 04
Registrar / DNS	Trusted script URL	Active session	Unauthorized trades

REIMBURSED $3.9M total	STOLEN DIRECTLY about $2.7M	SCRIPT ORIGIN DebugBear address	ROOT CAUSE Registrar-level DNS compromise (DebugBear)
CREDENTIALS No evidence of password / seed / key theft	TECHNIQUE Session riding in the page	UNKNOWN Registrar entry method	CONTROLS SRI, CSP, registry lock

VERIFIED REPORTING

SecurityWeek reports a $3.9 million Polymarket reimbursement total, including about $2.7 million stolen directly and additional trading-halt losses. Researchers linked malicious JavaScript to a DebugBear address. DebugBear reported registrar-level DNS compromise; it had not identified its targeted customer by name, and Polymarket had not named its provider. [6]

TECHNICAL EXPLANATION

The reported script acted through logged-in trading sessions. Researchers saw no evidence of passwords, seed phrases or private keys being sent to the attacker. This illustrates session riding: hostile code may exercise permissions already available in the page without stealing a password. The registrar entry method remains undisclosed. [6]

WHAT TO INVESTIGATE

MAG212 guidance: retain script responses and hashes, browser/network captures where available, DNS history, registrar audit records and application transaction logs. Correlate affected sessions with script exposure and transaction timestamps. Do not label the reimbursement total as the amount directly stolen.

RESPONSE AND PREVENTION

Remove the compromised dependency and revoke affected sessions. Inventory third-party JavaScript on sensitive pages; eliminate unnecessary scripts. Pin static assets with Subresource Integrity when feasible and use a tested Content Security Policy. CSP host allowlists alone cannot save an already trusted script origin. Protect registrar access with strong MFA, restricted changes and monitoring; use registry locks where supported.

bash - pin and monitor a third-party script — SYNTHETIC OUTPUT — READ-ONLY

```
# 1. Hash the exact bytes you reviewed and approved
analyst@srv:~$ curl -s https://cdn.vendor.example/rum/v3/agent.js -o agent.js
analyst@srv:~$ openssl dgst -sha384 -binary agent.js | openssl base64 -A; echo
u1kX3mQ...redacted...9bQ==
# 2. Compare with what production serves now (scheduled job)
analyst@srv:~$ curl -s https://cdn.vendor.example/rum/v3/agent.js | openssl dgst -sha384 -binary |
openssl base64 -A
Hn7pZ0c...redacted...Lw==        <- differs from approved hash: alert + review
# 3. Registrar protection status (look for server*Prohibited = registry lock)
analyst@srv:~$ whois example.com | grep -i "status"
Domain Status: clientTransferProhibited
```

Sources: SecurityWeek incident reporting [6]. Full references on page 31. Commands are read-only examples, not recorded executions.

05 / FIELD GUIDE | GUI + CLI

Polymarket: the browser session is the target



AI-generated conceptual editorial illustration. Not incident photography or a captured product screen.

INDICATORS, ARTIFACTS & KEY FACTS

TYPE	VALUE	NOTE
Loss	`$3.9M reimbursed`	not all stolen
Loss	`~ $2.7M direct theft`	separate metric
Origin	`DebugBear-linked address`	per researchers
Vector	`registrar-level DNS`	per DebugBear
Not seen	`password / seed / key exfil`	per researchers
Undisclosed	`registrar entry method`	do not infer
Evidence	`script bytes + SHA-384`	preserve offline
Evidence	`DNS / NS history`	registrar logs

GUI / INVESTIGATION WORKSPACE

DevTools - Network (JS) - trade.example.com — ILLUSTRATIVE UI / SYNTHETIC DATA

Filter: JS **3rd-party only** Preserve log Disable cache

NAME	DOMAIN	SIZE	INTEGRITY=	APPROVED HASH	STATUS
app.4f9c.js	trade.example.com	412 kB	sha384-...	match	1ST PARTY
agent.js	cdn.vendor.example	38 kB	(none)	MISMATCH	CHANGED
checkout.js	pay.provider.example	61 kB	sha384-...	match	PINNED
chat-widget.js	widget.support.example	120 kB	(none)	n/a	UNREVIEWED
analytics.js	stats.example.net	22 kB	(none)	n/a	REMOVE?

5 scripts / 3 third-party without SRI — Synthetic page, example domains

CLI / READ-ONLY INVESTIGATION

HTML + HTTP headers - SRI and CSP — CONFIG

```
<script
src="https://cdn.vendor.example/rum/v3/agent.js"
  integrity="sha384-REPLACE_WITH_REAL_HASH"
  crossorigin="anonymous"></script>
// Roll out CSP in report-only mode first
Content-Security-Policy-Report-Only:
  script-src 'self' https://cdn.vendor.example;
  object-src 'none'; report-to csp-ep
```

bash - DNS evidence snapshot — READ-ONLY

```
analyst@srv:~$ dig +short NS example.com
analyst@srv:~$ dig +short A cdn.vendor.example
analyst@srv:~$ dig +trace cdn.vendor.example | tail
-5
analyst@srv:~$ curl -s
https://rdap.org/domain/example.com \
  | jq ".status, .events"
# Today's answer does not prove past safety
```

REMEDIATE / EXECUTION CHECKLIST

01 Remove script; purge caches

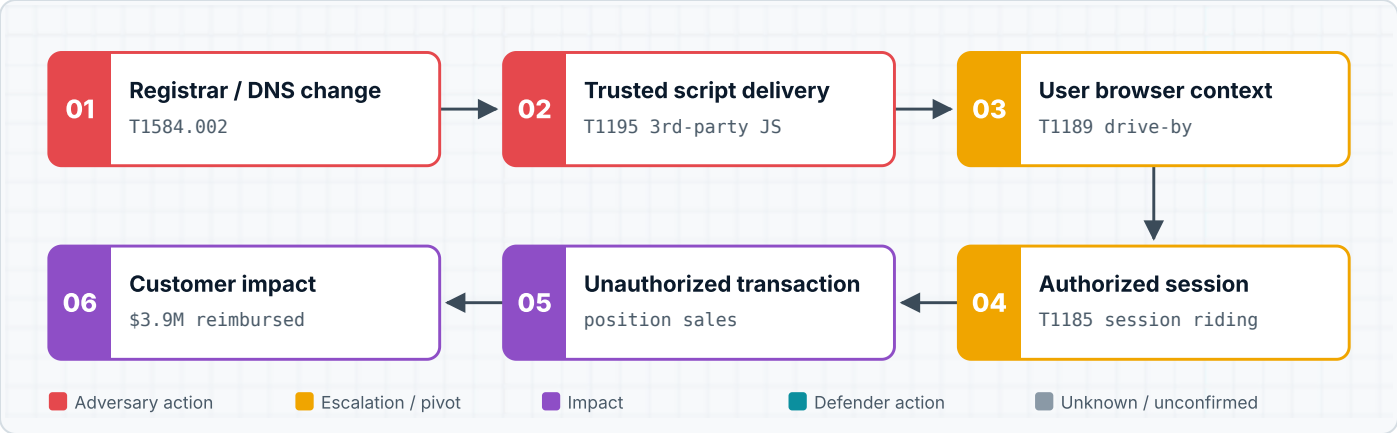
02 Invalidate affected sessions

03 Recover DNS; reconcile losses

GUI panels are original teaching reconstructions populated with synthetic data (example hostnames, RFC 5737 documentation IP ranges, CONTOSO domain). They are not findings about any environment. Published indicators are reproduced from the cited sources and defanged where they are network destinations. Validate every command in a lab before production use.

05 / ATLAS | MITRE ATT&CK MAPPING | DETECTION

Polymarket: the browser session is the target



Reported relationship; the exact registrar access method is undisclosed.

EXPOSURE WINDOW
Script URL, cache lifetime, page inclusion and time window. Never load the suspicious script in a live authenticated browser.

SESSION ABUSE
Script exposure vs authenticated requests, position sales and transfers. No password theft does not exclude fraud through the session.

DNS EVIDENCE
Registrar change logs, nameservers and historical DNS. A good lookup today does not prove what was served then.

ATT&CK MAPPING (MAG212 ANALYSIS)

ID	TECHNIQUE → WHAT TO LOOK FOR
T1584.002	**Compromise Infrastructure: DNS Server** Registrar change log, NS / A record history
T1195	**Supply Chain Compromise** Third-party script hash drift
T1189	**Drive-by Compromise** Malicious JS loaded on sensitive pages
T1185	**Browser Session Hijacking** Trades from valid sessions, odd timing
T1657	**Financial Theft** Transfers reconciled vs session logs

Python - transactions inside the expo... PYTHON

```python
import csv, datetime as dt
start = dt.datetime.fromisoformat("2026-10-
28T00:00+00:00")
end   = dt.datetime.fromisoformat("2026-10-
30T00:00+00:00")
for r in csv.DictReader(open("trades.csv")):
    t = dt.datetime.fromisoformat(r["utc_time"])
    if start <= t <= end and r["script_exposed"]
== "1":
        print(r["session_id"], r["account"],
r["action"], r["amount"])
# Window bounds above are placeholders: use
confirmed times
```

REMEDIATE / ORDERED RESPONSE

1. Remove affected third-party code and purge controlled caches after preserving evidence. Coordinate with the provider; keep the precise asset and exposure timestamps.
2. Invalidate affected sessions and investigate transactions. Reconcile actual theft separately from pause-related losses and reimbursements. Notify users with confirmed scope.
3. Recover registrar control, revoke unauthorized credentials and review nameservers, DNS records and email routing. Strengthen MFA and change controls; monitor out-of-band.
4. Reduce third-party code on sensitive flows. Use SRI for immutable assets, monitor script changes, test CSP report-only before enforcement. DNSSEC alone does not stop registrar abuse.

CLOSURE CRITERIA Clean asset delivery verified across caches, sessions revoked, DNS control restored and transaction reconciliation complete.

Evidence: SecurityWeek incident reporting [6]. Incident details are source-derived; ATT&CK mappings, queries and checklists are MAG212 defensive analysis. No environment was assessed.

06 / AI + CYBER / CUSTOMER DATA | DATA EXPOSURE | PHISHING RISK

Chainalysis: identity exposure after intrusion

STAGE 01	STAGE 02	STAGE 03	STAGE 04
Internal intrusion	Profile information	Targeted pretexts	Identity safeguards

NOTICE DATE	ACTOR	EXPOSED	NOT FOUND
October 30	External, "offensive agentic techniques"	User IDs, names, emails, org roles	Investigation content, credentials, session tokens
ENTRY METHOD	**SEPARATE FROM**	**MAIN RISK**	**BASIS**
Not published	Binance theft	Role-aware phishing	Investigation to date

VERIFIED REPORTING

In its October 30 notice, Chainalysis said an external actor used offensive agentic techniques to access internal systems and tools. Accessed customer information included user IDs, names, email addresses and organization roles. It found no evidence of access to investigation content, authentication credentials or session tokens, based on its investigation to date. [7]

TECHNICAL EXPLANATION

The company explicitly separates this incident from the Binance theft. It has not published a detailed exploit sequence or a technical basis that lets outsiders reconstruct the degree of automation. Attribute the AI characterization to Chainalysis rather than presenting a specific autonomous attack workflow as proven. [7]

WHAT TO INVESTIGATE

MAG212 analysis: identity and role information can support convincing phishing, even if passwords were not exposed. Review unusual password resets, unexpected authentication prompts, changed account recovery details and suspicious contacts referencing real organization roles. Preserve messages and timestamps for correlation with vendor updates.

RESPONSE AND PREVENTION

Enable strong MFA and verify sensitive requests through an established channel. Inform affected staff of the particular data categories disclosed without overstating the impact. Ask the vendor for the organization-specific affected-record breakdown and update the incident register as findings evolve. Separate absence of evidence from a guarantee that an exposure was impossible.

PowerShell (Microsoft Graph) - affected users: sign-ins and MFA methods — SYNTHETIC OUTPUT — READ-ONLY

```
PS C:\> Connect-MgGraph -Scopes "AuditLog.Read.All","UserAuthenticationMethod.Read.All"
PS C:\> $users = Get-Content .\chainalysis-affected-users.txt   # list supplied by the vendor
PS C:\> foreach ($u in $users) {
  Get-MgAuditLogSignIn -Filter "userPrincipalName eq '$u'" -Top 25 |
    Select-Object CreatedDateTime,UserPrincipalName,IpAddress,AppDisplayName,@{n="Err";e=
{$_.Status.ErrorCode}} }
CreatedDateTime      UserPrincipalName        IpAddress      AppDisplayName       Err
11/01/2026 08:14:22  a.khan@contoso.com       198.51.100.24  Microsoft Teams        0
11/01/2026 02:51:09  a.khan@contoso.com       203.0.113.77   Office 365 Exchange 50074
PS C:\> Get-MgUserAuthenticationMethod -UserId a.khan@contoso.com | Select-Object
Id,AdditionalProperties
```

Sources: Chainalysis customer notice [7]. Full references on page 31. Commands are read-only examples, not recorded executions.

06 / FIELD GUIDE | GUI + CLI

Chainalysis: identity exposure after intrusion



AI-generated conceptual editorial illustration. Not incident photography or a captured product screen.

INDICATORS, ARTIFACTS & KEY FACTS

TYPE	VALUE	NOTE
Exposed	user IDs, names, emails	per notice
Exposed	organization roles	per notice
Not found	credentials, session tokens	to date
Not found	investigation content	to date
Unknown	entry chain, IOCs	none published
Risk	role-specific pretexts	prospective
Watch	reset / MFA prompts	identity telemetry
Track	vendor notice version	scope may change

GUI / INVESTIGATION WORKSPACE

Entra admin center - Sign-in logs (reconstruction) — ILLUSTRATIVE UI / SYNTHETIC DATA

Date: last 7 days **User:** in affected list [Add filter] [Download]

- Overview
- Users
- **Sign-in logs**
- Audit logs
- Authentication methods
- Risky users

DATE (UTC)	USER	APPLICATION	IP ADDRESS	MFA	RESULT
11/01 08:14	a.khan	Microsoft Teams	198.51.100.24	Authenticator	Success
11/01 02:51	a.khan	Office 365 Exchange	203.0.113.77	denied by user	Failure 50074
10/31 19:30	b.osei	Self-service password reset	203.0.113.77	SMS	Interrupted
10/31 11:02	c.ruiz	Azure Portal	198.51.100.9	FIDO2	Success

2 events from an unfamiliar IP against affected users — Synthetic tenant

CLI / READ-ONLY INVESTIGATION

Sentinel KQL - recovery changes — KQL

```
AuditLogs
| where OperationName in (
   "User registered security info",
   "Reset password (self-service)",
   "Reset user password")
| extend UPN = tostring(TargetResources[0]
   .userPrincipalName)
| where UPN in~ (affected_users)
```

Sentinel KQL - MFA fatigue signal — KQL

```
SigninLogs
| where UserPrincipalName in~ (affected_users)
| where ResultType in ("50074","500121")
| summarize Prompts = count(),
    IPs = make_set(IPAddress)
    by UserPrincipalName, bin(TimeGenerated,1h)
| where Prompts >= 3
```

REMEDIATE / EXECUTION CHECKLIST

01 Verify requests through known IT

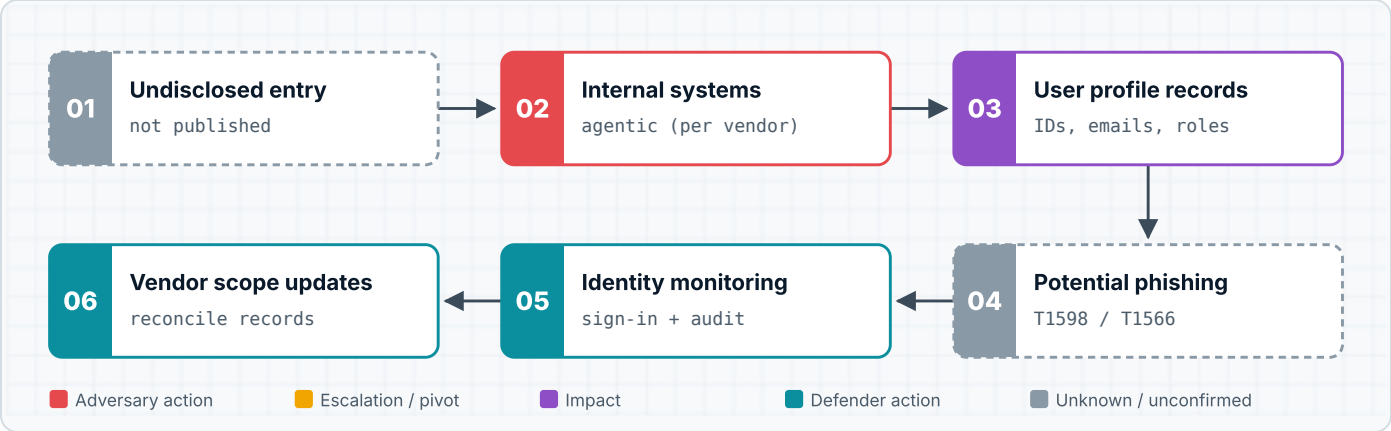
02 Enable phishing-resistant MFA

03 Track final vendor scope

GUI panels are original teaching reconstructions populated with synthetic data (example hostnames, RFC 5737 documentation IP ranges, CONTOSO domain). They are not findings about any environment. Published indicators are reproduced from the cited sources and defanged where they are network destinations. Validate every command in a lab before production use.

06 / ATLAS | MITRE ATT&CK MAPPING | DETECTION

Chainalysis: identity exposure after intrusion



Entry method is unknown; phishing is a prospective risk, not a confirmed follow-on attack.

DIRECT EXPOSURE
Get your organization-specific affected-record list. Using the service is not the same as a confirmed accessed record.

FOLLOW-ON RISK
Unexpected resets, MFA prompts, role-specific impersonation and changed recovery settings. Not every spam message is connected.

EVIDENCE LIMIT
No public exploit chain or IOC set. Do not invent malware hashes or a scan that clears all customers.

ATT&CK MAPPING (MAG212 ANALYSIS)

ID	TECHNIQUE → WHAT TO LOOK FOR
T1589	**Gather Victim Identity Information** Names, emails and roles now in attacker hands
T1598	**Phishing for Information** Messages citing real roles or projects
T1566	**Phishing** Lures to affected mailboxes
T1621	**MFA Request Generation** Bursts of 50074 / denied prompts
T1098	**Account Manipulation** New security info / recovery methods

KQL - define the affected list once

```kql
// Upload the vendor list as a watchlist
let affected_users =
  _GetWatchlist("chainalysis_affected")
  | project tolower(SearchKey);
SigninLogs
| where tolower(UserPrincipalName) in
(affected_users)
| where RiskLevelDuringSignIn in
("medium","high")
| project TimeGenerated, UserPrincipalName,
    IPAddress, AppDisplayName, ResultType
```

REMEDIATE / ORDERED RESPONSE

1. Open an internal incident record; retain the vendor notification, affected users and confirmed data categories. Record what remains pending.
2. Tell exposed users which targeted pretexts to watch for. Verify support and account changes through a known independent channel.
3. Enable strong MFA; investigate anomalous sessions before deciding whether blanket revocation is justified. If compromise is found, revoke sessions and restore recovery controls.
4. Track vendor updates and reconcile the final record list with your investigation. Apply notification obligations through the organization incident process.

CLOSURE CRITERIA Affected records reconciled, account alerts assessed and final vendor findings recorded.

Evidence: Chainalysis customer notice [7]. Incident details are source-derived; ATT&CK mappings, queries and checklists are MAG212 defensive analysis. No environment was assessed.

07 / CRYPTO + INCIDENT RESPONSE | NATION-STATE | FINANCIAL THEFT

Binance theft: attribution and asset tracing

STAGE 01	STAGE 02	STAGE 03	STAGE 04
Asset theft	Published addresses	Chain monitoring	Disruption + reporting

AMOUNT More than $6 billion	DATE Around October 29	ATTRIBUTION North Korea / TraderTraitor (FBI)	ALERT FBI, October 31
OBSERVED Conversion mainly into ETH	EXPECTED DeFi + instant-swap laundering	ENTRY METHOD Not disclosed	AUDIENCE VASPs, custodians

VERIFIED REPORTING

The FBI attributed the theft of more than $6 billion from Binance around October 29 to North Korean activity tracked as TraderTraitor. Its October 31 alert lists associated addresses and says actors were converting assets primarily into Ethereum, with further laundering through DeFi and instant-swap services expected. [8]

TECHNICAL EXPLANATION

The FBI alert establishes an official attribution and provides addresses for disruption. It does not disclose the initial compromise method. Do not invent a private-key leak, malicious transaction-signing mechanism or employee infection as the confirmed entry point. Distinguish observed conversion from anticipated laundering stages. [8]

WHAT TO INVESTIGATE

MAG212 guidance: virtual-asset service providers should integrate the current official address list into their monitoring and response processes. Correlate chain, transaction hash, block time, address, amount and asset type; the same address representation must be interpreted on the correct network. Preserve source provenance and update time.

RESPONSE AND PREVENTION

For organizations operating wallets, use transaction limits, independent approval, segregated signing roles and monitored withdrawal policies. These are general protective controls, not evidence of which safeguard failed here. Escalate matching activity through established incident and legal processes; avoid relying on an old static address list as a complete view of laundering.

Python - screen deposits/withdrawals against the official list — SYNTHETIC OUTPUT — PYTHON

```python
import csv
def norm(a):  # EVM hex is case-insensitive; Bitcoin base58 is NOT
    return a.lower() if a.startswith("0x") else a
iocs = {norm(l.strip()) for l in open("fbi_2026-10-31_addresses.txt") if l.strip()}
for r in csv.DictReader(open("ledger_export.csv")):
    if norm(r["from_addr"]) in iocs or norm(r["to_addr"]) in iocs:
        print(r["network"], r["tx_hash"], r["block_time_utc"], r["asset"], r["amount"], sep=" | ")
ethereum | 0x9f3c...redacted | 2026-10-30T04:12:55Z | ETH | 12.500000
# Record list retrieval date + SHA-256 of the list file in the case notes
```

Sources: FBI attribution and address list [8]. Full references on page 31. Commands are read-only examples, not recorded executions.

07 / FIELD GUIDE | GUI + CLI

Binance theft: attribution and asset tracing



AI-generated conceptual editorial illustration. Not incident photography or a captured product screen.

INDICATORS, ARTIFACTS & KEY FACTS

TYPE	VALUE	NOTE
Source	FBI alert, Oct 31	official list
Actor	TraderTraitor (DPRK)	FBI attribution
Amount	> $6B	around Oct 29
Observed	conversion to ETH	primary
Expected	DeFi / instant swaps	not yet observed
Undisclosed	initial access	do not infer
Fields	network, tx hash, block time	per match
Rule	right chain, right format	avoid false hits

GUI / INVESTIGATION WORKSPACE

MAG212 Tracing Case Board - case VASP-2026-1101 — ILLUSTRATIVE UI / SYNTHETIC DATA

Indicator list: FBI 2026-10-31 (retrieved 11/01 09:00 UTC) [Refresh list] [Export evidence]

- Matches (2)
- Watch-only
- Cleared
- Escalations
- Provenance

NETWORK	TX HASH	BLOCK TIME UTC	ASSET	AMOUNT	DIRECTION	STATUS
ethereum	0x9f3c...a1d2	10-30 04:12	ETH	12.50	inbound	MATCH
ethereum	0x41be...77c0	10-30 06:48	USDT	48,200	inbound	MATCH
bitcoin	5e1f...0c9a	10-29 22:15	BTC	0.75	inbound	NO MATCH
tron	c7a0...e512	10-30 01:03	USDT	9,900	outbound	WATCH

2 matches escalated to compliance — Synthetic hashes, no real addresses

CLI / READ-ONLY INVESTIGATION

bash - fetch tx detail (Etherscan ... — READ-ONLY

```
analyst@srv:~$ curl -s
"https://api.etherscan.io/v2/api?chainid=1\
&module=proxy&action=eth_getTransactionByHash\
&txhash=$TX&apikey=$KEY" | jq ".result |
{from,to,value}"
# Keep raw JSON + retrieval time as evidence
analyst@srv:~$ sha256sum fbi_2026-10-
31_addresses.txt
```

Case record - required fields — TEMPLATE

```
network        : ethereum
tx_hash        : 0x...
block_time_utc : 2026-10-30T04:12:55Z
asset / amount : ETH / 12.5
indicator_src  : FBI alert 2026-10-31
analyst_reason : direct counterparty match
# One wallet hit != attribution of a user
```

REMEDIATE / EXECUTION CHECKLIST

01 Preserve transaction evidence

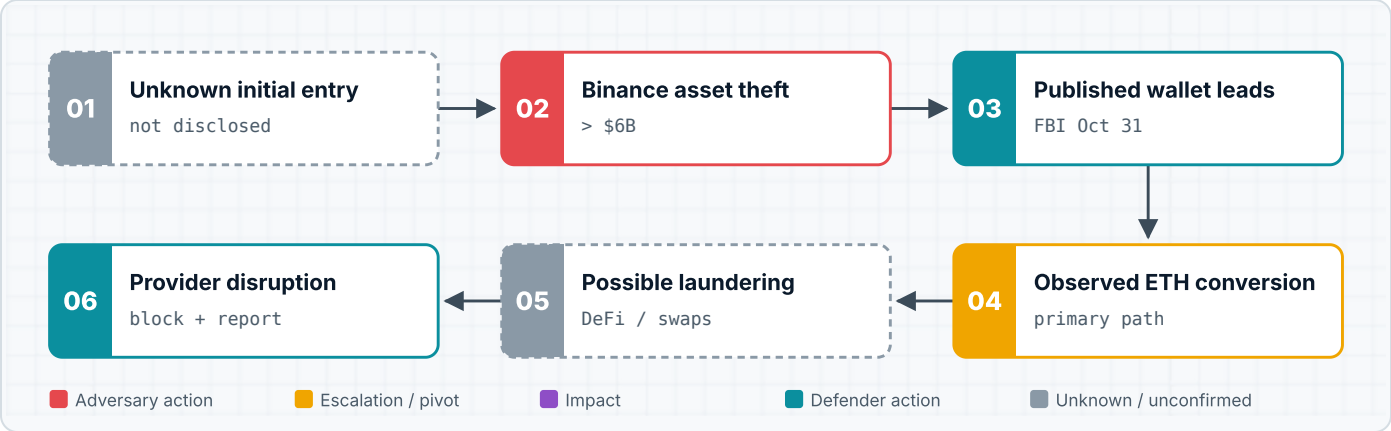
02 Escalate through provider process

03 Monitor updated official indicators

GUI panels are original teaching reconstructions populated with synthetic data (example hostnames, RFC 5737 documentation IP ranges, CONTOSO domain). They are not findings about any environment. Published indicators are reproduced from the cited sources and defanged where they are network destinations. Validate every command in a lab before production use.

07 / ATLAS | MITRE ATT&CK MAPPING | DETECTION

Binance theft: attribution and asset tracing



Unknown entry and anticipated laundering are explicitly labelled; no exploit is invented.

SERVICE-PROVIDER EXPOSURE
Do your custody or transaction services intersect official indicators? An unrelated crypto holder is not affected.

TRANSACTION EVIDENCE
Correct blockchain and exact tx hash. Record asset, amount, timestamp and counterparty; use specialist analytics.

ATTRIBUTION LIMIT
The FBI alert attributes the theft but not the access chain. A wallet interaction does not prove compromise.

ATT&CK MAPPING (MAG212 ANALYSIS)

ID	TECHNIQUE → WHAT TO LOOK FOR
T1657	**Financial Theft** Outflows to listed addresses
(n/a)	**Initial access not disclosed** Do not map an entry technique
Control	**Multi-party approval** Independent signers for large withdrawals
Control	**Velocity limits** Per-hour / per-asset withdrawal caps
Control	**List hygiene** Retrieve current list; log version + hash

SQL - ledger screening job (any wareho... SQL

```sql
SELECT l.network, l.tx_hash, l.block_time_utc,
       l.asset, l.amount, i.list_version
FROM ledger l
JOIN official_indicators i
  ON i.network = l.network
 AND i.address_norm IN (l.from_norm, l.to_norm)
WHERE l.block_time_utc >= TIMESTAMP '2026-10-29'
ORDER BY l.block_time_utc;
```

REMEDIATE / ORDERED RESPONSE

1. Preserve relevant transaction and account records with source provenance. Escalate matches to the provider security and compliance function.
2. Follow official blocking/reporting guidance for your role. Do not attempt recovery transfers to unverified helpers or contact attacker-controlled services.
3. For wallet operators, validate signing controls, transaction limits and independent approvals - general controls, not assumptions about the breach cause.
4. Monitor updated official indicators and reconcile affected assets and customer balances. Document confirmed loss and uncertainty.

CLOSURE CRITERIA Indicator matches investigated, records preserved and required reporting completed; asset recovery may remain unresolved.

Evidence: FBI attribution and address list [8]. Incident details are source-derived; ATT&CK mappings, queries and checklists are MAG212 defensive analysis. No environment was assessed.

08 / PATCH VALIDATION | EXPLOITATION ATTEMPTED | CONTAINERS

GitLab: patch the runtime, not just the label

STAGE 01	STAGE 02	STAGE 03	STAGE 04
Application upgrade	Custom Ruby remains	Exposure persists	Verify deployed runtime

CVE	COMPONENT	ATTEMPT	OUTCOME
CVE-2026-97229	Ruby regex engine via GitLab CI config	October 20-21 (Alpine Linux)	No evidence of RCE; not ruled out
FIXED RUBY	**REBUILT IMAGE**	**TRAP**	**INTEGRITY**
3.3.13 / 3.4.12 / 4.0.9+	alpinelinux/gitlab:19.3.6 (Oct 27)	Earlier copies of the same tag vulnerable	Packages and signing keys unaffected

VERIFIED REPORTING

Alpine Linux says an attacker attempted CVE-2026-97229 exploitation on October 20-21. It found no evidence of successful code execution but could not rule it out and rebuilt the instance. Alpine reported that checked repositories, packages and releases were not tampered with and signing keys were not on the GitLab server. [9]

TECHNICAL EXPLANATION

The flaw involved Ruby regular-expression processing used by GitLab CI configuration. Alpine upgraded GitLab but its custom image still used an unpatched Ruby. Alpine identifies fixed Ruby versions as 3.3.13, 3.4.12 and 4.0.9 or later. Its rebuilt October 27 alpinelinux/gitlab:19.3.6 image included Ruby 3.3.13; earlier copies of that tag remained vulnerable. [9]

WHAT TO INVESTIGATE

MAG212 guidance: inventory actual runtime versions inside the deployed image, not merely the application version in a dashboard. Check image digest, build provenance and dependency inventory. Correlate unusual CI submissions, worker crashes, 502 responses and missing logs. Suspicious crashes are leads, not proof that remote execution succeeded.

RESPONSE AND PREVENTION

For custom builds, verify both application and runtime corrections; rebuild and redeploy immutable images. Assess exposure during the vulnerable interval and preserve CI, application and outbound-network telemetry. Rotate relevant secrets if compromise cannot be excluded. Define a patch acceptance test that proves the production artifact contains the corrected dependency.

bash - prove which Ruby the running GitLab actually uses — SYNTHETIC OUTPUT — READ-ONLY

```
analyst@srv:~$ docker inspect --format "{{.Config.Image}} {{.Image}}" gitlab
alpinelinux/gitlab:19.3.6 sha256:7d1e...redacted      <- tag is the same; digest tells the truth
analyst@srv:~$ docker exec gitlab ruby -e 'print RUBY_VERSION'; echo
3.3.12
# Branch-aware check against the fixed releases
analyst@srv:~$ have=$(docker exec gitlab ruby -e 'print RUBY_VERSION')
analyst@srv:~$ case "$have" in 3.3.*) need=3.3.13;; 3.4.*) need=3.4.12;; 4.0.*) need=4.0.9;; *)
need=unknown;; esac
analyst@srv:~$ [ "$(printf "%s\n" "$need" "$have" | sort -V | head -1)" = "$need" ] && echo "OK $have"
|| echo "VULNERABLE $have < $need"
VULNERABLE 3.3.12 < 3.3.13
```

Sources: Alpine Linux incident analysis [9]. Full references on page 31. Commands are read-only examples, not recorded executions.

08 / FIELD GUIDE | GUI + CLI

GitLab: patch the runtime, not just the label



AI-generated conceptual editorial illustration. Not incident photography or a captured product screen.

INDICATORS, ARTIFACTS & KEY FACTS

TYPE	VALUE	NOTE
CVE	CVE-2026-97229	Ruby regex / CI config
Window	Oct 20-21	attempt at Alpine
Fixed	Ruby 3.3.13	3.3 branch
Fixed	Ruby 3.4.12	3.4 branch
Fixed	Ruby 4.0.9+	4.0 branch
Image	alpinelinux/gitlab:19.3.6	rebuilt Oct 27
Leads	CI submissions, crashes, 502s	benign causes exist
Gap	missing logs	lowers confidence

GUI / INVESTIGATION WORKSPACE

MAG212 Runtime Verification - production registry — ILLUSTRATIVE UI / SYNTHETIC DATA

Policy: deployed digest must contain fixed Ruby | Rescan SBOM | Block deploy

- GitLab
- Runners
- Registry
- Policies
- History

SERVICE	TAG	DIGEST	APP	RUBY	BUILT	GATE
gitlab-web	19.3.6	sha256:7d1e...	19.3.6	3.3.12	10-24	FAIL
gitlab-web (rebuild)	19.3.6	sha256:c09a...	19.3.6	3.3.13	10-27	PASS
gitlab-sidekiq	19.3.6	sha256:7d1e...	19.3.6	3.3.12	10-24	FAIL
runner-01	19.3.0	sha256:5aa2...	runner	n/a	10-15	N/A

Same tag, two digests: only the Oct 27 build passes — Synthetic registry

CLI / READ-ONLY INVESTIGATION

bash - SBOM / scanner check — READ-ONLY

```
analyst@srv:~$ syft alpinelinux/gitlab:19.3.6 -o
table | grep -i ruby
analyst@srv:~$ trivy image --scanners vuln \
  alpinelinux/gitlab:19.3.6 | grep -i ruby
# Omnibus installs: bundled Ruby path
analyst@srv:~$ /opt/gitlab/embedded/bin/ruby -v
analyst@srv:~$ sudo gitlab-rake gitlab:env:info |
grep -i ruby
```

bash - behavioural leads — READ-ONLY

```
analyst@srv:~$ sudo grep -c " 502 " \
  /var/log/gitlab/nginx/gitlab_access.log
analyst@srv:~$ sudo grep -iE "SIGSEGV|worker.*
(died|killed)" \
  /var/log/gitlab/puma/puma_stderr.log | tail
analyst@srv:~$ sudo zgrep -h "ci/lint"
/var/log/gitlab/nginx/*.gz | wc -l
# Combine signals; each alone is benign
```

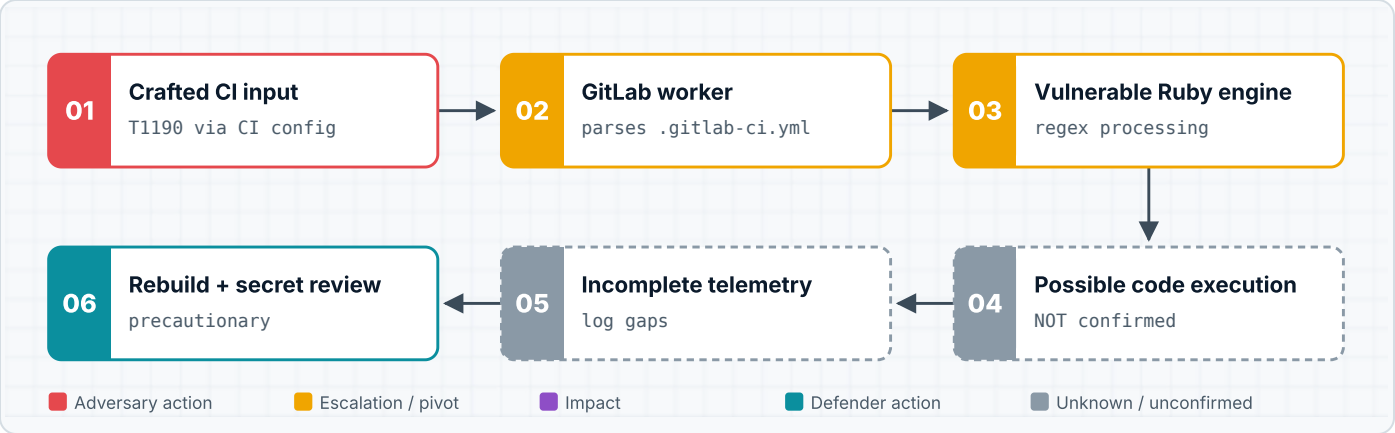
REMEDIATE / EXECUTION CHECKLIST

- 01 Deploy fixed application + runtime
- 02 Assess runners, tokens and secrets
- 03 Verify package and release integrity

GUI panels are original teaching reconstructions populated with synthetic data (example hostnames, RFC 5737 documentation IP ranges, CONTOSO domain). They are not findings about any environment. Published indicators are reproduced from the cited sources and defanged where they are network destinations. Validate every command in a lab before production use.

08 / ATLAS | MITRE ATT&CK MAPPING | DETECTION

GitLab: patch the runtime, not just the label



Possible code execution was not confirmed by Alpine; the rebuild was precautionary.

EXPOSURE
GitLab version, container digest, build date and the Ruby the service uses. A new tag or UI version does not prove the runtime was fixed.

BEHAVIOURAL LEADS
Abnormal CI submissions, worker crashes and HTTP 502s with account creation. Each has benign causes; investigate the combination.

CONFIDENCE LIMIT
Which application, background-job and outbound logs survived? Missing evidence proves neither compromise nor safety.

ATT&CK MAPPING (MAG212 ANALYSIS)

ID	TECHNIQUE → WHAT TO LOOK FOR
T1190	**Exploit Public-Facing Application** Crafted CI configuration submissions
T1499.004	**Application Exploitation (DoS side-effect)** Worker crashes, HTTP 502 bursts
T1136	**Create Account** New accounts near crash timestamps
Control	**Digest pinning** Deploy by sha256, never by tag
Control	**Patch acceptance test** Fail build if Ruby < fixed release

CI gate - .gitlab-ci.yml patch acceptan... YAML

```yaml
verify-runtime:
  stage: test
  image: $CI_REGISTRY_IMAGE@$IMAGE_DIGEST
  script:
    - >-
      ruby -e 'v=Gem::Version.new(RUBY_VERSION);
      m=
{"3.3"=>"3.3.13","3.4"=>"3.4.12","4.0"=>"4.0.9"}
[RUBY_VERSION[0,3]];
      abort("VULNERABLE") unless m && v >=
Gem::Version.new(m)'
# Deploy only the digest that passed this job
```

REMEDIATE / ORDERED RESPONSE

1. Preserve CI configuration, audit records, image metadata and host/network logs. Restrict access if active exploitation is suspected.
2. Deploy a verified fixed application and runtime image. For custom Ruby builds, compare the actual runtime with Alpine's fixed releases; never rely on the tag alone.
3. Rebuild potentially compromised infrastructure, replace relevant secrets and assess runners, tokens and repositories. Independently verify package or release integrity.
4. Test normal pipelines, enforce image-digest tracking and add a dependency-level patch acceptance check. Ensure outbound telemetry survives container replacement.

CLOSURE CRITERIA Running runtime verified, image digest recorded, secrets assessed and pipeline/release integrity reviewed.

Evidence: Alpine Linux incident analysis [9]. Incident details are source-derived; ATT&CK mappings, queries and checklists are MAG212 defensive analysis. No environment was assessed.

09 / GOVERNANCE | RANSOMWARE REPORTED | GOVERNANCE

GMO: security warnings need accountable owners

STAGE 01	STAGE 02	STAGE 03	STAGE 04
Internal warning	Unresolved exposure	Customer impact	Verified remediation

COMPANY GMO (Japan)	SERVICE Onamae.com Desktop Cloud	IMPACT Information exposure + unauthorized transactions	ADMISSION Internal warnings not acted on
RANSOMWARE Reported; payment refused	PROMISE Compensation + security changes	ENTRY CHAIN Not public	LESSON Known risk is not reduced risk

VERIFIED REPORTING

AP reports that GMO apologized after breaches of Onamae.com Desktop Cloud caused information exposure and unauthorized transactions. Its founder acknowledged that internal warnings had not led to adequate action. GMO promised compensation and security changes; AP reported ransomware involvement and that the company refused payment. [10]

TECHNICAL EXPLANATION

The public reporting reviewed for this edition does not establish a detailed initial-access chain. This story supports a governance lesson rather than a speculative malware reconstruction: knowing a weakness exists does not reduce exposure until mitigation is implemented and verified.

WHAT TO INVESTIGATE

MAG212 guidance: audit unresolved critical findings for named ownership, affected assets, exposure, deadlines and evidence of closure. Track exception renewals and escalation to leadership. Compare tickets marked resolved with the deployed configuration and monitored behaviour; a closed ticket alone is not assurance.

RESPONSE AND PREVENTION

Create an escalation rule for overdue internet-facing and identity-tier risks. Document compensating controls and their expiration dates. During an incident, maintain a single approved factual timeline and explain what remains under investigation. Exercise customer notification, restoration and transaction reconciliation procedures before the next crisis.

PowerShell - overdue critical findings from the risk register export — SYNTHETIC OUTPUT — READ-ONLY

```
PS C:\> $today = Get-Date
PS C:\> Import-Csv .\risk-register.csv |
  Where-Object { $_.Status -ne "Closed" -and $_.Severity -in "Critical","High" -and
[datetime]$_.DueDate -lt $today } |
  Select-Object FindingId,Asset,BusinessOwner,TechOwner,DueDate,@{n="DaysOver";e={($today-
[datetime]$_.DueDate).Days}} |
  Sort-Object DaysOver -Descending | Format-Table -AutoSize
FindingId Asset              BusinessOwner TechOwner DueDate     DaysOver
RISK-0417 vdi-gw (internet)  (none)        n.sato    2026-06-30       125
RISK-0388 idp-admin-portal   k.ito         r.mori    2026-09-15        48
# A finding with no business owner is an escalation by itself
```

Sources: AP reporting from Tokyo [10]. Full references on page 31. Commands are read-only examples, not recorded executions.

09 / FIELD GUIDE | GUI + CLI

GMO: security warnings need accountable owners



AI-generated conceptual editorial illustration. Not incident photography or a captured product screen.

INDICATORS, ARTIFACTS & KEY FACTS

TYPE	VALUE	NOTE
Service	`Onamae.com Desktop Cloud`	per AP
Impact	`data exposure`	per AP
Impact	`unauthorized transactions`	per AP
Malware	`ransomware involvement`	reported
Decision	`payment refused`	reported
Root issue	`warnings not acted on`	founder admission
Unknown	`initial access chain`	not public
Customer	`compensation promised`	verify via official channel

GUI / INVESTIGATION WORKSPACE

Risk Register - Critical & High (board view) — ILLUSTRATIVE UI / SYNTHETIC DATA

Filter: open, past due | Escalate to CISO | Request evidence

- **Overdue (3)**
- Due in 30 days
- Exceptions
- Closed - unverified
- Closed - verified

ID	FINDING	ASSET	OWNER	DUE	EVIDENCE	STATE
RISK-0417	MFA not enforced on gateway	vdi-gw	(none)	06-30	none	125 DAYS OVER
RISK-0388	Shared admin account	idp-admin	k.ito	09-15	ticket only	48 DAYS OVER
RISK-0402	Backups not restore-tested	cloud-bkp	y.abe	10-01	none	32 DAYS OVER
RISK-0350	Legacy TLS on API	api-edge	h.kato	08-01	scan 10-20	VERIFIED

3 overdue findings / 1 without a business owner — Synthetic register

CLI / READ-ONLY INVESTIGATION

Jira JQL - overdue internet-facing risks — JQL

```
project = SECRISK
AND priority in (Critical, High)
AND statusCategory != Done
AND duedate < now()
AND labels in (internet-facing, identity)
ORDER BY duedate ASC
```

Risk record schema (required fields) — TEMPLATE

```
finding_id, affected_assets
business_owner, technical_owner
remediation_deadline
compensating_control, control_expiry
verification_evidence, verified_by
residual_risk, escalation_decision_date
```

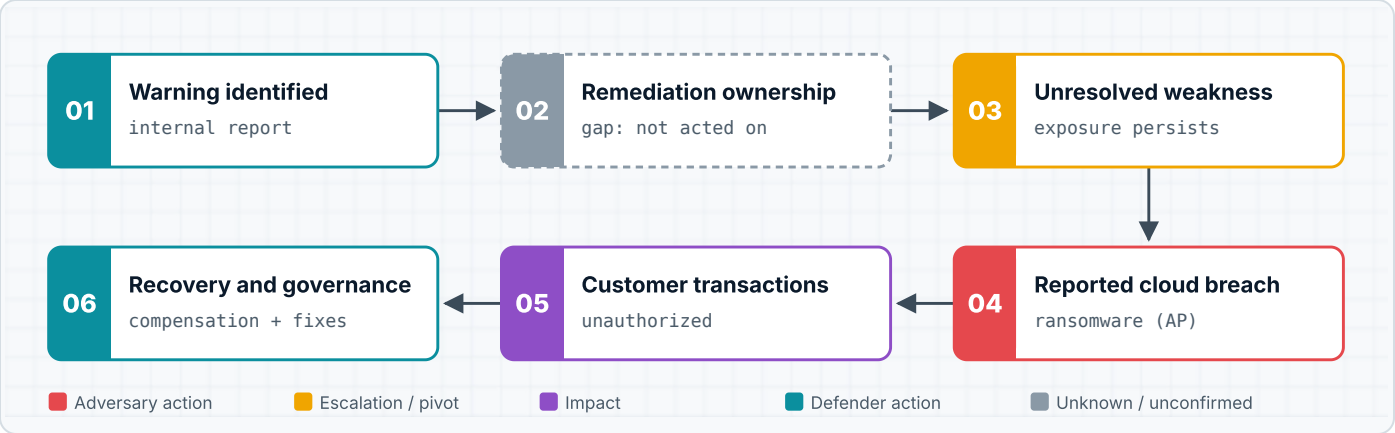
REMEDIATE / EXECUTION CHECKLIST

- **01** Secure affected customer accounts
- **02** Reconcile disputed transactions
- **03** Validate corrective actions

GUI panels are original teaching reconstructions populated with synthetic data (example hostnames, RFC 5737 documentation IP ranges, CONTOSO domain). They are not findings about any environment. Published indicators are reproduced from the cited sources and defanged where they are network destinations. Validate every command in a lab before production use.

09 / ATLAS | MITRE ATT&CK MAPPING | DETECTION

GMO: security warnings need accountable owners



Governance relationship only; this is not a forensic reconstruction of the GMO intrusion.

ORGANIZATIONAL EXPOSURE
Did your organization use the named service, and did it receive a direct provider notification? Reading the news is not exposure.

CUSTOMER EVIDENCE
Provider notices, sign-in history and transaction records vs authorized activity. Preserve discrepancies and case references.

TECHNICAL LIMIT
No full exploit chain or validated IOC set. Generic malware scans cannot resolve provider-side exposure.

ATT&CK MAPPING (MAG212 ANALYSIS)

ID	TECHNIQUE → WHAT TO LOOK FOR
Gov-1	**Named business owner** Every Critical/High finding has one
Gov-2	**Deadline + escalation** Auto-escalate at due date + 0 days
Gov-3	**Compensating control expiry** Exceptions expire and must be re-approved
Gov-4	**Evidence-based closure** Scan, config export or test result attached
Gov-5	**Board visibility** Overdue internet-facing risks reported monthly

Python - closed but unverified findi... PYTHON

```python
import csv
for r in csv.DictReader(open("risk-
register.csv")):
    if r["Status"] == "Closed" and not
r["VerificationEvidence"]:
        print("UNVERIFIED", r["FindingId"],
r["Asset"],
              r["TechOwner"], r["ClosedDate"])
# Re-open or attach evidence: a closed ticket
# is not proof the control is deployed
```

REMEDIATE / ORDERED RESPONSE

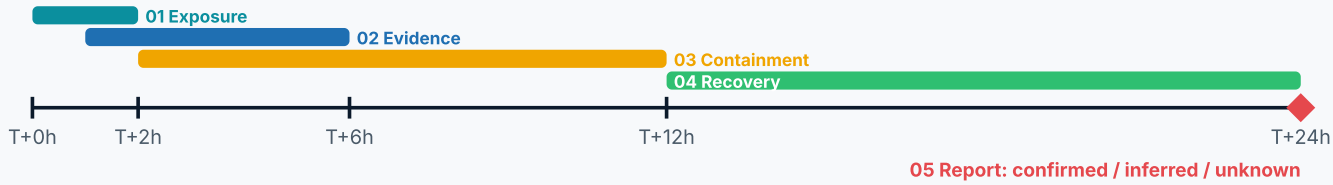
1. Contact the provider through verified channels, secure affected accounts and preserve records of disputed transactions. Avoid links in unsolicited compensation messages.
2. Investigate local endpoint or identity compromise if telemetry supports it; do not assume a provider breach infected every customer device.
3. For your own operations, give each serious finding a business owner, technical owner, deadline and independently verifiable closure condition. Escalate expired exceptions.
4. Exercise restoration, customer communications and reconciliation. Require evidence that the stated control is deployed before marking a finding resolved.

CLOSURE CRITERIA Customer impact reconciled, account controls restored and corrective actions independently verified.

Evidence: AP reporting from Tokyo [10]. Incident details are source-derived; ATT&CK mappings, queries and checklists are MAG212 defensive analysis. No environment was assessed.

OPERATIONS T+0 TO T+24 H

The first 24 hours



01 · 0-2 h

Establish exposure

Inventory Ivanti models and builds, Cursor extension sources and versions, GitLab runtime versions, AD CS enrollment endpoints and third-party scripts on sensitive pages. Record deployed artifacts and assign a technical owner.

02 · 1-6 h

Preserve evidence

Use a trusted collection path. Record timestamps, time zone, hostname, operator and hashes. Keep originals immutable. Gather process trees, identity events, CA issuance records, DNS changes and egress telemetry before retention expires.

03 · 2-12 h

Contain the access path

Isolate confirmed execution points. Revoke affected sessions and unauthorized certificates. Remove compromised scripts or packages. Coordinate certificate and directory actions with PKI and identity owners.

04 · 12-24 h

Recover and verify

Rebuild when integrity cannot be established. Verify fixed dependencies inside the production image. Rotate exposed credentials from clean systems. Confirm alerts, log retention, backups and user access.

05 · 24 h

Report accurately

State what is confirmed, inferred and unknown. Keep stolen amounts separate from reimbursements. Attribute AI claims to the source. Record publication date and event date.

PowerShell - Windows read-only triage collector (run as admin, from trusted media) READ-ONLY

```
PS C:\> $case = "C:\IR\$env:COMPUTERNAME-$(Get-Date -f yyyyMMdd-HHmm)"; New-Item -ItemType Directory $case | Out-Null
PS C:\> Get-CimInstance Win32_Process | Select-Object ProcessId,ParentProcessId,Name,CommandLine,CreationDate | Export-Csv "$case\proc.csv" -NoType
PS C:\> Get-NetTCPConnection -State Established | Select-Object LocalAddress,LocalPort,RemoteAddress,RemotePort,OwningProcess | Export-Csv "$case\net.csv" -NoType
PS C:\> Get-ScheduledTask | Select-Object TaskPath,TaskName,State,@{n="Exec";e={$_.Actions.Execute -join ";"}} | Export-Csv "$case\tasks.csv" -NoType
PS C:\> Get-ItemProperty HKLM:\Software\Microsoft\Windows\CurrentVersion\Run, HKCU:\Software\Microsoft\Windows\CurrentVersion\Run | Out-File "$case\run.txt"
PS C:\> wevtutil epl Security "$case\Security.evtx"
PS C:\> Get-ChildItem $case | Get-FileHash -Algorithm SHA256 | Export-Csv "$case\hashes.csv" -NoType
```

bash - Linux read-only triage collector READ-ONLY

```
analyst@srv:~$ C=/root/ir/$(hostname)-$(date -u +%Y%m%dT%H%MZ); mkdir -p "$C"
analyst@srv:~$ ps -eo pid,ppid,user,lstart,cmd --forest > "$C/ps.txt"
analyst@srv:~$ ss -tupan > "$C/ss.txt"; ip -s addr > "$C/ip.txt"
analyst@srv:~$ systemctl list-unit-files --state=enabled > "$C/units.txt"
analyst@srv:~$ cp -a /etc/cron* /var/spool/cron "$C/" 2>/dev/null
analyst@srv:~$ journalctl --since "48 hours ago" -o short-iso > "$C/journal.txt"
analyst@srv:~$ sha256sum "$C"/* > "$C/SHA256SUMS"
```

CHAIN OF CUSTODY - MINIMUM RECORD

FIELD	EXAMPLE
Collector / operator	`j.doe (IR-on-call)`
Host + time zone	`web01 / UTC`
Start - end (UTC)	`2026-11-02T14:05Z - 14:19Z`
Method + media	`script v1.3, USB #IR-07`
Hash manifest	`SHA256SUMS / hashes.csv`
Storage + access	`evidence share, WORM, 2 people`

Command notes: utilities, privileges and paths vary by platform. Read-only collection still creates normal audit activity and may change access timestamps; wevtutil export requires administrative rights. Review scripts before use; these are investigation aids, not automated remediation.

REFERENCES | 10 PRIMARY SOURCES | 6 TECHNICAL REFERENCES

Sources and verification

INCIDENT SOURCES

[1] **Rapid7 technical advisory**
https://www.rapid7.com/blog/post/etr-additional-ivanti-connect-secure-zero-days-exploited-in-the-wild-cve-2026-111357-cve-2026-118527/

[2] **Cursor incident notice**
https://cursor.com/blog/claude-code-extension-security-notice

[3] **Microsoft package analysis**
https://www.microsoft.com/en-us/security/blog/2026/10/30/inside-malicious-claude-code-branded-cursor-extension-persistent-ssh-backdoor/

[4] **Microsoft intrusion research**
https://www.microsoft.com/en-us/security/blog/2026/10/30/help-desk-impersonation-ad-cs-abuse-investigating-esc1-esc8-intrusion-paths/

[5] **Unit 42 campaign analysis**
https://origin-unit42.paloaltonetworks.com/ai-campaign-targets-linux-infrastructure/

[6] **SecurityWeek incident reporting**
https://www.securityweek.com/polymarket-returns-3-9-million-to-traders-as-debugbear-cites-dns-hijack/

[7] **Chainalysis customer notice**
https://www.chainalysis.com/blog/security-incident-october-30/

[8] **FBI attribution and address list**
https://www.fbi.gov/investigate/cyber/alerts/2026/north-korea-responsible-for-6-billion-binance-attack

[9] **Alpine Linux incident analysis**
https://www.alpinelinux.org/posts/GitLab-security-incident-CVE-2026-97229.html

[10] **AP reporting from Tokyo**
https://apnews.com/article/hack-cyberattack-security-gmo-japan-1a14fac1015f2bbca62a07358615af19

TECHNICAL REFERENCES FOR COMMANDS, QUERIES AND MAPPINGS

R1 **MITRE ATT&CK Enterprise matrix (technique IDs used in the atlases)**
https://attack.mitre.org/

R2 **Microsoft KB5005413: mitigating NTLM relay attacks on AD CS**
https://support.microsoft.com/kb/5005413

R3 **Microsoft Learn: certutil command reference**
https://learn.microsoft.com/windows-server/administration/windows-commands/certutil

R4 **Microsoft Graph PowerShell: Get-MgAuditLogSignIn**
https://learn.microsoft.com/powershell/module/microsoft.graph.reports/get-mgauditlogsignin

R5 **MDN: Subresource Integrity**
https://developer.mozilla.org/docs/Web/Security/Subresource_Integrity

R6 **SigmaHQ: Sigma rule specification**
https://github.com/SigmaHQ/sigma-specification

PUBLICATION NOTES

Editorial standard: original summaries and defensive analysis, with no copied publisher photography. Incident facts come from the numbered sources; MITRE ATT&CK mappings, detection queries, GUI reconstructions and checklists are MAG212 analysis.

Synthetic data: hostnames, users, hashes, serial numbers and IP addresses in GUI panels and sample outputs are invented (example.com / CONTOSO / RFC 5737 ranges). Published indicators are reproduced from sources and defanged.

Status: prepared for publication; not yet posted to cyber.mag212.com. Recheck evolving advisories before deployment. Copyright in original MAG212 layout and analysis does not imply ownership of source reporting or vendor trademarks.