INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-19 06:45 UTC

Claw Chain: Critical OpenClaw Vulnerabilities (CVE-2026-44112, 44113, 44115, 44118) Enable Data Theft, Privilege Escalation, and Persistent Access

CVE VULNERABILITY | **CRITICAL** | CVSS 9.0

SCC Item ID	SCC-CVE-2026-0191
Type	CVE Vulnerability
CVE ID	CVE-2026-44112, CVE-2026-44113, CVE-2026-44115, CVE-2026-44118
Severity	CRITICAL
CVSS Base Score	9.0
EPSS Score	0.0003 (9th percentile)
Affected Products	OpenClaw AI Agent (specific versions not confirmed from available source data)
Published	2 days ago
Discovery Source	Serper

Executive Summary

Four critical vulnerabilities in the OpenClaw AI agent platform, collectively called 'Claw Chain' (CVE-2026-44112, CVE-2026-44113, CVE-2026-44115, CVE-2026-44118), could allow attackers to steal data, escalate privileges, and maintain persistent access on affected systems. Cyera Research discovered the flaws; secondary sources report thousands of servers are potentially exposed. Organizations running OpenClaw in production should treat this as an urgent patching and isolation priority. Official CVSS scores and affected version details are pending NVD publication; verify current exposure against vendor advisory.

Technical Analysis

The 'Claw Chain' cluster comprises four CVEs in OpenClaw, an AI agent platform. The vulnerabilities map to CWE-200 (information exposure), CWE-284 (improper access control), and CWE-269 (improper privilege management). MITRE ATT&CK techniques associated with the chain include T1190 (exploit public-facing application), T1068 (exploitation for privilege escalation), T1098 (account manipulation), T1530 (data from cloud storage), and T1078 (valid accounts). A reported CVSS base score of 9.0 (critical) is cited in vendor advisory and secondary research but has not yet been independently verified by NVD. NVD entries for these CVEs were

not publicly available as of the time this advisory was prepared. EPSS score is 0.031% (9th percentile), indicating low current exploitation probability, though chainable critical flaws in AI agent infrastructure can attract rapid weaponization. A fifth CVE, CVE-2026-32922, covering a separate privilege escalation issue in OpenClaw, was documented by ARMO Security; its relationship to the Claw Chain cluster is unconfirmed. Exact affected versions and patch availability are not confirmed in available source data. Confidence level: medium, verify all technical details and patch applicability directly against vendor advisories and NVD once published before acting.

Action Checklist

1. Step 1: Containment, Identify all OpenClaw AI agent instances in your environment. Isolate internet-facing deployments immediately; restrict inbound access to trusted IP ranges or disable external access until patches are confirmed. Check for CVE-2026-44112, 44113, 44115, and 44118 against your deployed version via the Cyera Research advisory and vendor release notes.
2. Step 2: Detection, Review OpenClaw service logs for anomalous authentication events, unexpected privilege changes, and unusual data access patterns. Map log activity against T1190 (unexpected external connections to the agent API), T1068 (process or role elevation not triggered by authorized users), T1098 (account additions or permission changes), and T1530 (bulk data retrieval from storage). SIEM rules should alert on privilege escalation sequences and unexpected persistence mechanisms in AI agent process trees.
3. Step 3: Eradication, Apply the vendor-issued patch for OpenClaw once confirmed and tested. Rotate all credentials and API keys associated with OpenClaw service accounts. Remove any unauthorized accounts or tokens identified during detection review. Confirm remediation against Cyera Research's published technical details.
4. Step 4: Recovery, After patching, validate that no unauthorized accounts, scheduled tasks, or persistence mechanisms remain. Monitor OpenClaw logs for 30 days post-remediation for re-exploitation indicators. Monitor NVD for official CVSS scores and affected version details; update your risk assessment once published.
5. Step 5: Post-Incident, Audit AI agent platform permissions using least-privilege principles (NIST SP 800-53 AC-6). Review whether AI agent infrastructure is appropriately segmented from sensitive data stores and cloud storage (T1530 exposure). Evaluate whether AI/ML platform deployments are included in your vulnerability management program scope; this cluster highlights a gap common in organizations that treat AI tooling as lower-risk than enterprise software.

IR / Forensic Enrichment

Triage Priority	IMMEDIATE
Escalation Criteria	Escalate to CISO and legal counsel immediately if OpenClaw log review under Step 2 reveals confirmed bulk data retrieval (T1530) from storage systems containing PII, PHI, or regulated financial data, as this triggers breach notification assessment under GDPR, HIPAA, or applicable state privacy law, or if privilege escalation artifacts indicate attacker-controlled accounts were active for more than 24 hours prior to detection.

Recovery Notes	After patching all CVE-2026-44112/44113/44115/44118 affected OpenClaw instances, validate recovery by running the Cyera Research technical indicators (once published) against your patched environment to confirm the specific vulnerability mechanisms are closed — do not rely solely on version number confirmation. Monitor OpenClaw API access logs and OS-level authentication logs continuously for 30 days post-patch, specifically watching for recurrence of the T1098 account manipulation and T1068 privilege escalation patterns identified during detection, as attackers who achieved persistence via the Claw Chain before containment may have implanted mechanisms that survive patching. Re-assess the full scope of data accessible to OpenClaw service accounts and implement network segmentation between the OpenClaw agent layer and sensitive data stores as a durable architectural control to limit blast radius of any future AI agent platform vulnerabilities.
Forensic Artifacts	OpenClaw application access logs: contain source IPs, session tokens, API endpoint paths, and data volumes per request — the primary artifact for scoping CVE-2026-44112 data theft (T1530) and CVE-2026-44118 unauthorized external access (T1190); preserve before any log rotation or instance shutdown | OpenClaw configuration files and credential stores (agent config directory, .env files, API key JSON blobs, environment variables): targeted directly by CVE-2026-44112 data theft vector and may contain rotatable secrets that were exfiltrated; hash and preserve before rotation | OS-level authentication logs (/var/log/auth.log or /var/log/secure on Linux; Windows Security Event Log Event IDs 4624, 4648, 4672, 4720, 4728): record OS-layer privilege escalation and account creation artifacts left by CVE-2026-44113 and CVE-2026-44115 exploitation chains (T1068, T1098) | Scheduled task and cron job listings from OpenClaw host (Linux: /etc/cron*, crontab -l per user, /etc/systemd/system/; Windows: schtasks output, HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Run): capture attacker-planted persistence mechanisms installed during the Claw Chain attack sequence before eradication removes them | Network flow logs or packet captures from the OpenClaw API port during the exposure window: required to identify external IPs that interacted with the vulnerable agent API (T1190), quantify data volumes transferred out (T1530), and determine whether exploitation was targeted or opportunistic given the reported thousands of exposed servers

Per-Action IR Details

Step 1: Containment — Identify all OpenClaw AI agent instances in your environment. Isolate internet-facing deployments immediately; restrict inbound access to trusted IP ranges or disable external access until patches are confirmed. Check for CVE-2026-44112, 44113, 44115, and 44118 against your deployed version via the Cyera Research advisory (cyera.com/blog) and vendor release notes.

NIST Phase: Containment

Reference: NIST 800-61r3 §3.3 — Containment Strategy

Controls: NIST IR-4 (Incident Handling), NIST SC-7 (Boundary Protection), CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 4.4 (Implement and Manage a Firewall on Servers), CIS 4.5 (Implement and Manage a Firewall on End-User Devices)

Compensating: Run 'netstat -tlnp | grep ' or 'ss -tlnp' on Linux hosts to identify all listening OpenClaw agent processes and their bound interfaces. On Windows, use 'netstat -ano | findstr LISTENING' and cross-reference PIDs with 'tasklist'. Use iptables or Windows Firewall ('netsh advfirewall firewall add rule') to block inbound connections to OpenClaw API ports from non-trusted CIDRs immediately. Enumerate all hosts running OpenClaw by querying your asset inventory or scanning with 'nmap -p --open ' to find undocumented instances.

Evidence: Before isolating, capture a full network connection snapshot per host ('ss -tlnp -a' or 'netstat -ano') to record all active sessions to the OpenClaw API at time of containment — this establishes the pre-isolation blast radius. Export OpenClaw service process list including parent-child relationships ('ps auxf' on Linux or Sysinternals Process Explorer on Windows) to baseline what was running. Preserve OpenClaw application logs (default location varies by

deployment — check the OpenClaw installation directory and any configured log output paths in the agent configuration file) without rotation before isolating the instance.

Step 2: Detection — Review OpenClaw service logs for anomalous authentication events, unexpected privilege changes, and unusual data access patterns. Map log activity against T1190 (unexpected external connections to the agent API), T1068 (process or role elevation not triggered by authorized users), T1098 (account additions or permission changes), and T1530 (bulk data retrieval from storage). SIEM rules should alert on privilege escalation sequences and unexpected persistence mechanisms in AI agent process trees.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis

Controls: NIST IR-5 (Incident Monitoring), NIST AU-6 (Audit Record Review, Analysis, and Reporting), NIST AU-2 (Event Logging), NIST SI-4 (System Monitoring), CIS 8.2 (Collect Audit Logs)

Compensating: Without a SIEM, use grep or awk against OpenClaw application logs to surface authentication anomalies: 'grep -iE "(failed|unauthorized|forbidden|privilege|escalat|token|admin|root)" /path/to/openclaw/logs/*.log | sort | uniq -c | sort -rn'. For T1068 (privilege escalation via CVE-2026-44115 or 44113), query Linux auth logs with 'grep -E "sudo|su |NOPASSWD|uid=0" /var/log/auth.log' and cross-reference timestamps with OpenClaw process activity. For T1530 bulk data retrieval (CVE-2026-44112), calculate data volume per session from OpenClaw access logs using 'awk '{sum += $NF} END {print sum}' access.log' (adjust field index to bytes-transferred column). Deploy the free Sigma rule converter (sigma-cli) against OpenClaw logs using community Sigma rules mapped to T1098 and T1068 for offline log scanning. On Windows, query Security Event Log for Event ID 4672 (Special Privileges Assigned) and Event ID 4720 (User Account Created) in the timeframe OpenClaw was exposed.

Evidence: Collect the full OpenClaw authentication and API access logs covering the period of exposure — these will contain the source IPs, session tokens, and endpoint paths targeted during exploitation of T1190 (CVE-2026-44118 external access vector). Capture Linux '/var/log/auth.log' or '/var/log/secure' and Windows Security Event Log (Event IDs 4624, 4625, 4648, 4672, 4720, 4728) for the same window to correlate OS-level privilege changes with OpenClaw API activity associated with CVE-2026-44113 and CVE-2026-44115 escalation chains. Pull OpenClaw agent configuration files and any token/credential store files (check the agent's config directory for API key files, .env files, or credential JSON blobs) — attackers exploiting CVE-2026-44112 for data theft may have accessed or exfiltrated these directly.

Step 3: Eradication — Apply the vendor-issued patch for OpenClaw once confirmed and tested. Rotate all credentials and API keys associated with OpenClaw service accounts. Remove any unauthorized accounts or tokens identified during detection review. Confirm remediation against Cyera Research's published technical details.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST SI-2 (Flaw Remediation), NIST IR-4 (Incident Handling), NIST IA-5 (Authenticator Management), CIS 7.3 (Perform Automated Operating System Patch Management), CIS 7.4 (Perform Automated Application Patch Management), CIS 5.2 (Use Unique Passwords)

Compensating: Before applying the OpenClaw vendor patch, take a filesystem snapshot or VM snapshot of the affected instance to preserve forensic state. Enumerate all API keys and service account tokens in the OpenClaw configuration directory and any connected environment variables ('printenv | grep -iE "key|token|secret|pass"' on Linux; 'Get-ChildItem Env: | Select-String -Pattern "key|token|secret|pass"' on PowerShell) — revoke and rotate every one of these regardless of whether compromise is confirmed, as CVE-2026-44112 specifically enables data theft that could include credential material. Validate patch integrity by comparing the downloaded package hash against the vendor-published checksum before installation. Use 'diff' or a file integrity tool (AIDE or Tripwire free edition) to confirm patch modified only expected files.

Evidence: Before patching, capture a full list of local user accounts and service accounts on the OpenClaw host ('cat /etc/passwd', 'getent passwd', or 'net user /domain' on Windows) to identify accounts added by attackers exploiting CVE-2026-44113 or CVE-2026-44115 for persistence via T1098. Export the OpenClaw agent's current role and permission assignments from its configuration or admin interface — CVE-2026-44115 privilege escalation may have

left elevated roles assigned to attacker-controlled identities that persist after session termination. Dump all active and recent API tokens from the OpenClaw token store before rotation so they can be cross-referenced against external access logs for scope-of-compromise determination.

Step 4: Recovery — After patching, validate that no unauthorized accounts, scheduled tasks, or persistence mechanisms remain. Monitor OpenClaw logs for 30 days post-remediation for re-exploitation indicators. Confirm CVSS scores and affected version ranges against NVD entries once published, and re-assess your exposure accordingly.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST IR-4 (Incident Handling), NIST SI-7 (Software, Firmware, and Information Integrity), NIST AU-11 (Audit Record Retention), CIS 5.3 (Disable Dormant Accounts), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: Enumerate all scheduled tasks and cron jobs on OpenClaw hosts for attacker-planted persistence associated with Claw Chain exploitation: on Linux run 'crontab -l -u ' and 'ls -la /etc/cron*' and 'systemctl list-units --type=service | grep -v systemd'; on Windows run 'schtasks /query /fo LIST /v | findstr /i "openclaw"' and check 'HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Run' and 'HKCU\SOFTWARE\Microsoft\Windows\CurrentVersion\Run' registry keys for OpenClaw-adjacent entries. Deploy osquery with a query against 'crontab', 'scheduled_tasks', and 'startup_items' tables to continuously monitor for re-emergence of persistence mechanisms during the 30-day watch period. Set up a simple log-watch script ('tail -F /path/to/openclaw/access.log | grep -iE "admin|root|escalat|token"') as a lightweight continuous alert proxy if no SIEM is available.

Evidence: After patching, run a file integrity baseline on the OpenClaw installation directory using AIDE ('aide --init' then 'aide --check') or sha256sum against all binaries and config files — compare against pre-patch snapshot to confirm no attacker-modified files survive the patch. Check for OpenClaw plugin or extension directories that may have had malicious modules installed via the Claw Chain persistence vector (T1098 account manipulation could include installing rogue agent plugins or hooks). Verify the NVD entries for CVE-2026-44112, 44113, 44115, and 44118 once published to confirm the exact version ranges and validate that your patched version falls outside the affected range — retain this NVD confirmation as audit documentation.

Step 5: Post-Incident — Audit AI agent platform permissions using least-privilege principles (NIST SP 800-53 AC-6). Review whether AI agent infrastructure is appropriately segmented from sensitive data stores and cloud storage (T1530 exposure). Evaluate whether AI/ML platform deployments are included in your vulnerability management program scope — this cluster highlights a gap common in organizations that treat AI tooling as low-risk.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity

Controls: NIST AC-6 (Least Privilege), NIST IR-4 (Incident Handling), NIST RA-3 (Risk Assessment), NIST SI-5 (Security Alerts, Advisories, and Directives), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process), CIS 3.3 (Configure Data Access Control Lists)

Compensating: Conduct a permission audit of the OpenClaw service account using 'id ' and 'sudo -l -U ' on Linux to confirm the agent runs with minimal OS privileges — it should not have sudo, wheel group membership, or access to sensitive data directories beyond its operational scope. Map OpenClaw's configured data source integrations (check the agent config for database connection strings, S3 bucket ARNs, cloud storage credentials, or API endpoints) against what is operationally required — T1530 exploitation of CVE-2026-44112 was enabled by excessive data access permissions. Add OpenClaw and any other AI/ML agent frameworks (LangChain, AutoGPT, CrewAI, etc.) to your monthly vulnerability scan scope using OpenVAS or Nessus Essentials (free tier) so future CVE clusters in this product class are caught proactively.

Evidence: Produce a lessons-learned document capturing: (1) the timeline from Cyera Research disclosure to organizational awareness and containment for this Claw Chain event, (2) which OpenClaw instances were discovered via the incident versus pre-existing asset inventory, and (3) whether AI/ML platform components were in scope for the

existing vulnerability management program prior to this incident — these gaps are the primary post-incident finding this advisory surfaces. Retain all log collections, network snapshots, and forensic artifacts from Steps 1-4 for a minimum of 90 days (or per your retention policy under NIST AU-11) in case regulatory notification obligations emerge as the NVD entries and breach scope are confirmed.

Detection Guidance

Query authentication and authorization logs on OpenClaw hosts for: (1) privilege escalation events not correlated with approved change tickets; (2) new account creation or role assignment by service accounts; (3) bulk read operations against cloud storage buckets or internal data stores. In SIEM, correlate T1190 indicators (unexpected inbound API calls from external IPs) with subsequent T1068 or T1098 activity within the same session window, this sequence is consistent with the reported chaining pattern. Behavioral indicators include: OpenClaw processes spawning child processes with elevated privileges, unexpected outbound data transfers from agent hosts, and new persistence entries (cron jobs, startup scripts) on OpenClaw server nodes. Note: specific IOC signatures (hashes, IPs, domains) were not present in available source data. Detection rules should focus on behavioral patterns above until vendor or community IOCs are published.

Framework Mappings

MITRE-ATTACK

- **T1068** — Exploitation for Privilege Escalation
- **T1190** — Exploit Public-Facing Application
- **T1098** — Account Manipulation
- **T1530** — Data from Cloud Storage
- **T1078** — Valid Accounts

NIST-800-53R5

- **AC-6** — Least Privilege
- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SI-7** — Software, Firmware, and Information Integrity
- **AC-2** — Account Management
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **AC-3** — Access Enforcement
- **SC-28** — Protection of Information at Rest

OWASP-TOP10-2021

- **A01:2021** — Broken Access Control

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control

CIS-V8

- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **5.4** — Restrict Administrator Privileges to Dedicated Administrator Accounts
- **6.8** — Define and Maintain Role-Based Access Control

SOC2-TSC

- **CC6.1** — The entity implements logical access security software, infrastructure, and architectures over protected information assets
- **CC6.3** — Authorizes, modifies, or removes access

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1068	Exploitation for Privilege Escalation	Privilege-Escalation
T1190	Exploit Public-Facing Application	Initial-Access
T1098	Account Manipulation	Persistence
T1530	Data from Cloud Storage	Collection
T1078	Valid Accounts	Defense-Evasion

Sources

Source	URL	Tier
	https://www.rescana.com/post/claw-chain-critical-openclaw-vulnerabi...	T3
Four OpenClaw Flaws Enable Data Theft, Privilege Escalation, and ...	https://thehackernews.com/2026/05/four-openclaw-flaws-enable-data-t...	T3
Cyera Research Unveil Four Chainable Vulnerabilities in OpenClaw	https://www.cyera.com/blog/claw-chain-cyera-research-unveil-four-ch...	T3

Source	URL	Tier
4 vulnerabilities in OpenClaw AI agent put thousands of servers at risk	https://www.scworld.com/brief/four-vulnerabilities-in-openclaw-ai-a...	T3
CVE-2026-32922: Critical Privilege Escalation in OpenClaw	https://www.armosec.io/blog/cve-2026-32922-openclaw-privilege-escal...	T3
NVD	https://nvd.nist.gov/vuln/detail/CVE-2026-44112, CVE-2026-44113, CV...	T1

Generated 2026-05-19 06:45 UTC by TJS Security Command Center