INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-16 07:14 UTC

China-Nexus APT Maintained Year-Long Persistent Access to US Research Networks via Stolen REDCap Credentials

THREAT CAMPAIGN | **HIGH** | CVSS 7.5

SCC Item ID	SCC-CAM-2026-0476
Type	Threat Campaign
Severity	HIGH
CVSS Base Score	7.5
Affected Products	REDCap (Research Electronic Data Capture), research data management platform; multiple unnamed US research institutions
Published	2026-06-15T13:00:45
Discovery Source	Rss

Executive Summary

A China-nexus APT sustained undetected access to multiple US research institutions for approximately one year by exploiting stolen REDCap credentials, exfiltrating sensitive scientific research data. Google's Threat Intelligence Group discovered and disrupted the operation. The primary business risk is the long-term loss of proprietary research, potential national security implications, and significant reputational harm to affected institutions and their funding relationships.

Technical Analysis

The campaign involved credential-based persistent access to REDCap, a widely deployed secure web application for clinical and translational research data management used across thousands of institutions globally. No CVE is assigned to this campaign; the attack vector relied on stolen credentials rather than an unpatched software vulnerability. Relevant weaknesses map to CWE-200 (Exposure of Sensitive Information), CWE-287 (Improper Authentication), and CWE-522 (Insufficiently Protected Credentials). MITRE ATT&CK techniques observed include: T1078/T1078.004 (Valid Accounts/Cloud Accounts), T1005 (Data from Local System), T1074 (Data Staged), T1041 (Exfiltration Over C2 Channel), T1567 (Exfiltration Over Web Service), T1021 (Remote Services), T1087 (Account Discovery), T1213 (Data from Information Repositories), T1555 (Credentials from Password Stores), T1556 (Modify Authentication Process), T1539 (Steal Web Session Cookie), T1083 (File and Directory Discovery), T1114 (Email Collection), and T1190 (Exploit Public-Facing Application). The approximately 12-month dwell time indicates deliberate low-noise tradecraft designed to blend

credential use with legitimate access patterns. Attribution is assessed as China-nexus at medium confidence; no official CISA or government confirmation has been issued as of the configuration date. Source: Dark Reading (T3); no CVE assigned.

Action Checklist

1. Step 1: Containment, Immediately audit all active REDCap sessions and suspend any accounts showing anomalous login times, source IPs, or access volumes inconsistent with normal researcher behavior. Force a global session invalidation for REDCap and require re-authentication with MFA. Reference NIST AC-12 (Session Termination) and AC-17 (Remote Access).
2. Step 2: Detection, Review REDCap authentication logs and web server access logs for the past 12+ months. Query for: accounts with login activity outside normal business hours, logins from unexpected geographic IP ranges or anonymizing infrastructure (Tor, VPS providers, datacenter ASNs), single accounts accessing unusually large volumes of records, and repeated successful logins after prior failed attempts. Cross-reference against AU-2 (Event Logging) log categories. Enable or verify CIS 8.2 (Collect Audit Logs) compliance across the REDCap infrastructure.
3. Step 3: Eradication, Reset all REDCap account credentials institution-wide. Revoke and reissue API tokens. Enforce MFA for all REDCap accounts, prioritizing externally accessible instances. Review and harden REDCap configuration per CIS 4.6 (Securely Manage Enterprise Assets and Software) and apply NIST AC-3 (Access Enforcement) by auditing role-based access to confirm least privilege. Disable any dormant or unrecognized accounts per CIS 5.3 (Disable Dormant Accounts).
4. Step 4: Recovery, After credential rotation, monitor REDCap access logs daily for a minimum of 30 days for re-emergence of anomalous patterns. Validate that MFA is enforced on all accounts using CIS 6.3 (Require MFA for Externally-Exposed Applications) and CIS 6.5 (Require MFA for Administrative Access). Confirm no unauthorized administrative accounts were created during the dwell period. Cross-reference NIST AU-6 (Audit Record Review, Analysis, and Reporting) cadence requirements.
5. Step 5: Post-Incident, Conduct a privileged access review against NIST AC-6 (Least Privilege) and AC-2 (Account Management) to eliminate over-provisioned research accounts. Implement a formal credential monitoring program aligned with D3-CRO (Credential Rotation) and D3-CH (Credential Hardening). Establish behavioral baselines for REDCap usage to enable anomaly detection going forward. Engage Google Threat Intelligence Group or equivalent for threat hunt services covering the full 12-month dwell window. Review third-party and federated identity integrations for similar credential exposure paths.

IR / Forensic Enrichment

Triage Priority	IMMEDIATE
Escalation Criteria	Escalate immediately to institutional CISO, legal counsel, and federal law enforcement (FBI Cyber Division / CISA) if analysis of `redcap_log_view` confirms exfiltration of federally funded research data, PII, or HIPAA-covered subject data — all of which trigger mandatory breach notification obligations under FISMA, HIPAA, and/or sponsor agency reporting requirements; additionally escalate if the threat hunt reveals lateral movement beyond REDCap into adjacent institutional systems.

Recovery Notes	After credential rotation and MFA enforcement, monitor REDCap authentication and record-access logs daily for a minimum of 30 days, specifically watching for logins from datacenter ASNs or geographic regions inconsistent with the institution's researcher population — this APT rotated infrastructure and may attempt re-entry via a previously undetected compromised account or a newly phished credential. Validate that all REDCap API tokens were regenerated and that no token issued during the 12-month dwell period remains active, as API tokens provide programmatic access that bypasses interactive MFA controls and represent the highest-risk persistent access vector in this campaign. Confirm system integrity of the REDCap application files themselves via checksum comparison against a known-good installation to rule out web shell or application-level backdoor placement during the dwell period.
Forensic Artifacts	REDCap MySQL `redcap_log_view` table: records every individual research record viewed or exported by each user, with timestamp and IP — this is the definitive artifact for scoping the volume and sensitivity of data accessed by the APT over the 12-month dwell period and is unique to REDCap's audit architecture | REDCap MySQL `redcap_log_event` table filtered on action types 'LOGIN', 'LOGIN_FAIL', and 'API_USE': captures the full authentication timeline including failed attempts preceding successful stolen-credential logins, and API token usage events that may represent automated bulk-exfiltration sessions distinct from interactive logins | Apache/Nginx web server access logs for REDCap's `/api/` endpoint: POST requests to the REDCap API with valid tokens generate minimal interactive footprint but high data-transfer volume — large response body sizes (bytes transferred field) in access logs are a primary indicator of bulk record exfiltration via API credential abuse in this campaign | Federated identity provider authentication logs (Shibboleth IdP event logs or Active Directory Security Event Log Event ID 4624/4625): the stolen REDCap credentials may have originated from a phishing campaign or upstream IdP compromise — IdP logs covering the same 12-month window may reveal the initial credential theft event and whether other institutional systems beyond REDCap were accessed with the same stolen identity | Network flow/NetFlow or proxy logs for outbound connections from the REDCap server's IP to non-institutional destinations: China-nexus APT actors in similar campaigns have staged exfiltrated data to cloud storage providers (AWS S3, Alibaba Cloud OSS) or actor-controlled VPS infrastructure — large outbound data transfers from the REDCap host to external IPs during off-hours are a high-confidence exfiltration indicator specific to this campaign's methodology

Per-Action IR Details

Step 1: Containment — Immediately audit all active REDCap sessions and suspend any accounts showing anomalous login times, source IPs, or access volumes inconsistent with normal researcher behavior. Force a global session invalidation for REDCap and require re-authentication with MFA. Reference NIST AC-12 (Session Termination) and AC-17 (Remote Access).

NIST Phase: Containment

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

Controls: NIST AC-12 (Session Termination), NIST AC-17 (Remote Access), CIS 6.3 (Require MFA for Externally-Exposed Applications)

Compensating: Use REDCap's built-in 'Browse Users' admin panel (Control Center → User Rights) to export the full active user/session table to CSV. For IP geolocation without a SIEM, pipe REDCap's MySQL auth log through a free MaxMind GeoLite2 lookup script: `python3 geoip_lookup.py redcap_auth.log --asn-flag` to flag datacenter ASNs (AS14061 DigitalOcean, AS16276 OVH, AS13335 Cloudflare) and Tor exit nodes against the public Tor exit node list (https://check.torproject.org/torbulkexitlist). Force global PHP session invalidation by truncating the REDCap session table: `mysql -u redcap_user -p redcap_db -e 'TRUNCATE TABLE redcap_sessions;'`.

Evidence: Before invalidating any session, capture all volatile session state: export the REDCap MySQL `redcap_log_event` table filtered to the past 12 months (`SELECT * FROM redcap_log_event WHERE ts >= DATE_SUB(NOW(), INTERVAL 12 MONTH) INTO OUTFILE '/tmp/redcap_log_event_snapshot.csv'`), dump the `redcap_sessions` table in full, and capture active web server connections via `ss -tnp` or `netstat -ano` on the REDCap host to preserve any live attacker session state. Also capture the PHP session file directory (typically `/var/lib/php/sessions/`) before truncation — session files may contain encoded user/IP data useful for attribution. This China-nexus APT sustained access for ~12 months, so session artifacts may reflect long-lived or periodically refreshed stolen credential reuse rather than a single session.

Step 2: Detection — Review REDCap authentication logs and web server access logs for the past 12+ months. Query for: accounts with login activity outside normal business hours, logins from unexpected geographic IP ranges or anonymizing infrastructure (Tor, VPS providers, datacenter ASNs), single accounts accessing unusually large volumes of records, and repeated successful logins after prior failed attempts. Cross-reference against AU-2 (Event Logging) log categories. Enable or verify CIS 8.2 (Collect Audit Logs) compliance across the REDCap infrastructure.

NIST Phase: Detection Analysis

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

Controls: NIST AU-2 (Event Logging), NIST AU-6 (Audit Record Review, Analysis, And Reporting), CIS 8.2 (Collect Audit Logs)

Compensating: Without a SIEM, query REDCap's MySQL `redcap_log_event` table directly for the 12-month window: `SELECT user, ip, ts, page, object_type FROM redcap_log_event WHERE ts >= DATE_SUB(NOW(), INTERVAL 12 MONTH) ORDER BY user, ts;` — then pipe output into a Python Pandas script to flag: (1) logins between 22:00–06:00 local time, (2) per-account record-access counts exceeding 2 standard deviations from mean, (3) IPs not matching the institution's registered ASN(s). Supplement with Apache/Nginx access log `grep` for REDCap's `/api/` endpoint: `grep 'POST /redcap/api/' /var/log/apache2/access.log | awk '{print $1, $7, $9}' | sort | uniq -c | sort -rn` to identify high-volume API credential abuse. Cross-reference source IPs against AbuseIPDB free tier (1,000 checks/day).

Evidence: Preserve all log files before analysis modifies access timestamps: archive the full Apache/Nginx access and error logs (`/var/log/apache2/` or `/var/log/nginx/`), the REDCap MySQL `redcap_log_event` and `redcap_log_view` tables (the latter records individual record views and is critical for scoping data exfiltration volume), and any REDCap API access logs. Hash all archived files with SHA-256 before analysis (`sha256sum *.log > evidence_hashes.txt`). Because this campaign relied on stolen credentials rather than a vulnerability exploit, the primary forensic signal will be in authentication and data-access logs rather than system call or process artifacts — the `redcap_log_view` table is the definitive record of which research records the APT actor accessed over the 12-month dwell period.

Step 3: Eradication — Reset all REDCap account credentials institution-wide. Revoke and reissue API tokens. Enforce MFA for all REDCap accounts, prioritizing externally accessible instances. Review and harden REDCap configuration per CIS 4.6 (Securely Manage Enterprise Assets and Software) and apply NIST AC-3 (Access Enforcement) by auditing role-based access to confirm least privilege. Disable any dormant or unrecognized accounts per CIS 5.3 (Disable Dormant Accounts).

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST AC-3 (Access Enforcement), NIST AC-6 (Least Privilege), CIS 4.6 (Securely Manage Enterprise Assets and Software), CIS 5.3 (Disable Dormant Accounts), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts)

Compensating: Bulk-reset REDCap passwords via MySQL: `UPDATE redcap_auth SET password = MD5(CONCAT(RAND(), user_id)), temp_pwd = 1 WHERE 1;` (forces password-change prompt on next login per REDCap's temp_pwd flag). Revoke all API tokens: `UPDATE redcap_user_information SET api_token = NULL WHERE api_token IS NOT NULL;` and log affected usernames before execution. For accounts inactive >45 days, generate a disable list: `SELECT username, last_login FROM redcap_user_information WHERE last_login < DATE_SUB(NOW(), INTERVAL 45 DAY) OR last_login IS NULL;` then set `user_suspended_time = NOW()` for each. REDCap's built-in Two-Factor Authentication module (Control Center → Two-Factor Authentication) supports TOTP

via Google Authenticator at no cost.

Evidence: Before resetting credentials and revoking API tokens, capture: the full `redcap_user_information` table (contains api_token, last_login, user_suspended_time, and role assignments), the `redcap_user_rights` table (project-level permissions per user — essential for scoping which research projects were accessible to compromised accounts), and any externally configured SSO/LDAP integration configs (`/var/www/html/redcap/webtools2/ldap/` or equivalent). API tokens in REDCap are functionally equivalent to persistent credentials and may have been independently harvested by the APT; document all tokens present before revocation as these represent additional potential exfiltration vectors beyond interactive logins.

Step 4: Recovery — After credential rotation, monitor REDCap access logs daily for a minimum of 30 days for re-emergence of anomalous patterns. Validate that MFA is enforced on all accounts using CIS 6.3 (Require MFA for Externally-Exposed Applications) and CIS 6.5 (Require MFA for Administrative Access). Confirm no unauthorized administrative accounts were created during the dwell period. Cross-reference NIST AU-6 (Audit Record Review, Analysis, and Reporting) cadence requirements.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST AU-6 (Audit Record Review, Analysis, And Reporting), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 6.5 (Require MFA for Administrative Access), CIS 5.1 (Establish and Maintain an Inventory of Accounts)

Compensating: Establish a daily cron job to export and diff the REDCap `redcap_log_event` table for new anomalous logins: `0 6 * * * mysql -u redcap_ro -p redcap_db -e "SELECT user, ip, ts FROM redcap_log_event WHERE ts >= DATE_SUB(NOW(), INTERVAL 1 DAY) AND (HOUR(ts) 22)" >> /var/log/redcap_anomaly_daily.log`. Validate MFA enrollment completeness by querying `SELECT username FROM redcap_twofactor_auth WHERE twofactor_auth_code IS NULL OR twofactor_auth_enabled = 0;` and generating a weekly compliance report for the CISO. For admin account integrity, weekly diff the `redcap_user_information` table against the pre-incident user inventory snapshot captured during eradication.

Evidence: During the 30-day monitoring window, collect and retain daily snapshots of REDCap login events and any new API token issuances — this APT demonstrated a 12-month dwell capability, and re-entry via a previously unidentified compromised credential or a newly phished account is a realistic re-compromise vector. Flag any successful login originating from a new IP not seen in the post-rotation baseline, as China-nexus actors commonly rotate VPS/proxy infrastructure after a disruption. Retain all monitoring logs with SHA-256 hashes for potential regulatory disclosure purposes if PII or HIPAA-covered research data was confirmed exfiltrated.

Step 5: Post-Incident — Conduct a privileged access review against NIST AC-6 (Least Privilege) and AC-2 (Account Management) to eliminate over-provisioned research accounts. Implement a formal credential monitoring program aligned with D3-CRO (Credential Rotation) and D3-CH (Credential Hardening). Establish behavioral baselines for REDCap usage to enable anomaly detection going forward. Engage Google Threat Intelligence Group or equivalent for threat hunt services covering the full 12-month dwell window. Review third-party and federated identity integrations for similar credential exposure paths.

NIST Phase: Post Incident

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

Controls: NIST AC-2 (Account Management), NIST AC-6 (Least Privilege), CIS 6.1 (Establish an Access Granting Process), CIS 6.2 (Establish an Access Revoking Process), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: For behavioral baselining without a SIEM, script a weekly statistical profile from `redcap_log_event`: compute per-user mean and standard deviation of daily record-access counts, login hours, and source IP count over the first 60 post-incident days using a Python script with `scipy.stats` — store baselines as a JSON reference file and alert when any daily observation exceeds 3 standard deviations. For credential exposure monitoring, enroll institutional email domains in Have I Been Pwned's free domain monitoring (https://haveibeenpwned.com/DomainSearch) and configure SpyCloud or DeHashed API alerts for institutional email patterns if budget allows. For federated identity, audit Shibboleth or InCommon federation attribute release policies to confirm REDCap SP is not over-receiving identity

attributes that could be leveraged for privilege escalation.

Evidence: For the threat hunt covering the full 12-month dwell window, provide Google TIG or the engaged hunt team with: the complete archived `redcap_log_event` and `redcap_log_view` MySQL exports, web server access logs for the full dwell period, any network flow/NetFlow data retained from the institution's perimeter, DNS query logs for the REDCap server's outbound resolutions (which may reveal C2 or data staging infrastructure if the APT used the REDCap server as a pivot), and authentication logs from any federated identity provider (Shibboleth IdP, Active Directory, or LDAP) that REDCap integrated with — the credential theft initial vector may have originated upstream of REDCap itself.

Detection Guidance

Focus detection efforts on REDCap authentication and access logs spanning at least 12 months retroactively. Key behavioral indicators include: (1) successful logins from IP addresses in datacenter, VPS, or unexpected foreign ASN ranges not consistent with your researcher population; (2) accounts accessing data outside contracted research hours or from multiple geographically distant IPs within short timeframes (impossible travel); (3) large-volume record exports or API calls by accounts with no prior history of bulk access; (4) account discovery activity such as enumeration of user lists or project membership; (5) modifications to authentication configuration or account permissions not tied to a change request. At the network layer, look for unusual outbound data transfers from REDCap hosting infrastructure, particularly to cloud storage or file-sharing endpoints, consistent with T1567 (Exfiltration Over Web Service). Cross-reference NIST AU-3 (Content of Audit Records) fields: event type, timestamp, source IP, user identity, and data volume. Apply D3-LAM (Local Account Monitoring) for any REDCap service or OS-level accounts. If SIEM is deployed, build detection rules for consecutive successful logins by the same credential from non-overlapping IP subnets within a 24-hour window. No public IOC hashes, domains, or IP addresses have been released by Google or any government authority as of the configuration date; monitor CISA and Google Threat Intelligence Group advisories for IOC release.

Indicators of Compromise

Type	Value	Context	Confidence
DOMAIN	`none released`	No IOCs have been publicly released by Google Threat Intelligence Group or CISA as of the configuration date. Monitor official channels for future disclosure.	LOW

Framework Mappings

MITRE-ATTACK

- **T1074** — Data Staged
- **T1005** — Data from Local System
- **T1021** — Remote Services
- **T1087** — Account Discovery
- **T1213** — Data from Information Repositories
- **T1078.004** — Cloud Accounts

- **T1567** — Exfiltration Over Web Service
- **T1555** — Credentials from Password Stores
- **T1041** — Exfiltration Over C2 Channel
- **T1556** — Modify Authentication Process
- **T1190** — Exploit Public-Facing Application
- **T1083** — File and Directory Discovery
- **T1114** — Email Collection
- **T1078** — Valid Accounts
- **T1539** — Steal Web Session Cookie

NIST-800-53R5

- **AC-17** — Remote Access
- **AC-3** — Access Enforcement
- **CM-7** — Least Functionality
- **IA-2** — Identification and Authentication (Organizational Users)
- **CA-7** — Continuous Monitoring
- **SC-7** — Boundary Protection
- **SI-4** — System Monitoring
- **AC-6** — Least Privilege
- **IA-5** — Authenticator Management
- **SI-7** — Software, Firmware, and Information Integrity
- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SI-2** — Flaw Remediation
- **AC-2** — Account Management
- **SC-28** — Protection of Information at Rest
- **IA-8** — Identification and Authentication (Non-Organizational Users)

OWASP-TOP10-2021

- **A01:2021** — Broken Access Control
- **A07:2021** — Identification and Authentication Failures
- **A04:2021** — Insecure Design

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control
- **164.312(d)** — Person or Entity Authentication
- **164.308(a)(5)(ii)(D)** — Password Management

CIS-V8

- **6.3** — Require MFA for Externally-Exposed Applications
- **6.4** — Require MFA for Remote Network Access

- **6.5** — Require MFA for Administrative Access
- **5.2** — Use Unique Passwords
- **8.2** — Collect Audit Logs

SOC2-TSC

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

ISO-27001-2022

- **A.5.34** — Privacy and protection of personal information

NIST-CSF-2

- **DE.CM-01** — Networks and network services are monitored

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1074	Data Staged	Collection
T1005	Data from Local System	Collection
T1021	Remote Services	Lateral-Movement
T1087	Account Discovery	Discovery
T1213	Data from Information Repositories	Collection
T1078.004	Cloud Accounts	Defense-Evasion
T1567	Exfiltration Over Web Service	Exfiltration
T1555	Credentials from Password Stores	Credential-Access
T1041	Exfiltration Over C2 Channel	Exfiltration
T1556	Modify Authentication Process	Credential-Access
T1190	Exploit Public-Facing Application	Initial-Access
T1083	File and Directory Discovery	Discovery
T1114	Email Collection	Collection
T1078	Valid Accounts	Defense-Evasion
T1539	Steal Web Session Cookie	Credential-Access

Sources

Source	URL	Tier
Security News	https://www.darkreading.com/threat-intelligence/china-nexus-actor-u...	T3
Research Electronic Data Capture (REDCap) - PMC - NIH	https://pmc.ncbi.nlm.nih.gov/articles/PMC5764586/	T1
REDCap: Multiple Cross-Site Scripting (XSS) Vulnerabilities	https://www.levelblue.com/blogs/spiderlabs-blog/redcap-multiple-cro...	T3
About - REDCap	https://projectredcap.org/about/	T3
How REDCap Supports Multi-Site Clinical Research - JAF Consulting	https://jafconsulting.com/how-redcap-supports-multi-site-clinical-r...	T3

DISCLAIMER

This intelligence report is produced by Tech Jacks Solutions Security Command Center (SCC) for informational purposes only. It does not constitute professional security advice, legal counsel, or an incident response engagement. The information herein is derived from publicly available sources and AI-assisted analysis; while every effort is made to ensure accuracy, Tech Jacks Solutions makes no warranties regarding completeness or timeliness. Organizations should conduct their own validation and consult qualified security professionals before taking action based on this report. Tech Jacks Solutions is not liable for any damages resulting from the use of this information.

Generated 2026-06-16 07:14 UTC by TJS Security Command Center