INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-06 08:40 UTC

ShinyHunters Abuses Canvas Native APIs to Exfiltrate 280 Million Education Records Across 8,800+ Institutions

DATA BREACH | **HIGH** | CVSS 9.5

SCC Item ID	SCC-DBR-2026-0114
Type	Data Breach
Severity	HIGH
CVSS Base Score	9.5
Affected Products	Instructure Canvas LMS (cloud-based SaaS); 8,809+ educational institutions globally including colleges, K-12 school districts, and online education platforms
Published	2026-05-05T17:20:23
Discovery Source	Rss

Executive Summary

ShinyHunters, a known extortion group, claims to have stolen approximately 280 million student, teacher, and staff records from Instructure's Canvas LMS by abusing the platform's own data access and provisioning APIs rather than exploiting a software vulnerability. Instructure has confirmed a cybersecurity incident. The breach affects an estimated 8,800+ institutions globally, exposing PII at a scale that creates immediate downstream risk of targeted phishing, identity fraud, and social engineering against students, staff, and minors.

Technical Analysis

ShinyHunters reportedly exfiltrated data from Instructure Canvas by abusing legitimate platform capabilities: the Data Access Platform (DAP), provisioning reports, and user-facing APIs. This is a living-off-the-land technique applied to SaaS, no code-level vulnerability was required. Access was likely obtained through compromised credentials or over-permissioned API keys, allowing bulk enumeration and exfiltration at scale. No CVE has been assigned; the failure is rooted in improper access control and authorization enforcement (CWE-269: Privilege Misuse, CWE-284: Improper Access Control, CWE-285: Improper Authorization, CWE-200: Exposure of Sensitive Information). MITRE ATT&CK techniques include T1078 (Valid Accounts), T1213 (Data from Information Repositories), T1530 (Data from Cloud Storage), and T1567 (Exfiltration Over Web Service). Exposed data reportedly includes names, email addresses, institutional affiliations, and potentially academic records. As of the latest available reporting (May 2026), Instructure has not publicly disclosed a remediation

patch; the focus is on access control review and credential rotation by affected institutions. Institutions should treat all Canvas API keys and admin credentials as potentially compromised.

Action Checklist

1. Containment: Immediately audit all active Canvas API keys, OAuth tokens, and DAP credentials; revoke any that are unused, over-permissioned, or unrecognized. Disable or restrict DAP and bulk provisioning report access for non-essential accounts pending Instructure guidance.
2. Detection: Review Canvas audit logs and DAP query logs for anomalous bulk data pulls, unusual provisioning report generation, or API calls from unexpected IP addresses or service accounts. Check for high-volume GET requests against /api/v1/accounts, /api/v1/users, and provisioning report endpoints. Contact Instructure support to request a tenant-specific access log review.
3. Eradication: Rotate all Canvas admin credentials, API keys, and integration tokens. Remove or scope-limit any third-party integrations with broad Canvas data access. Apply least-privilege principles to all Canvas roles and API scopes. Follow Instructure's published incident guidance when released.
4. Recovery: Validate that no residual unauthorized access sessions remain active in Canvas admin panels. Monitor for newly created Canvas accounts or API keys post-rotation. Confirm with Instructure that your tenant's exposure scope is understood before resuming normal DAP or bulk reporting workflows.
5. Post-Incident: This incident exposes a control gap common in SaaS environments: over-permissioned API access and insufficient monitoring of bulk data export capabilities. Conduct a full SaaS access review across your portfolio. Implement CASB or SaaS security posture management (SSPM) tooling to alert on bulk data exports. Update your vendor risk assessments for Instructure and similar LMS platforms.

IR / Forensic Enrichment

Triage Priority	IMMEDIATE
Escalation Criteria	Escalate to institutional legal counsel, privacy officer, and Instructure's incident response team immediately if Canvas audit logs or Instructure tenant analysis confirms any bulk export of student PII from your institution's tenant, as FERPA breach notification obligations and applicable state student privacy laws (e.g., SOPIPA, NY Ed Law §2-d) are triggered; additionally escalate if post-rotation monitoring detects any new unauthorized API keys or admin sessions, indicating persistent access beyond the initial credential abuse.
Recovery Notes	Before resuming DAP or bulk provisioning report workflows, obtain written confirmation from Instructure via your support ticket specifying the earliest and latest timestamps of unauthorized access to your tenant and which report types or API endpoints were queried — this scopes your FERPA notification obligation. Monitor Canvas audit logs daily for a minimum of 30 days post-rotation for anomalous API key creation, unexpected admin logins, or DAP job executions initiated by service accounts, as ShinyHunters is a persistent extortion actor known to maintain access footholds. Validate that all third-party integrations (e.g., SIS sync connectors, Clever, Illuminate) have been re-authorized with scoped OAuth permissions before restoring automated provisioning workflows.

Forensic Artifacts	Canvas Admin Audit Log (JSON export from Admin > Settings > Logging): Filter for event types 'api_key_created', 'login', 'sis_batch_created', 'grade_change', and 'user_created' in the 90 days prior to incident disclosure — these are the event types that would reflect ShinyHunters' API-based bulk enumeration of /api/v1/users and /api/v1/accounts endpoints. | Canvas Data Access Platform (DAP) Query History: Full job execution log including job_id, initiated_by credential, query_type (e.g., 'provisioning_report', 'sis_export'), row_count_returned, execution_timestamp, and destination storage location — directly reflects the bulk data exfiltration mechanism attributed to this incident. | Canvas Developer Keys Inventory Snapshot (Admin > Developer Keys): Exported list of all active and inactive API tokens including creation date, associated user/service account, granted OAuth scopes, and last-activity timestamp — establishes which credentials were potentially used or compromised during the ShinyHunters campaign. | Web Proxy / Firewall Egress Logs: Outbound HTTPS connection logs to *.instructure.com and *.canvaslms.com from campus or datacenter IP ranges, filtered for HTTP response sizes exceeding 500KB from /api/v1/reports, /api/v1/users, /api/v1/accounts, and DAP endpoints — large response bodies indicate successful bulk record retrieval consistent with the 280M record exfiltration claim. | IdP / SSO Authentication Logs for Canvas Service Accounts: SAML assertion logs or OAuth token issuance logs from your identity provider (e.g., Shibboleth, Azure AD, Okta) for all accounts holding Canvas admin or DAP roles — abnormal authentication times, unusual originating IPs, or token grants to unrecognized relying parties may indicate upstream credential compromise that enabled the API abuse.

Per-Action IR Details

Containment — Immediately audit all active Canvas API keys, OAuth tokens, and DAP credentials; revoke any that are unused, over-permissioned, or unrecognized. Disable or restrict DAP and bulk provisioning report access for non-essential accounts pending Instructure guidance.

NIST Phase: Containment

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

Controls: NIST IR-4 (Incident Handling), NIST AC-2 (Account Management), NIST AC-17 (Remote Access), CIS 5.3 (Disable Dormant Accounts), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts), CIS 6.2 (Establish an Access Revoking Process)

Compensating: Export all Canvas API tokens via Admin > Developer Keys panel and cross-reference against your last known-good integration inventory (spreadsheet acceptable). For each unrecognized key, note the associated user account, creation date, and last-used timestamp shown in the Canvas admin UI. Revoke via Admin > Developer Keys > toggle off, then document in a change log. For DAP, navigate to Admin > Data Access Platform and disable any active credentials not tied to an approved integration. No SIEM required — the Canvas admin console surfaces this natively.

Evidence: Before revoking, screenshot or CSV-export the full Canvas Developer Keys list (Admin > Developer Keys) capturing key name, created-by user, creation date, and last-activity timestamp. Export Canvas Data Access Platform credential list including associated account names and last query timestamps. Preserve the Canvas audit log export (Admin > Settings > Logging) for the 90 days prior to the incident date, specifically filtering on 'api_key_created', 'login', and 'grade_change' event types. This establishes the credential inventory baseline before you alter it.

Detection — Review Canvas audit logs and DAP query logs for anomalous bulk data pulls, unusual provisioning report generation, or API calls from unexpected IP addresses or service accounts. Check for high-volume GET requests against /api/v1/accounts, /api/v1/users, and provisioning report endpoints. Contact Instructure support to request a tenant-specific access log review.

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-12 (Audit Record Generation), NIST SI-4 (System Monitoring), CIS 8.2 (Collect Audit Logs)

Compensating: Use the Canvas Data Services export or Admin > Logging UI to pull raw audit events. Pipe the exported JSON/CSV into a local Python script or jq query to count GET requests per service account per hour against /api/v1/accounts, /api/v1/users, /api/v1/accounts//users, and /api/v1/reports endpoints. Flag any account exceeding 1,000 user-record retrievals in a single session or any IP not matching your institution's known egress ranges. Example jq filter: `jq '[.[] | select(.url | test("/api/v1/users|/api/v1/accounts|/api/v1/reports"))] | group_by(.pseudonym_id) | map({account: .[0].pseudonym_id, count: length})' canvas_audit.json`. For DAP specifically, query the DAP console for jobs executed in the 90 days prior, noting job type, row counts returned, and destination credentials.

Evidence: Capture the Canvas audit log export covering 90 days pre-incident, preserving raw JSON with fields: user_id, pseudonym_id, url, http_method, created_at, remote_ip, and request_id. Retrieve DAP query history logs showing job_id, initiated_by, query_type, row_count_returned, and execution_timestamp — request this from Instructure support if not self-serviceable. Collect any web proxy or firewall logs showing outbound HTTPS connections from campus networks to instructure.com API endpoints, focusing on response body sizes that indicate bulk record retrieval (responses >1MB from /api/v1/users or /api/v1/reports). If your institution uses SSO/SAML with Canvas, pull IdP authentication logs for the service accounts associated with the suspicious API keys to determine if credentials were compromised upstream.

Eradication — Rotate all Canvas admin credentials, API keys, and integration tokens. Remove or scope-limit any third-party integrations with broad Canvas data access. Apply least-privilege principles to all Canvas roles and API scopes. Follow Instructure's published incident guidance when released.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST IR-4 (Incident Handling), NIST AC-2 (Account Management), NIST IA-5 (Authenticator Management), NIST CM-7 (Least Functionality), NIST SI-2 (Flaw Remediation), CIS 4.6 (Securely Manage Enterprise Assets and Software), CIS 5.1 (Establish and Maintain an Inventory of Accounts), CIS 5.2 (Use Unique Passwords)

Compensating: Systematically rotate credentials in this order to avoid locking out critical workflows: (1) Canvas root admin account password via institution SSO provider or direct Canvas admin reset; (2) all Developer Keys via Admin > Developer Keys > delete and re-issue with scoped permissions; (3) LTI integration secrets via Admin > Settings > Apps — document each integration's data scope before removal; (4) DAP credentials via the DAP console credential reset function. For third-party integrations (e.g., Clever, Illuminate, PowerSchool sync), review each integration's requested OAuth scopes and downscope or remove write/read-all-users permissions. Document every change in a credential rotation log with timestamp, actor, and old-key fingerprint for forensic continuity. No enterprise tooling required — Canvas admin UI supports all of these actions natively.

Evidence: Before rotation, preserve a complete export of: all existing Developer Key names, associated user accounts, OAuth scope strings granted, and creation timestamps. For each third-party LTI/OAuth integration, record the granted scope list from Admin > Settings > Apps. Capture the Canvas roles and permissions matrix (Admin > Permissions) showing which custom roles had 'Manage SIS data', 'Manage developer keys', or 'View all grades' permissions — these represent the attack surface ShinyHunters likely exploited. This evidence baseline confirms the pre-incident permission state for regulatory notification and supports scope-of-exposure analysis.

Recovery — Validate that no residual unauthorized access sessions remain active in Canvas admin panels. Monitor for newly created Canvas accounts or API keys post-rotation. Confirm with Instructure that your tenant's exposure scope is understood before resuming normal DAP or bulk reporting workflows.

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-6 (Audit Record Review, Analysis, and Reporting), NIST CP-4 (Contingency Plan Testing), CIS 5.1 (Establish and Maintain an Inventory of Accounts), CIS 6.1 (Establish an Access Granting Process)

Compensating: Enable Canvas's built-in 'New Developer Key Created' and 'Admin Login' audit events and configure email alerts via Admin > Notifications if your Canvas instance supports it. As a manual compensating control, assign

one analyst to run a daily Canvas audit log query for 30 days post-rotation filtering on event types: 'api_key_created', 'login' from external IPs, and 'sis_batch_created' (which flags new provisioning report jobs). Use a simple Python cron job or scheduled task to pull the Canvas API endpoint GET /api/v1/audit/authentication/logins daily and diff against your known-good account list. Before re-enabling DAP, obtain written confirmation from Instructure (via your support ticket) stating your tenant's last confirmed unauthorized access timestamp and which data tables or report types were queried.

Evidence: Preserve post-rotation Canvas audit logs daily for 30 days as your recovery monitoring baseline — specifically log entries for account/user creation events, new Developer Key issuance, and any DAP job execution. Document the Instructure support case number and any tenant-specific exposure confirmation they provide in writing, as this will be required for FERPA breach notification scoping. If any new unauthorized API keys or admin accounts are detected post-rotation, immediately re-escalate to containment phase and treat as evidence of persistent access.

Post-Incident — This incident exposes a control gap common in SaaS environments: over-permissioned API access and insufficient monitoring of bulk data export capabilities. Conduct a full SaaS access review across your portfolio. Implement CASB or SaaS security posture management (SSPM) tooling to alert on bulk data exports. Update your vendor risk assessments for Instructure and similar LMS platforms.

NIST Phase: Post Incident

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

Controls: NIST IR-4 (Incident Handling), NIST IR-8 (Incident Response Plan), NIST RA-3 (Risk Assessment), NIST SI-5 (Security Alerts, Advisories, and Directives), NIST CA-7 (Continuous Monitoring), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: For teams without CASB/SSPM budget: build a lightweight SaaS API monitoring framework using free tools. Deploy a scheduled Python script using each SaaS platform's audit API (Canvas, Google Workspace, Microsoft 365, Zoom) to pull daily access logs and alert on bulk export events — define 'bulk' as >500 records in a single API response or report job. Use osquery on any on-premise systems that hold SaaS integration credentials to detect credential file access: `SELECT * FROM file_events WHERE path LIKE '%canvas%' OR path LIKE '%lms%'`. For vendor risk reassessment, add the following questions to your Instructure and LMS vendor reviews: (1) Does the platform support API rate limiting per token? (2) Are DAP/bulk export capabilities separately permissioned from standard API access? (3) Does the vendor provide tenant-level SIEM-exportable audit logs? Document lessons-learned findings referencing the ShinyHunters DAP abuse vector specifically, and update your SaaS onboarding checklist to require least-privilege API scoping and bulk export monitoring as baseline controls.

Evidence: Compile a post-incident artifact package including: the full credential rotation log, Instructure support case correspondence confirming your tenant's exposure scope, the pre-incident Canvas permissions matrix, the DAP query history, and the audit log exports from the 90-day investigation window. This package supports FERPA breach notification filings (required if student PII was confirmed exposed), state-level student privacy law obligations (e.g., SOPIPA, PPRA), and any institutional cyber insurance claim. Retain all artifacts per your institution's records retention policy — minimum 3 years recommended for breach-related documentation.

Detection Guidance

Review Canvas audit log exports for the following behavioral indicators: bulk or sequential API calls to user enumeration endpoints (/api/v1/accounts/{id}/users, /api/v1/courses/{id}/enrollments); repeated provisioning report generation outside normal scheduling windows; DAP queries returning unusually large result sets; API authentication events from unfamiliar IP ranges or service account names not associated with known integrations. If your institution uses a SIEM, ingest Canvas audit logs and alert on: single API key generating more than a defined threshold of user-record queries within a short window; provisioning report downloads from IPs not matching your institution's known admin ranges; API activity at unusual hours relative to your institution's operational baseline. No public IOCs (IPs, domains, hashes) have been confirmed in available reporting as of May 2026.

Framework Mappings

MITRE-ATTACK

- **T1213** — Data from Information Repositories
- **T1567** — Exfiltration Over Web Service
- **T1530** — Data from Cloud Storage
- **T1078** — Valid Accounts
- **T1619** — Cloud Storage Object Discovery
- **T1190** — Exploit Public-Facing Application

NIST-800-53R5

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

OWASP-TOP10-2021

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

CIS-V8

- **5.4** — Restrict Administrator Privileges to Dedicated Administrator Accounts
- **6.8** — Define and Maintain Role-Based Access Control
- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **14.2** — Train Workforce Members to Recognize Social Engineering Attacks

SOC2-TSC

- **CC6.1** — The entity implements logical access security software, infrastructure, and architectures over protected information assets
- **CC7.4** — Responds to identified security incidents

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control

- **164.308(a)(5)(i)** — Security Awareness and Training
- **164.308(a)(6)(ii)** — Response and Reporting

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.34** — Privacy and protection of personal information
- **A.5.23** — Information security for use of cloud services

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1213	Data from Information Repositories	Collection
T1567	Exfiltration Over Web Service	Exfiltration
T1530	Data from Cloud Storage	Collection
T1078	Valid Accounts	Defense-Evasion
T1619	Cloud Storage Object Discovery	Discovery
T1190	Exploit Public-Facing Application	Initial-Access

Sources

Source	URL	Tier
Security News	https://www.bleepingcomputer.com/news/security/instructure-hacker-c...	T3
Instructure confirms cybersecurity incident - K-12 Dive	https://www.k12dive.com/news/instructure-confirms-cybersecurity-inc...	T3
Hackers steal students' data during breach at education tech giant ...	https://techcrunch.com/2026/05/05/hackers-steal-students-data-durin...	T2
ShinyHunters: Instructure breach affects 275 million teachers and ...	https://mashable.com/article/instructure-canvas-edtech-data-breach-...	T3
Instructure (Canvas LMS) Data Breach Investigation	https://chimicles.com/instructure-canvas-data-breach-investigation/	T3

DISCLAIMER

Generated 2026-05-06 08:40 UTC by TJS Security Command Center