INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-06 09:03 UTC

Educational tech firm Instructure data breach may have impacted 9,000 schools

DATA BREACH | HIGH | CVSS 7.5

SCC Item ID	SCC-DBR-2026-0116
Type	Data Breach
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Instructure Canvas Learning Platform, user personal data; scope potentially includes students, educators, and administrators across approximately 9,000 schools
Published	1 day ago
Discovery Source	Serper

Executive Summary

Instructure, the company behind the Canvas learning management system, confirmed unauthorized access to user personal data affecting a platform used by approximately 9,000 schools across K-12 and higher education. A threat actor has reportedly leaked data from the breach; the full scope of exposed records and attack vector remain unconfirmed pending Instructure's investigation. Organizations dependent on Canvas face potential exposure of student, educator, and administrator personal data, with significant regulatory and reputational risk given the involvement of minors and educational records.

Technical Analysis

Instructure confirmed a cybersecurity incident resulting in unauthorized access to and exfiltration of user personal data from the Canvas learning management system. MITRE ATT&CK techniques mapped to the incident include T1078 (Valid Accounts), T1213 (Data from Information Repositories), and T1530 (Data from Cloud Storage). No CVE has been assigned; this is a breach/incident, not a discrete software vulnerability. No CWE identifiers have been confirmed. The attack vector, affected system components, specific data types compromised, and record count have not been officially confirmed by Instructure as of reporting. A threat actor has reportedly leaked data externally. No patch or vendor advisory is applicable in the traditional sense; response centers on access review, credential rotation, and monitoring. Source confidence is medium, reporting derives from SecurityWeek, K-12 Dive, and Security Affairs; no independently verified official Instructure disclosure statement was available in the source data.

Action Checklist

1. Containment, If your institution uses Canvas, immediately audit active API tokens, OAuth integrations, and third-party LTI connections for anomalous access. Suspend or revoke tokens that cannot be attributed to known, authorized use.
2. Detection, Review identity provider (IdP) and Canvas access logs for logins from unfamiliar IP addresses, unusual bulk data exports, or access to the Canvas Data 2 (Instructure Data Services) pipeline. Focus on admin and privileged accounts first.
3. Eradication, Force password resets and MFA re-enrollment for all Canvas administrator, instructor, and integration service accounts. Rotate API keys and OAuth credentials for all third-party integrations connected to your Canvas instance.
4. Recovery, Verify that no unauthorized LTI tools, webhook configurations, or data pipeline connections remain active. Monitor Canvas audit logs and IdP sign-in reports for at least 30 days post-reset for residual unauthorized access.
5. Post-Incident, Assess whether your institution's Canvas configuration enforces least-privilege data access for third-party integrations. Review data retention and export permissions for Canvas Data 2. Evaluate whether FERPA breach notification obligations apply and consult your privacy counsel.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to institutional leadership, legal counsel, and your state education agency if Canvas Data 2 pipeline logs confirm bulk export of student records (names, emails, enrollment data) during the breach window, if FERPA notification obligations are triggered, or if your institution lacks the administrative access to audit Canvas API tokens and LTI configurations without Instructure support engagement.
Recovery Notes	Post-containment recovery must include a verified inventory comparison of LTI tools and Developer Keys against a pre-incident baseline — Canvas does not natively alert on new LTI installations, making manual diffing essential to confirm no backdoor integrations remain. Monitor Canvas authentication audit logs and IdP sign-in reports daily for the first two weeks, then weekly through 30 days post-reset, specifically watching for logins from ASNs or geographies that appeared in the suspicious access window. If Instructure publishes an updated incident scope or attack vector confirmation, reassess whether the initial containment actions were sufficient or whether the credential rotation scope needs to be expanded.

Forensic Artifacts	Canvas Data 2 (Instructure Data Services) pipeline job execution history — records which credentials initiated exports, what schema tables were queried, export file sizes, and timestamps; anomalous jobs outside scheduled windows or from unrecognized credentials are the primary exfiltration indicator for this breach. | Canvas authentication audit log (API: GET /api/v1/audit/authentication/accounts/:account_id) — captures every login event with source IP, user-agent, and account role; bulk admin logins from unfamiliar IPs or automated user-agent strings during the breach window indicate credential compromise or automated scraping. | IdP sign-in logs (Azure AD, Okta, or Google Workspace) for all Canvas-linked accounts — cross-reference source IP, device fingerprint, and MFA challenge result against known institutional patterns; successful MFA bypass or logins from new devices on admin accounts indicate stolen session tokens or credential stuffing. | Canvas Developer Keys export (Admin > Developer Keys) and OAuth grant log — enumerates all active API tokens, their last-use timestamps, and the accounts they were issued to; rogue or unclaimed tokens with recent last-use timestamps during the breach window indicate an attacker maintained persistent API access. | Canvas LTI tool installation list (Admin > Settings > Apps, or API: GET /api/v1/accounts/:account_id/lti_apps) — unauthorized LTI tools can act as persistent data exfiltration channels by receiving Canvas user data on every launch event; any tool installed within 90 days of breach disclosure without a documented change request should be treated as a candidate persistence mechanism.

Per-Action IR Details

Containment — If your institution uses Canvas, immediately audit active API tokens, OAuth integrations, and third-party LTI connections for anomalous access. Suspend or revoke tokens that cannot be attributed to known, authorized use.

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 5.1 (Establish and Maintain an Inventory of Accounts), CIS 6.2 (Establish an Access Revoking Process)

Compensating: Export the Canvas Admin API token list via: GET /api/v1/users/:user_id/tokens using an admin token and pipe through jq to extract created_at, last_used_at, and purpose fields. Flag any token with last_used_at within the breach window that lacks a documented purpose. For OAuth apps, navigate to Admin > Developer Keys in the Canvas UI and cross-reference each active key against your IT asset inventory. Revoke unmatched keys immediately. Document all revoked tokens with timestamps in a shared incident log (Google Sheet or wiki page suffices for a 2-person team).

Evidence: Before revoking, export and preserve the full Canvas Developer Keys list (Admin > Developer Keys > export or API: GET /api/v1/dev_keys), the OAuth token grant log, and the LTI tool installation list (Admin > Settings > Apps). Capture Canvas audit log entries under Admin > Logging showing token creation and last-use timestamps. These records establish the pre-remediation access baseline and may reveal the attacker's persistence mechanism if a rogue OAuth app or LTI tool was installed during the breach.

Detection — Review identity provider (IdP) and Canvas access logs for logins from unfamiliar IP addresses, unusual bulk data exports, or access to the Canvas Data 2 (Instructure Data Services) pipeline. Focus on admin and privileged accounts first.

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: For IdP logs (Azure AD, Okta, Google Workspace): export sign-in logs for all Canvas-linked accounts and filter for source IPs not matching your institution's known IP ranges. Azure AD CLI: az monitor activity-log list --filter

"eventTimestamp ge 2025-01-01" | grep canvas. For Canvas-native logs, use the Canvas Admin API: GET /api/v1/audit/authentication/accounts/:account_id to pull authentication events. For Canvas Data 2 (Instructure Data Services) pipeline access, review the Instructure Data Services portal access logs and check for API credential usage outside of scheduled sync windows. Cross-reference any bulk export events against known ETL job schedules documented by your data team.

Evidence: Capture IdP sign-in logs showing authentication events for Canvas admin and privileged accounts covering at least 90 days prior to breach disclosure. Preserve Canvas authentication audit logs (GET /api/v1/audit/authentication) with full IP, user-agent, and timestamp fields. Extract Canvas Data 2 pipeline job execution history from the Instructure Data Services portal, focusing on export jobs that ran outside normal scheduled windows or from unexpected initiating credentials. Anomalous bulk exports from the Canvas Data 2 schema tables (users, enrollments, pseudonyms) are the primary forensic indicator of data exfiltration in this breach.

Eradication — Force password resets and MFA re-enrollment for all Canvas administrator, instructor, and integration service accounts. Rotate API keys and OAuth credentials for all third-party integrations connected to your Canvas instance.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST IR-4 (Incident Handling), NIST IA-5 (Authenticator Management), NIST AC-2 (Account Management), NIST SI-2 (Flaw Remediation), CIS 5.2 (Use Unique Passwords), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 6.5 (Require MFA for Administrative Access)

Compensating: Use the Canvas Admin API to enumerate all admin-role accounts: GET /api/v1/accounts/:account_id/admins and trigger password resets via your IdP's bulk reset feature (Okta: okta-cli users reset-password --all-matching-filter; Azure AD: PowerShell Get-AzureADUser | where {$_.assignedRoles -match 'canvas'} | Set-AzureADUserPassword). For service accounts driving LTI or SIS integrations, regenerate API keys directly in Canvas Admin > Developer Keys and update credentials in the downstream integration config immediately after rotation. Document each rotated credential with the old key hash (never the value), rotation timestamp, and responsible party in your incident log.

Evidence: Before forcing resets, capture a snapshot of all current Canvas admin account last-login timestamps and MFA enrollment status via your IdP's admin console. Export the Developer Keys list one final time to record which keys existed pre-eradication. Preserve any Canvas audit log entries showing credential use during the suspected breach window — these records may be overwritten or aged out and are needed to establish the full scope of access before you invalidate the attacker's foothold.

Recovery — Verify that no unauthorized LTI tools, webhook configurations, or data pipeline connections remain active. Monitor Canvas audit logs and IdP sign-in reports for at least 30 days post-reset for residual unauthorized access.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST IR-4 (Incident Handling), NIST AU-6 (Audit Record Review, Analysis, and Reporting), NIST CM-7 (Least Functionality), NIST SI-7 (Software, Firmware, and Information Integrity), CIS 2.3 (Address Unauthorized Software), CIS 4.6 (Securely Manage Enterprise Assets and Software)

Compensating: Generate a verified baseline of authorized LTI tools using Canvas API: GET /api/v1/accounts/:account_id/lti_apps and diff against your pre-incident inventory documented during containment. For webhooks, review Canvas Admin > Settings > Notification Preferences and any configured Data Services subscriptions in the Instructure Data Services portal. Set a recurring weekly cron job or calendar reminder for a 2-person team to pull and review Canvas authentication audit logs (GET /api/v1/audit/authentication) and IdP sign-in anomaly reports for the full 30-day window. Flag any login from a previously unseen ASN or country code as a residual access indicator.

Evidence: During recovery verification, capture a post-remediation snapshot of the LTI app list, Developer Keys, and Data Services pipeline credentials as a clean baseline. Preserve IdP conditional access or sign-in logs daily during the 30-day monitoring period — if a threat actor retained a secondary credential or session token that survived the reset, the first re-access attempt will appear here before it appears anywhere else.

Post-Incident — Assess whether your institution's Canvas configuration enforces least-privilege data access for third-party integrations. Review data retention and export permissions for Canvas Data 2. Evaluate whether FERPA breach notification obligations apply and consult your privacy counsel.

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 AU-11 (Audit Record Retention), NIST RA-2 (Security Categorization), NIST SI-12 (Information Management and Retention), CIS 3.2 (Establish and Maintain a Data Inventory), CIS 3.4 (Enforce Data Retention), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: Conduct a Canvas Data 2 schema review: document which tables (users, enrollments, pseudonyms, access_tokens) your institution's pipeline exports and map each to the minimum data set required for your use case. Use the Instructure Data Services portal to audit which credentials have export access and remove any that are not tied to an active, documented business need. For FERPA assessment, compile a data inventory of record types exposed (student names, email addresses, enrollment data, login credentials) and cross-reference against your institution's FERPA notification policy — notification to parents/eligible students is required when education records are disclosed without consent, and the 34 CFR Part 99 timeline begins from the date your institution confirms exposure, not Instructure's disclosure date. This step requires legal review — flag for privacy counsel immediately.

Evidence: Assemble the full incident artifact package for lessons-learned review: the pre- and post-remediation Canvas Developer Keys exports, IdP sign-in anomaly reports from the detection window, Canvas Data 2 pipeline job history, the LTI tool baseline comparison, and all credential rotation records. These artifacts collectively document the scope of exposure, the response timeline, and the authorization state before and after remediation — all of which are required inputs for a FERPA breach notification determination and any future regulatory inquiry.

Detection Guidance

Review Canvas audit logs for bulk data access events, especially involving the Canvas Data 2 pipeline or gradebook/roster exports at scale. In your IdP (e.g., Okta, Azure AD, Shibboleth), query for Canvas logins from new geolocations, impossible travel events, or service accounts authenticating outside business hours. Flag any API calls using tokens that do not map to known integrations. If your institution uses a SIEM, alert on Canvas admin actions (role changes, data exports, integration additions) performed by accounts not in your authorized admin group. No specific IOC hashes, IPs, or domains have been confirmed in available reporting; detection must rely on behavioral and access anomalies rather than known-bad indicators at this time.

Framework Mappings

MITRE-ATTACK

- **T1078** — Valid Accounts
- **T1213** — Data from Information Repositories
- **T1530** — Data from Cloud Storage

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.34** — Privacy and protection of personal information

HIPAA-SECURITY

- **164.308(a)(6)(ii)** — Response and Reporting

SOC2-TSC

- **CC7.4** — Responds to identified security incidents
- **CC6.3** — Authorizes, modifies, or removes access

NIST-CSF-2

- **RS.CO-03** — Recovery activities and progress communicated

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1078	Valid Accounts	Defense-Evasion
T1213	Data from Information Repositories	Collection
T1530	Data from Cloud Storage	Collection

Sources

Source	URL	Tier
	https://securityaffairs.com/191686/cyber-crime/educational-tech-fir...	T3
Edtech Firm Instructure Discloses Data Breach Amid Hacker Leak ...	https://www.securityweek.com/edtech-firm-instructure-discloses-data...	T3
Educational tech firm Instructure data breach may have impacted ...	https://x.com/TheCyberSecHub/status/2051574429487185967	T3
Educational tech firm Instructure data breach may have impacted ...	https://cdotimes.com/2026/05/05/educational-tech-firm-instructure-d...	T3
Instructure confirms cybersecurity incident - K-12 Dive	https://www.k12dive.com/news/instructure-confirms-cybersecurity-inc...	T3

Generated 2026-05-06 09:03 UTC by TJS Security Command Center