INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-19 15:56 UTC

CB Financial Services / Community Bank: Unauthorized AI Application Exposes Non-Public Customer Data

DATA BREACH | **HIGH** | CVSS 7.5

SCC Item ID	SCC-DBR-2026-0134
Type	Data Breach
Severity	HIGH
CVSS Base Score	7.5
Affected Products	CB Financial Services, Inc. / Community Bank (wholly-owned subsidiary)
Published	2026-05-11
Discovery Source	Sec 8K

Executive Summary

On May 5, 2026, Community Bank (a wholly-owned subsidiary of CB Financial Services, Inc., NASDAQ: CBFV) discovered that an unauthorized AI application accessed and exposed non-public customer information (NPI) through an improper API call made by an internal actor. The Bank filed an SEC 8-K under Item 1.05 (Material Cybersecurity Incidents) on May 7, 2026, and engaged external cybersecurity advisors to investigate. Business risk is significant: regulatory exposure under GLBA and applicable state privacy laws, potential customer notification obligations, and reputational harm to a community banking institution where customer trust is a primary competitive differentiator.

Technical Analysis

This is an insider/operational misuse incident, not a software vulnerability exploitation. No CVE has been assigned. An unauthorized AI-based software application was used internally, and secondary sources indicate the mechanism involved an unauthorized API call that exposed non-public customer information (NPI). The Bank has not publicly confirmed the full scope of affected customers or data categories in available filing text. Applicable CWEs: CWE-284 (Improper Access Control), the AI application accessed customer data without authorization; CWE-359 (Exposure of Private Personal Information to an Unauthorized Actor), NPI was exposed through the unauthorized application. Relevant MITRE ATT&CK techniques: T1213 (Data from Information Repositories), unauthorized access to customer data stores; T1078 (Valid Accounts), the actor likely used legitimate credentials to access systems; T1530 (Data from Cloud Storage), possible if the AI tool accessed cloud-hosted data repositories. No patch exists; remediation is procedural and architectural. The Bank states it

secured the affected information and engaged external advisors; investigation is ongoing as of filing date.

Action Checklist

1. Step 1: Containment. Immediately audit all AI/ML tools and third-party applications with access to customer data repositories. Revoke API tokens and access credentials for any application not formally approved through your AI governance or technology risk management process. Isolate any systems identified as having been accessed by unauthorized tooling.
2. Step 2: Detection. Query API gateway logs, SIEM, and DLP systems for anomalous data access patterns tied to unrecognized application identifiers or service accounts. Look for high-volume read operations against customer data tables or NPI repositories from non-standard applications, unusual OAuth grant events, and API calls originating from tools not in your approved software inventory. Review IAM audit logs for T1078 (Valid Accounts) indicators, legitimate credentials used in unexpected contexts.
3. Step 3: Eradication. Remove all unauthorized AI applications and revoke associated API keys, OAuth tokens, and service account credentials. Enforce an approved software inventory policy (application allowlisting) for any tool that can access customer data. Require formal security review and business justification before any AI tool is granted data access.
4. Step 4: Recovery. Validate that NPI repositories are accessible only by approved, inventoried applications. Confirm DLP controls are actively monitoring for NPI exfiltration. Re-audit IAM permissions on customer data stores and apply least-privilege principles. Monitor API gateway and data access logs for 30 days post-remediation for residual anomalous activity.
5. Step 5: Post-Incident. Conduct a formal AI/ML tool governance review. Implement a shadow AI detection program (or expand existing program) to identify unsanctioned tools. Update acceptable use and technology risk policies to explicitly address unsanctioned AI tools. Review GLBA Safeguards Rule compliance posture, particularly §314.4(b) risk assessment and §314.4(f) service provider oversight requirements. Develop or refine an AI-specific incident response playbook.

IR / Forensic Enrichment

Triage Priority	IMMEDIATE
Escalation Criteria	Escalate to external cybersecurity counsel and GLBA-required regulatory notification if forensic analysis confirms NPI fields (SSN, account numbers, or financial transaction data) were exfiltrated or accessible beyond the bank's network boundary, if the internal actor held privileged access to additional customer data systems, or if the 30-day SEC 8-K follow-up window requires a material update on containment status — CB Financial Services has already filed under Item 1.05, triggering ongoing disclosure obligations.
Recovery Notes	Validate recovery by confirming zero traffic from the revoked OAuth client_id in API gateway logs for a minimum of 72 hours post-revocation, and by running a full permission audit against all NPI data stores to confirm only inventoried service accounts retain access. Given the GLBA and SEC 8-K disclosure implications, monitor API gateway and IAM audit logs for the full 30-day post-remediation window and retain all monitoring outputs as regulatory evidence. Any recurrence of unauthorized data access patterns during the monitoring window should immediately trigger re-escalation to external cybersecurity advisors and legal counsel to assess updated breach notification obligations.

Forensic Artifacts	API gateway access logs (AWS API Gateway, Kong, Apigee, or nginx): filter on the unauthorized AI application's client_id and OAuth token across the full log retention window — these logs will show every NPI endpoint called, the volume of records returned, and the precise time window of unauthorized access. | OAuth authorization server logs (Azure AD, Okta, or on-prem identity provider): the grant event log for the token issued to the unauthorized AI application will contain the internal actor's user identity (`sub` claim), the scopes granted (confirming what data the tool was authorized to access), and the client registration details — critical for establishing whether the internal actor intentionally provisioned access. | Database query audit logs (PostgreSQL `pg_audit`, MySQL general query log, or Oracle Unified Auditing): SELECT statements executed by the unauthorized AI application's service account against customer tables will define the exact data scope — export with timestamps, query text, and row counts to support GLBA breach notification scope assessment. | Internal actor's workstation artifacts: browser history and locally cached credentials for the unauthorized AI SaaS platform (check Chrome/Firefox profile directories at `%LOCALAPPDATA%\Google\Chrome\User Data\Default\History` or `~/.mozilla/firefox/*.default/places.sqlite`), plus any downloaded NPI exports or temporary files created by the AI tool (check `%TEMP%`, `~/Downloads/`, and recycle bin). | DLP system alert history and policy match logs: any NPI pattern matches (SSN regex, account number patterns) triggered during the unauthorized access window — these logs define whether data was merely accessed in-place or transmitted to an external endpoint, which is the determinative factor for GLBA breach notification obligations and the SEC 8-K Item 1.05 materiality assessment.

Per-Action IR Details

Step 1: Containment — Immediately audit all AI/ML tools and third-party applications with access to customer data repositories. Revoke API tokens and access credentials for any application not formally approved through your AI governance or technology risk management process. Isolate any systems identified as having been accessed by unauthorized tooling.

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 all active OAuth tokens and API keys from your identity provider (e.g., run `az ad app list --all` for Azure AD or `gcloud auth list` for GCP) and cross-reference against a manually maintained approved-tools register. For on-prem environments without IAM tooling, enumerate service accounts via PowerShell: `Get-ADServiceAccount -Filter * | Select-Object Name, Enabled, LastLogonDate`. Immediately disable any account or token created within 90 days that lacks a documented business justification. A 2-person team can divide ownership: one person handles token revocation, the other documents the audit trail.

Evidence: Before revoking any token or credential, capture a forensic snapshot of: (1) API gateway access logs showing the unauthorized AI application's client_id and all endpoints it called against the NPI data repository — pay particular attention to high-volume GET requests to customer record endpoints; (2) OAuth authorization server logs showing the grant event that issued the token to the unauthorized application, including grant_type, scope, and the internal actor's user identity; (3) IAM audit logs (e.g., AWS CloudTrail `GetSecretValue`, `AssumeRole`, or Azure AD Sign-in logs) timestamped around the May 5, 2026 discovery date; (4) a full export of the API key/token metadata (creation date, last-used timestamp, associated scopes) before revocation so the timeline of unauthorized access can be reconstructed.

Step 2: Detection — Query API gateway logs, SIEM, and DLP systems for anomalous data access patterns tied to unrecognized application identifiers or service accounts. Look for high-volume read operations against customer data tables or NPI repositories from non-standard applications, unusual OAuth grant events, and

API calls originating from tools not in your approved software inventory. Review IAM audit logs for T1078 (Valid Accounts) indicators — legitimate credentials used in unexpected contexts.

NIST Phase: Detection Analysis

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

Controls: NIST SI-4 (System Monitoring), NIST AU-6 (Audit Record Review, Analysis, and Reporting), NIST AU-2 (Event Logging), NIST IR-5 (Incident Monitoring), CIS 8.2 (Collect Audit Logs), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: Without a SIEM, parse API gateway logs (Apache/nginx access logs or AWS API Gateway CloudWatch logs) using command-line tooling: `awk '{print $1, $7}' access.log | sort | uniq -c | sort -rn` to surface high-frequency caller IPs and URI paths hitting NPI endpoints. For OAuth anomaly detection, extract authorization server logs and run: `grep -E 'grant_type|client_id' auth.log | sort | uniq -c` to identify non-standard clients. Use osquery on endpoints where the unauthorized AI tool may have been installed: `SELECT name, path, start_time FROM processes WHERE name LIKE '%ai%' OR name LIKE '%llm%'`. Map findings to MITRE ATT&CK T1078 (Valid Accounts) and document the earliest observed timestamp of the unauthorized client_id in gateway logs to establish the breach window.

Evidence: Capture before scoping is complete: (1) API gateway access logs filtered for the unauthorized application's client_id across the full 90-day retention window — look for URI patterns such as `/api/customers/`, `/api/accounts/`, or any endpoint returning NPI fields (SSN, account number, balance); (2) DLP system alerts or policy match logs for NPI keywords (e.g., regex matches on SSN patterns `\d{3}-\d{2}-\d{4}`) triggered by the unauthorized application's user-agent or source IP; (3) OAuth server logs for the specific grant event that authorized the AI tool — extract the `sub` claim to identify the internal actor's identity; (4) database query logs (e.g., PostgreSQL `pg_stat_activity` or MySQL general query log) showing SELECT statements against customer tables executed under the AI application's service account; (5) any endpoint detection telemetry or browser history on the internal actor's workstation showing installation or authentication to the unauthorized AI application.

Step 3: Eradication — Remove all unauthorized AI applications and revoke associated API keys, OAuth tokens, and service account credentials. Enforce an approved software inventory policy (application allowlisting) for any tool that can access customer data. Require formal security review and business justification before any AI tool is granted data access.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST SI-2 (Flaw Remediation), NIST CM-7 (Least Functionality), NIST CM-8 (System Component Inventory), NIST AC-2 (Account Management), CIS 2.1 (Establish and Maintain a Software Inventory), CIS 2.2 (Ensure Authorized Software is Currently Supported), CIS 2.3 (Address Unauthorized Software)

Compensating: Implement application allowlisting for data-tier access using built-in database controls: in PostgreSQL, revoke the service account's privileges (`REVOKE ALL PRIVILEGES ON ALL TABLES IN SCHEMA public FROM ;`) and audit remaining grants (`SELECT grantee, privilege_type FROM information_privileges WHERE table_schema='public'`). On Windows endpoints, use AppLocker or WDAC policies to block execution of unauthorized AI client binaries by publisher or file hash. For API gateways, enforce an explicit allowlist of approved client_id values and configure a default-deny policy for unregistered OAuth clients. Document each removed application in your change management log with the revocation timestamp to support GLBA incident documentation requirements.

Evidence: Before removing the unauthorized AI application: (1) preserve a forensic image or file hash inventory of the application binary and any local configuration files (e.g., `.env` files, `config.json`, or credential stores) that may contain embedded API keys or NPI samples; (2) capture a memory snapshot or process dump if the application is still running — it may hold decrypted NPI in memory or active session tokens; (3) collect any local logs generated by the AI application itself (check `%APPDATA%`, `~/.config/`, or application-specific log directories) for evidence of what data it queried and exfiltrated; (4) preserve a copy of the service account's privilege grants in the database and IAM system before revocation to document the scope of access the tool had.

Step 4: Recovery — Validate that NPI repositories are accessible only by approved, inventoried applications. Confirm DLP controls are actively monitoring for NPI exfiltration. Re-audit IAM permissions on customer data

stores and apply least-privilege principles. Monitor API gateway and data access logs for 30 days post-remediation for residual anomalous activity.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST IR-4 (Incident Handling), NIST AC-6 (Least Privilege), NIST AU-12 (Audit Record Generation), NIST SI-7 (Software, Firmware, and Information Integrity), CIS 3.3 (Configure Data Access Control Lists), CIS 6.1 (Establish an Access Granting Process), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: Without an enterprise DLP platform, implement NPI monitoring using a combination of: (1) database-level audit triggers in PostgreSQL (`CREATE RULE` or `pg_audit` extension) to log every SELECT against customer tables, alerting on queries returning more than 100 rows; (2) a YARA rule deployed via ClamAV on egress mail/file servers scanning for SSN and account number patterns; (3) a weekly cron job running `netstat -an | grep ESTABLISHED` on database hosts to detect unexpected outbound connections. For 30-day post-remediation monitoring, configure API gateway alerts on any client_id not in the approved whitelist using native gateway alerting (AWS API Gateway → CloudWatch metric filter on `4xx` with unknown client_id, or nginx `map` directive logging non-approved User-Agents).

Evidence: During recovery validation, capture and retain: (1) a timestamped export of current IAM permission sets on all NPI data stores, signed and stored as the post-remediation baseline for GLBA audit evidence; (2) API gateway access logs for the first 72 hours post-remediation confirming zero traffic from the revoked client_id or associated IP ranges — this constitutes proof of successful containment for the SEC 8-K follow-up; (3) DLP system confirmation logs showing NPI classification policies are active and generating alerts on test data; (4) database audit logs confirming only approved service accounts are executing queries against customer tables in the recovery window.

Step 5: Post-Incident — Conduct a formal AI/ML tool governance review. Establish or strengthen a shadow AI detection program. Update your acceptable use and technology risk policies to explicitly address unsanctioned AI tools. Review GLBA Safeguards Rule compliance posture, particularly §314.4(b) risk assessment and §314.4(f) service provider oversight requirements. Develop or refine an AI-specific incident response playbook.

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 SA-9 (External System Services), NIST PM-9 (Risk Management Strategy), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 2.1 (Establish and Maintain a Software Inventory)

Compensating: For shadow AI detection without enterprise tooling, implement a monthly DNS query log review on your recursive resolver to identify domains associated with AI services (e.g., `api.openai.com`, `api.anthropic.com`, `generativelanguage.googleapis.com`) — use `grep -E 'openai|anthropic|cohere|huggingface' dns_query.log | awk '{print $5}' | sort | uniq -c | sort -rn`. Deploy a Sigma rule against authentication logs to detect new OAuth application registrations by non-IT personnel. For GLBA §314.4(f) service provider oversight, create a quarterly spreadsheet-based inventory of all third-party tools with data access, capturing vendor name, data categories accessed, contractual data handling terms, and last security review date — sufficient for examiner review in the absence of a GRC platform.

Evidence: For the lessons-learned report and GLBA regulatory documentation, preserve and package: (1) the complete timeline of the incident from first unauthorized API call (from gateway logs) through SEC 8-K filing on May 7, 2026, with evidence supporting each timeline entry; (2) the internal actor's access provisioning records — how the unauthorized AI tool was granted API access, what approval (or lack thereof) existed, and which policy gap enabled it; (3) a data scope assessment documenting exactly which NPI fields (name, account number, SSN, transaction history, etc.) were accessible to the unauthorized application, based on the API endpoint definitions and database schema — required for GLBA breach notification scope determination; (4) the post-remediation IAM audit baseline and DLP confirmation logs as evidence of corrective action for regulators and external cybersecurity advisors engaged by CB Financial Services.

Detection Guidance

Focus detection efforts on API gateway logs and IAM audit trails. Look for: (1) API calls to customer data endpoints (account tables, CRM records, NPI repositories) from application IDs not present in your approved software inventory; (2) service account activity outside normal business hours or with atypically high data read volumes; (3) OAuth token grants or API key issuance to unregistered applications; (4) DLP alerts on NPI data categories (SSN, account numbers, contact information) accessed in bulk without a corresponding business workflow. MITRE T1213 behavioral pattern: large-volume, short-duration read operations against internal data repositories from a non-standard client. MITRE T1078 indicator: valid employee credentials used to authenticate an unrecognized application. If your environment uses a UEBA or IDR platform, create a rule for API client diversity anomaly: any new application identifier accessing customer data for the first time should generate an alert for manual review.

Framework Mappings

MITRE-ATTACK

- **T1213** — Data from Information Repositories
- **T1078** — Valid Accounts
- **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
- **AC-3** — Access Enforcement

OWASP-TOP10-2021

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

CIS-V8

- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process

SOC2-TSC

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

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control

ISO-27001-2022

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

MITRE ATT&CK Mapping

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

Sources

Source	URL	Tier
SEC EDGAR Filing Index	https://www.sec.gov/Archives/edgar/data/1605301/000160530126000021/...	T1
CB Financial Services, Inc. 8-K (Item 1.05)	https://www.sec.gov/Archives/edgar/data/1605301/000160530126000021/...	T1
CB Financial Services, Inc. Cybersecurity Incident Details	https://www.board-cybersecurity.com/incidents/tracker/cb-financial-...	T3
Community Bank Discloses Unauthorized Data Exposure via API Call	https://www.linkedin.com/posts/tim-erlin_us-bank-reports-itself-aft...	T3
[PDF] UNITED STATES SECURITIES AND EXCHANGE COMMISSION ...	https://www.classaction.org/media/community-bank-data-breach-sec-8-...	T3

Generated 2026-05-19 15:56 UTC by TJS Security Command Center