INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-23 06:38 UTC

App host Vercel says it was hacked and customer data stolen

DATA BREACH | HIGH | CVSS 7.5

SCC Item ID	SCC-DBR-2026-0099
Type	Data Breach
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Vercel platform (customer data); Context AI integration
Published	2 days ago
Discovery Source	Serper

Executive Summary

Vercel, a widely used cloud platform for web application hosting, disclosed a breach in April 2026 in which an attacker accessed and stole customer data by first compromising Context AI, a third-party AI tool integrated into Vercel's employee workflow. The attacker used credentials or session access obtained from the Context AI breach to gain entry into Vercel's systems. Vercel has described the stolen data as limited in scope, but full details on affected data types and customer counts remain incomplete, creating uncertainty for organizations that host applications or store customer data on the platform.

Technical Analysis

Attack vector: supply chain compromise via third-party AI tooling integration. The threat actor initially compromised Context AI, a tool with authorized access to Vercel's environment through employee usage. Using credentials or a hijacked session obtained from Context AI (mapped to CWE-287: Improper Authentication; CWE-522: Insufficiently Protected Credentials), the attacker pivoted into Vercel's internal systems and exfiltrated customer data. MITRE ATT&CK techniques observed: T1199 (Trusted Relationship) as initial access via the compromised vendor; T1078/T1078.004 (Valid Accounts: Cloud Accounts) for account hijacking; T1552 (Unsecured Credentials) for credential harvesting from the Context AI side; T1528 (Steal Application Access Token) as a possible session token abuse mechanism. CWE-1357 (Reliance on Insufficiently Trustworthy Component) reflects the structural risk of granting third-party AI tools access to production-adjacent systems without adequate isolation or monitoring. No CVE has been assigned. Patch status: not applicable, this is an access control and vendor trust failure, not a software vulnerability. Remediation centers on third-party access revocation, session invalidation, and identity hygiene. Severity rated High; CVSS base noted at 7.5 in source data. EPSS not applicable (no CVE). CISA KEV: not listed.

Action Checklist

1. Step 1: Containment. Immediately audit all third-party AI tool integrations (Context AI and similar) connected to your environment. Revoke OAuth tokens, API keys, and session access granted to Context AI or comparable tools. If any employee used Context AI with credentials that had access to customer data systems, treat those accounts as compromised and rotate credentials now.
2. Step 2: Detection. Review identity and access logs (IdP audit logs, SSO events, cloud provider access logs such as AWS CloudTrail, GCP Audit Logs, or Vercel's own access logs) for anomalous access from third-party tool service accounts or OAuth grant sources in the April 2026 timeframe. Look for unusual cloud account logins from unexpected locations, unexpected geographic access patterns, or access outside normal business hours. Check for unexpected token issuance or token use from non-human principals. If you are a Vercel customer, monitor Vercel's knowledge base bulletin (https://vercel.com/kb/bulletin/vercel-april-2026-security-incident) for updates on affected data types and customer notification status.
3. Step 3: Eradication. Remove or suspend Context AI integration across your organization pending vendor confirmation that their compromise is fully contained. Rotate all service account credentials and OAuth tokens associated with any third-party AI tooling that has access to employee accounts or production-adjacent systems. Enforce least-privilege scoping on any re-authorized integrations: read-only where possible, no access to customer data stores.
4. Step 4: Recovery. Validate that revoked tokens are no longer active (confirm in IdP and cloud provider consoles). Re-establish third-party tool access only after vendor attestation and with scoped, monitored permissions. Monitor customer-facing systems for signs of data misuse or downstream phishing using stolen data. If you are a Vercel customer, await Vercel's direct notification for confirmation of whether your account data was included in the exfiltration.
5. Step 5: Post-Incident. This incident exposes a specific control gap: third-party AI tools operating with broad employee-level access to systems containing customer data, without session isolation or behavioral monitoring. Conduct a full inventory of AI tool integrations and their permission scopes. Apply CIS Control 5 (Account Management) and CIS Control 12 (Network Infrastructure Management) to restrict lateral reach. Map existing vendor access against NIST SP 800-53 SA-9 (External Information System Services) controls. Require vendors with any data-path access to provide current SOC 2 Type II reports and incident notification SLAs before re-authorization.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate to legal, privacy counsel, and executive leadership immediately if Vercel's direct notification confirms your tenant's customer data (PII, email addresses, or any data subject to GDPR, CCPA, or HIPAA) was included in the exfiltration, triggering mandatory breach notification timelines (72 hours under GDPR Article 33; 30–60 days under CCPA/state equivalents depending on jurisdiction).

Recovery Notes	After token revocation and integration removal are confirmed, re-authorize third-party AI tool access only with explicitly scoped, least-privilege OAuth grants that exclude access to any data store containing customer PII — validate this by re-reviewing the grant scope list against your data classification inventory. Monitor Vercel Audit Logs and IdP sign-in logs daily for a minimum of 30 days post-containment for re-emergence of Context AI service account activity or anomalous access patterns matching the April 2026 attacker TTPs (T1078.004 cloud account abuse, T1528 token theft). If Vercel confirms customer data exfiltration that includes your tenant, extend monitoring to your customer-facing authentication systems for credential stuffing or targeted spear-phishing campaigns leveraging the stolen data for at least 90 days.
Forensic Artifacts	IdP OAuth application audit logs: full grant history for the Context AI application client_id, including authorization_code and refresh_token issuance events, granted scopes, and granting user UPNs — establishes attacker's access path from Context AI compromise into your environment | Vercel Audit Log export (Dashboard > Settings > Audit Log, CSV format): all team member actions, integration access events, environment variable reads, and deployment events for the March–April 2026 window — the primary artifact for determining whether your tenant's data was accessed via the attacker's Vercel-level access | AWS CloudTrail or GCP Audit Log events attributed to the Context AI service role or OAuth-derived session: specifically GetObject, ListBuckets, and any datastore query events during the breach window — maps what customer data the attacker could have reached through the compromised integration | Vercel project Environment Variables configuration snapshot (pre-eradication): documents which secrets and API keys were stored in Vercel projects accessible to the Context AI integration, establishing whether the attacker could have harvested additional credentials for lateral movement beyond the initial breach | Network proxy or DNS logs showing employee workstation or CI/CD runner connections to Context AI API endpoints (app.context.ai or equivalent): correlates which employees or automated pipelines were active users of Context AI during the breach window, scoping the credential rotation requirement

Per-Action IR Details

Step 1: Containment — Immediately audit all third-party AI tool integrations (Context AI and similar) connected to your environment. Revoke OAuth tokens, API keys, and session access granted to Context AI or comparable tools. If any employee used Context AI with credentials that had access to customer data systems, treat those accounts as compromised and rotate credentials now.

NIST Phase: Containment

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

Controls: NIST IR-4 (Incident Handling), NIST AC-2 (Account Management), NIST IA-4 (Identifier Management), CIS 5.1 (Establish and Maintain an Inventory of Accounts), CIS 6.2 (Establish an Access Revoking Process)

Compensating: Export OAuth grant lists from your IdP (Okta: Settings > API > Tokens; Azure AD: az ad app list --query; Google Workspace Admin Console > Security > API Controls) and pipe to a spreadsheet for manual audit. For token revocation without SIEM, use okta-cli or Azure CLI: 'az ad user revoke-sign-in-sessions --id ' for each account that authenticated to Context AI. On AWS, run 'aws iam list-access-keys --user-name ' and 'aws iam delete-access-key' for any key associated with the Context AI integration. Two-person task: one person revokes, second person verifies revocation in the IdP audit trail.

Evidence: Before revoking, export and preserve: (1) IdP audit logs showing all OAuth grant events for the Context AI application client_id — capture grant timestamps, scopes granted, and granting user UPNs; (2) Cloud provider access logs (AWS CloudTrail: filter on eventSource=signin.amazonaws.com and userAgent strings matching Context AI's API client; GCP Audit Logs: filter on protoPayload.authenticationInfo.principalEmail matching Context AI service account); (3) Vercel team access logs (Settings > Audit Log) for the April 2026 window — export the full CSV before any token revocation clears active session state; (4) Screenshot or export active OAuth applications list from your IdP before

removal to document scope of access granted.

Step 2: Detection — Review identity and access logs (IdP audit logs, SSO events, cloud provider access logs such as AWS CloudTrail, GCP Audit Logs, or Vercel's own access logs) for anomalous access from third-party tool service accounts or OAuth grant sources in the April 2026 timeframe. Look for T1078.004 indicators: cloud account logins from unexpected IPs, geographic anomalies, or access outside business hours. Check for T1528 indicators: unexpected token issuance or token use from non-human principals. If you are a Vercel customer, monitor Vercel's knowledge base bulletin (https://vercel.com/kb/bulletin/vercel-april-2026-security-incident) for updates on affected data types and customer notification status.

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: Without a SIEM, use jq or Python to parse CloudTrail JSON locally: filter on 'userIdentity.type=AssumedRole' and cross-reference against your known Context AI service role ARN. For Okta, use the System Log API with a date filter: 'GET /api/v1/logs?since=2026-03-01T00:00:00Z&until=2026-04-30T23:59:59Z&filter=client.id+eq+"<app_id>"'. For Azure AD, use: 'Get-AzureADAuditSignInLogs | Where-Object { $_.AppId -eq "<app_id>" }' in PowerShell. Flag any access events where the source IP does not match Context AI's documented egress IPs. The Sigma rule for T1528 (github.com/SigmaHQ/sigma, rule id: cloud/azure/azure_app_token_theft) can be manually applied as a grep pattern against exported sign-in logs.

Evidence: Before analysis is complete, preserve immutable copies of: (1) AWS CloudTrail logs from the Context AI service role ARN covering March–April 2026 — specifically GetObject and ListBucket calls against S3 buckets tagged with customer data classifications; (2) Vercel Audit Log CSV export (Vercel Dashboard > Settings > Audit Log) for all team members and integrations, filtered to the April 2026 breach window — look for deployments, environment variable reads, or project access from unfamiliar session tokens; (3) IdP sign-in logs showing Context AI OAuth app usage — capture the full token issuance chain (authorization_code grant → access_token → refresh_token) to establish whether the attacker used a stolen refresh token or a stolen session cookie; (4) Network flow logs or proxy logs showing outbound connections from employee workstations or CI/CD runners to Context AI endpoints during the compromised period.

Step 3: Eradication — Remove or suspend Context AI integration across your organization pending vendor confirmation that their compromise is fully contained. Rotate all service account credentials and OAuth tokens associated with any third-party AI tooling that has access to employee accounts or production-adjacent systems. Enforce least-privilege scoping on any re-authorized integrations: read-only where possible, no access to customer data stores.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST IR-4 (Incident Handling), NIST AC-6 (Least Privilege), NIST IA-5 (Authenticator Management), NIST SA-9 (External System Services), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts), CIS 6.1 (Establish an Access Granting Process)

Compensating: Generate a full list of third-party OAuth applications and service accounts with access to production systems using your IdP's API or admin console export. For each integration that touches employee accounts or Vercel-connected systems: (1) delete the OAuth application registration, not just revoke a single token — this prevents refresh token reuse; (2) rotate the associated service account password or API key and document the rotation in a change ticket; (3) audit Vercel project Environment Variables (Dashboard > Project > Settings > Environment Variables) for any stored Context AI API keys or secrets and delete them. Use a simple bash script with the Vercel CLI ('vercel env ls --token <token>') to enumerate all environment variables across projects and flag any referencing Context AI credentials.

Evidence: Before removing the Context AI integration, capture: (1) The full OAuth scope list granted to Context AI's application registration — this documents the blast radius (e.g., whether it had 'read:user', 'repo', or 'admin:org' scopes on GitHub; or 'openid profile email' plus custom scopes on your IdP); (2) Vercel integration configuration export showing which projects and environments Context AI had access to, and whether it had access to production vs. preview vs. development environment variables; (3) Any webhook configurations that Context AI registered against your Vercel projects — these can persist after OAuth revocation and represent a residual exfiltration channel (Vercel Dashboard > Project > Settings > Git > Deploy Hooks and Webhooks).

Step 4: Recovery — Validate that revoked tokens are no longer active (confirm in IdP and cloud provider consoles). Re-establish third-party tool access only after vendor attestation and with scoped, monitored permissions. Monitor customer-facing systems for signs of data misuse or downstream phishing using stolen data. If you are a Vercel customer, await Vercel's direct notification for confirmation of whether your account data was included in the exfiltration.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST IR-4 (Incident Handling), NIST AU-9 (Protection of Audit Information), NIST SI-7 (Software, Firmware, and Information Integrity), NIST SA-9 (External System Services), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: To validate token revocation without an enterprise IdP dashboard: use OAuth introspection endpoints if supported (RFC 7662 — POST to your IdP's /introspect with the previously issued token and confirm 'active: false' in the response). For AWS, confirm IAM key deletion by running 'aws iam get-access-key-last-used --access-key-id ' — a 404 or error confirms deletion. For downstream phishing monitoring tied to stolen Vercel customer data (names, emails, project names), set up Google Alerts or similar free monitoring on your organization's domain name combined with terms like 'Vercel' or your specific project names to detect phishing infrastructure. Monitor your email abuse@ inbox for reports of impersonation using stolen customer contact data.

Evidence: During recovery validation, document: (1) Token introspection responses or IdP screenshots confirming all Context AI-associated tokens show 'active: false' or equivalent revoked state — timestamp these confirmations for breach notification records; (2) Vercel's official notification to your account (email to account owner or security contact) confirming whether your tenant's data was in scope — retain this as a regulatory evidence artifact; (3) Any customer-reported phishing attempts or credential stuffing activity against your application that may indicate the stolen Vercel data (emails, project metadata) is being operationalized by the attacker downstream.

Step 5: Post-Incident — This incident exposes a specific control gap: third-party AI tools operating with broad employee-level access to systems containing customer data, without session isolation or behavioral monitoring. Conduct a full inventory of AI tool integrations and their permission scopes. Apply CIS Control 5 (Account Management) and CIS Control 12 (Network Infrastructure Management) to restrict lateral reach. Map existing vendor access against NIST SP 800-53 SA-9 (External Information System Services) controls. Require vendors with any data-path access to provide current SOC 2 Type II reports and incident notification SLAs before re-authorization.

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 SA-9 (External System Services), NIST SI-5 (Security Alerts, Advisories, and Directives), CIS 5.1 (Establish and Maintain an Inventory of Accounts), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 6.3 (Require MFA for Externally-Exposed Applications)

Compensating: Build a third-party AI tool access registry in a spreadsheet with columns: tool name, OAuth client_id, scopes granted, data stores reachable, last access date, SOC 2 status, incident notification SLA. Review and attest quarterly. For behavioral monitoring of re-authorized integrations without a SIEM, configure your IdP to send audit log webhooks to a free Slack channel or email alias — this gives a two-person team visibility into unusual token use (e.g., Context AI accessing systems at 3 AM) without requiring a commercial SIEM. For Vercel-specific hardening: enable Vercel's Audit Log streaming (available on Pro/Enterprise plans) to an S3 bucket or webhook for persistent retention

beyond the default dashboard window.

Evidence: For the post-incident lessons-learned record, compile and retain: (1) The complete timeline of Context AI's OAuth grant history — first grant date, scope changes over time, and revocation date — to quantify the exposure window for breach notification calculations; (2) The full list of Vercel projects, environment variables, and deployment secrets that were accessible to the Context AI integration, as documented before removal in Step 3 — this defines the data-at-risk boundary for any required regulatory notification; (3) Vercel's published incident report or KB bulletin content as of the lessons-learned date, to document what Vercel confirmed vs. what remains unconfirmed about data types and customer scope.

Detection Guidance

No confirmed IOCs have been published from Vercel or Context AI as of initial reporting. Detection should focus on behavioral indicators rather than static IOCs. In IdP and SSO audit logs, look for: OAuth token grants to Context AI or connected AI tool client IDs, followed by access to data stores or customer record systems. In cloud provider logs (CloudTrail, GCP Audit Logs, Azure Activity Logs), search for API calls to customer data resources originating from service principals or OAuth sessions tied to third-party tool integrations. Flag cloud account access patterns from new geographic locations, new devices, or outside established usage windows. Flag token issuance events not initiated by the account owner, or token use from IP ranges inconsistent with employee locations. If your organization uses Vercel, cross-reference your Vercel team membership and API key issuance logs against the April 2026 incident window. Monitor threat intelligence feeds for Context AI-specific IOCs as the vendor investigation matures; none were confirmed in initial disclosures.

Framework Mappings

MITRE-ATTACK

- **T1552** — Unsecured Credentials
- **T1078** — Valid Accounts
- **T1078.004** — Cloud Accounts
- **T1199** — Trusted Relationship
- **T1528** — Steal Application Access Token

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **IA-8** — Identification and Authentication (Non-Organizational Users)
- **SR-2** — Supply Chain Risk Management Plan

OWASP-TOP10-2021

- **A07:2021** — Identification and Authentication Failures
- **A04:2021** — Insecure Design

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
- **15.1** — Establish and Maintain an Inventory of Service Providers

SOC2-TSC

- **CC6.1** — The entity implements logical access security software, infrastructure, and architectures over protected information assets
- **CC9.2** — Manages risks associated with vendors and business partners

HIPAA-SECURITY

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

NIST-CSF-2

- **GV.SC-01** — Cybersecurity supply chain risk management program

ISO-27001-2022

- **A.5.21** — Managing information security in the ICT supply chain
- **A.5.23** — Information security for use of cloud services

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1552	Unsecured Credentials	Credential-Access
T1078	Valid Accounts	Defense-Evasion
T1078.004	Cloud Accounts	Defense-Evasion
T1199	Trusted Relationship	Initial-Access
T1528	Steal Application Access Token	Credential-Access

Sources

Source	URL	Tier
	https://techcrunch.com/2026/04/20/app-host-vercel-confirms-security...	T2
Vercel April 2026 security incident | Vercel Knowledge Base	https://vercel.com/kb/bulletin/vercel-april-2026-security-incident	T3

Source	URL	Tier
AI cloud company Vercel breached after employee grants AI tool ...	https://www.reddit.com/r/technology/comments/1sr2wwt/ai_cloud_compa...	T3
Vercel Employee's AI Tool Access Led to Data Breach - Dark Reading	https://www.darkreading.com/application-security/vercel-employees-a...	T3
Vercel Breach Tied to Context AI Hack Exposes Limited Customer ...	https://thehackernews.com/2026/04/vercel-breach-tied-to-context-ai-...	T3

Generated 2026-04-23 06:38 UTC by TJS Security Command Center