INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-17 08:05 UTC

Vertex AI SDK Bucket Squatting Enables No-Credential Model Poisoning and Cross-Tenant Code Execution

CVE VULNERABILITY | HIGH | CVSS 7.5

SCC Item ID	SCC-CVE-2026-0317
Type	CVE Vulnerability
CVE ID	CVE-2026-2473
Severity	HIGH
CVSS Base Score	7.5
EPSS Score	0.0041 (33th percentile)
Affected Products	Google Cloud Vertex AI SDK for Python (versions prior to 1.148.0), Google Cloud Storage, Google Cloud Vertex AI
Published	2026-06-16T15:05:41
Discovery Source	Rss

Executive Summary

A logic flaw in the Google Cloud Vertex AI Python SDK allowed attackers to hijack AI model uploads using only a target's public project ID, requiring no stolen credentials or prior access. By pre-registering a predictably named staging bucket, an attacker could substitute a malicious model payload that executed code inside Google's serving infrastructure and stole OAuth tokens capable of accessing cross-tenant cloud resources. Organizations using the Vertex AI Python SDK below version 1.148.0 to upload or serve models should treat this as an urgent patching item.

Technical Analysis

CVE-2026-2473 (CVSS 7.5, High) affects the Google Cloud Vertex AI SDK for Python prior to version 1.148.0. The flaw stems from a predictable staging bucket naming convention (CWE-345: Insufficient Verification of Data Authenticity; CWE-362: Race Condition). An unauthenticated attacker who knows a victim's public GCP project ID can pre-create the expected staging bucket name in their own project before the victim's SDK creates it. When the victim subsequently uploads a model, the SDK writes to the attacker-controlled bucket. The attacker replaces the model artifact with a malicious Python pickle payload (CWE-502: Deserialization of Untrusted Data), which executes upon model load inside Google's Vertex AI serving infrastructure. Executed code can steal OAuth tokens scoped to cross-tenant resources (T1528, T1552.007), enabling lateral movement beyond

the victim project (T1550.001). Additional mapped techniques: T1059.006 (Python execution), T1190 (exploit public-facing application), T1525 (implant container image), T1565.001 (stored data manipulation), T1195.001 (supply chain compromise). CWE-284 (Improper Access Control) reflects the absence of bucket ownership verification. Google patched the issue in SDK v1.148.0 by enforcing bucket ownership validation before write operations. Google's security activity in 2026 has included multiple predictable-resource-naming vulnerabilities in cloud platform defaults, indicating attention to this class of defect. No CISA KEV listing as of the configuration date. EPSS score 0.00414 (32.9th percentile), reflecting low observed exploitation activity to date.

Action Checklist

1. Step 1: Containment. Immediately identify all environments running google-cloud-aiplatform (Vertex AI Python SDK) below version 1.148.0. Suspend model upload and training jobs in those environments until the SDK is upgraded. Audit GCP project IDs for public exposure; restrict project ID visibility where IAM policy permits. Reference: NVD entry for CVE-2026-2473 (https://nvd.nist.gov/vuln/detail/cve-2026-2473) and Google Cloud official security advisory.
2. Step 2: Detection. Query GCP Cloud Storage audit logs (data access logs) for CreateBucket events on staging bucket names matching the Vertex AI default naming pattern (typically 'vertex-ai-staging-') originating from project IDs you do not own. Review Vertex AI model upload job logs for unexpected bucket references or model artifact checksums that do not match your internal build pipeline. Inspect GCP IAM audit logs for OAuth token issuance events (method: 'GenerateAccessToken') tied to service accounts used by Vertex AI serving jobs, particularly any cross-project resource access. MITRE techniques to hunt: T1528 (token theft), T1565.001 (model artifact manipulation), T1195.001 (supply chain staging).
3. Step 3: Eradication. Upgrade google-cloud-aiplatform to version 1.148.0 or later using: 'pip install --upgrade google-cloud-aiplatform>=1.148.0'. Validate the installed version with: 'pip show google-cloud-aiplatform'. For any model artifacts uploaded through an affected SDK version, treat all staging bucket contents as untrusted; re-upload clean, verified model artifacts through the patched SDK. Pre-delete or take ownership of predictably named staging buckets in your GCP project to prevent re-exploitation. Ensure the Vertex AI SDK upgrade is tracked in your software asset inventory and automated patch management process (per NIST SI-2, Flaw Remediation).
4. Step 4: Recovery. After upgrading, re-run model uploads and confirm staging buckets are created within your own GCP project namespace. Rotate service account keys and OAuth credentials associated with any Vertex AI serving jobs that ran models uploaded through an affected SDK version (credential rotation per NIST AC-2.1). Verify model artifact integrity against known-good checksums from your internal registry before re-deploying to production serving endpoints. Monitor Vertex AI serving job logs and IAM activity for 30 days post-remediation for anomalous cross-project API calls. Reference NIST SI-4 (System Monitoring) for post-remediation surveillance scope.
5. Step 5: Post-Incident. This vulnerability exposes a gap in ML pipeline supply chain controls: model artifact integrity is not verified at ingest. Implement cryptographic signing and verification for all model artifacts before upload (file integrity and provenance verification per NIST SI-7). Review all GCP default resource naming conventions used by your ML platform to identify other predictably named resources an attacker could pre-register. Apply NIST AC-6 (Least Privilege) to Vertex AI service account scopes, limiting cross-project OAuth token permissions to the minimum required. Establish a vulnerability management process that includes ML SDK dependencies in your standard patching cadence (NIST SI-2). Document this incident pattern for future architecture reviews of AI/ML pipeline designs on shared cloud platforms.

Detection Guidance

Primary log sources: GCP Cloud Storage Data Access audit logs and GCP IAM Activity logs. Query Cloud Storage logs for CreateBucket events where the bucket name matches the Vertex AI staging pattern ('vertex-ai-staging-') and the initiator project does not match your project. Query IAM logs for 'GenerateAccessToken' calls from Vertex AI service accounts followed by cross-project resource access within the same session window. In Vertex AI pipeline logs, flag any model load events where the artifact URI resolves to a bucket not owned by your project. Behavioral indicator: a model artifact that was not produced by your internal build or registry system appearing in a Vertex AI staging bucket. If your environment uses SIEM ingestion of GCP logs, create a correlation rule pairing a bucket write event to an external project with a subsequent model deployment event within the same pipeline run. For pickle-based payloads (CWE-502), endpoint-level detection on Vertex AI managed instances is not directly available to tenants; rely on artifact provenance controls and bucket ownership validation in the patched SDK. File integrity and provenance verification principles (per NIST SI-7, System Monitoring) apply to model artifact monitoring in your registry. No public IOCs (IPs, domains, hashes) have been identified for active exploitation of this vulnerability as of the configuration date.

Framework Mappings

MITRE-ATTACK

- **T1528** — Steal Application Access Token
- **T1550.001** — Application Access Token
- **T1059.006** — Python
- **T1190** — Exploit Public-Facing Application
- **T1525** — Implant Internal Image
- **T1565.001** — Stored Data Manipulation
- **T1552.007** — Container API
- **T1195.001** — Compromise Software Dependencies and Development Tools

NIST-800-53R5

- **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
- **SI-10** — Information Input Validation
- **SI-16** — Memory Protection
- **AC-3** — Access Enforcement
- **AT-2** — Literacy Training and Awareness

OWASP-TOP10-2021

- **A08:2021** — Software and Data Integrity Failures
- **A01:2021** — Broken Access Control

CIS-V8

- **16.10** — Apply Secure Design Principles in Application Architectures
- **2.5** — Allowlist Authorized Software
- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **6.3** — Require MFA for Externally-Exposed Applications
- **7.3** — Perform Automated Operating System Patch Management
- **7.4** — Perform Automated Application Patch Management
- **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

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control
- **164.312(d)** — Person or Entity Authentication
- **164.308(a)(5)(i)** — Security Awareness and Training

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
T1528	Steal Application Access Token	Credential-Access
T1550.001	Application Access Token	Defense-Evasion
T1059.006	Python	Execution
T1190	Exploit Public-Facing Application	Initial-Access
T1525	Implant Internal Image	Persistence
T1565.001	Stored Data Manipulation	Impact
T1552.007	Container API	Credential-Access

Technique ID	Technique Name	Tactic
T1195.001	Compromise Software Dependencies and Development Tools	Initial-Access

Sources

Source	URL	Tier
Security News	https://thehackernews.com/2026/06/google-vertex-ai-sdk-flaw-let-att...	T3
CVE-2026-2473 Detail - NVD	https://nvd.nist.gov/vuln/detail/cve-2026-2473	T1
CVE-2026-2473: Google Cloud Vertex AI RCE Vulnerability	https://www.sentinelone.com/vulnerability-database/cve-2026-2473/	T3
CVE-2026-2473 - CVE Details, Severity, and Analysis | Strobes VI	https://strobes.co/vi/cve/CVE-2026-2473	T3
Google Security Advisory	https://chromereleases.googleblog.com/	T1

Generated 2026-06-17 08:05 UTC by TJS Security Command Center