INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-13 18:25 UTC

APT41 Deploys Evasive Backdoor Targeting Credentials Across Major Cloud Platforms via Typosquatted C2

THREAT CAMPAIGN | HIGH | CVSS 9.5

SCC Item ID	SCC-CAM-2026-0174
Type	Threat Campaign
Severity	HIGH
CVSS Base Score	9.5
Affected Products	AWS, Google Cloud Platform, Microsoft Azure, Alibaba Cloud, specific services and versions unconfirmed in available source data
Published	2026-04-13T11:08:12
Discovery Source	Rss

Executive Summary

■■ CONFIDENCE NOTE: This campaign report is sourced from secondary threat intelligence; primary corroboration from CISA or authoritative vendor has not been confirmed. Treat with elevated scrutiny and monitor for authoritative validation.

APT41 (Double Dragon), a Chinese state-sponsored threat group, is reported to be operating a credential harvesting campaign targeting organizations using AWS, Google Cloud, Microsoft Azure, and Alibaba Cloud. The group is deploying a backdoor that disguises command-and-control traffic using typosquatted domains mimicking legitimate cloud service endpoints, making detection at the network layer difficult. If confirmed, this campaign represents a sustained intelligence collection operation with potential to compromise cloud-hosted data, identities, and downstream enterprise systems across all major cloud providers.

Technical Analysis

APT41 is reported to have deployed a novel backdoor targeting credential material across AWS, Google Cloud Platform, Microsoft Azure, and Alibaba Cloud. The campaign uses typosquatting of cloud API and service domains (CWE-1021: Improper Restriction of Rendered UI Layers, applicable to domain confusion) to blend C2 traffic with legitimate cloud API calls, complicating DNS and HTTP-layer detection. Primary weakness classes are CWE-522 (Insufficiently Protected Credentials) and CWE-287 (Improper Authentication), consistent with

credential harvesting objectives. Relevant MITRE ATT&CK techniques include: T1078/T1078.004 (Valid Accounts, Cloud Accounts), T1071.001 (Web Protocols for C2), T1528 (Steal Application Access Token), T1530 (Data from Cloud Storage), T1552.005 (Cloud Instance Metadata API), T1550.001 (Application Access Token reuse), T1568 (Dynamic Resolution, likely tied to typosquatted C2 infrastructure), T1583.001 (Domain acquisition for C2), and T1556 (Modify Authentication Process). No CVE is assigned. Affected specific service versions are unconfirmed in available source data.

Action Checklist

1. Containment: Audit all active cloud IAM credentials, service account keys, and OAuth tokens across AWS, GCP, Azure, and Alibaba Cloud. Rotate all credentials that have accessed unusual endpoints, atypical geolocations, or that show anomalous API call patterns. For credentials with no detected anomalies, prioritize those with broad permissions (Admin, PowerUser equivalent) for rotation within 30 days. Flag and disable accounts with anomalous API call patterns.
2. Detection: Query DNS resolver logs and cloud-native DNS services (Route 53 Resolver, Cloud DNS, Azure DNS) for lookups matching typosquatted patterns of known cloud service domains (e.g., variations on amazonaws.com, googleapis.com, azure.com, aliyuncs.com). Review CloudTrail (AWS), Cloud Audit Logs (GCP), Azure Monitor Activity Logs, and Alibaba ActionTrail for T1528 indicators: unexpected token issuances, metadata API calls from non-standard source IPs, and ListBuckets/ListObjects calls outside expected service accounts.
3. Eradication: Remove any unrecognized IAM roles, service principals, or API keys. Enforce conditional access policies requiring MFA for all cloud console and API access. Block outbound DNS and HTTPS to typosquatted domains identified through log review. No vendor patch is applicable; this is an identity and network control gap, not a software vulnerability with a patch ID.
4. Recovery: Validate that all rotated credentials have propagated and that no legacy keys remain active. Re-baseline cloud access patterns using SIEM or cloud-native CSPM tooling. Monitor for re-authentication attempts using revoked tokens (T1550.001). Confirm no persistent backdoor access exists via reviewing all federated identity configurations and third-party OAuth grants.
5. Post-Incident: Conduct a cloud identity hygiene review against CIS Benchmarks for AWS, Azure, and GCP. Implement DNS-layer monitoring with anomaly alerting for newly registered or typosquatted domains. Map gaps to NIST CSF PR.AC-4 (access permissions) and DE.CM-7 (monitoring for unauthorized connections). If this campaign is confirmed against your environment, treat as a potential espionage-motivated intrusion and engage IR counsel before disclosing findings externally.

IR / Forensic Enrichment

Triage Priority	IMMEDIATE

Escalation Criteria	Escalate immediately to executive leadership and legal/IR counsel if CloudTrail, GCP Audit Logs, Azure Monitor, or Alibaba ActionTrail confirm any of the following: successful resolution of typosquatted APT41 C2 domains from your environment, API calls to cloud metadata services (IMDS) from unexpected source IPs, unauthorized token issuances or federated identity authentications, or evidence that compromised credentials accessed buckets, blobs, or storage objects containing PII, PHI, or regulated data — as these conditions collectively indicate a confirmed APT41 intrusion with potential espionage impact and may trigger breach notification obligations under GDPR, HIPAA, or relevant state privacy laws.
Recovery Notes	After all credentials are rotated and typosquatted C2 domains are blocked at the DNS layer, maintain heightened monitoring of CloudTrail (AWS), Cloud Audit Logs (GCP), Azure Monitor, and Alibaba ActionTrail for a minimum of 30 days for re-authentication attempts using revoked tokens (T1550.001) and any new IAM entities created outside your approved provisioning pipeline, as APT41 is a persistent, well-resourced group known to re-establish access after initial eviction. Re-baseline all cloud API call patterns in your SIEM or CSPM tool against the new post-incident normal before reducing alert thresholds. Conduct a federated identity and OAuth grant audit quarterly for the 12 months following this incident to detect any persistence mechanisms that survived the initial eradication sweep.
Forensic Artifacts	AWS CloudTrail JSON logs (90-day window): filter on eventSource 'sts.amazonaws.com' for AssumeRoleWithWebIdentity and GetFederationToken events, and on eventName ListBuckets/ListObjects from service account principals not in approved baseline — these are the specific API call patterns APT41 would generate during credential harvesting and data enumeration (MITRE T1528, T1530) | Route 53 Resolver Query Logs / GCP Cloud DNS Logs / Azure DNS diagnostic logs: raw FQDN query records showing any lookups against typosquatted variants of amazonaws.com, googleapis.com, azure.com, or aliyuncs.com, which represent the APT41 C2 communication channel and are the primary network-layer indicator of compromise for this campaign | Cloud platform IAM credential reports (AWS) and service account key listings (GCP / Azure service principal credential exports): point-in-time snapshots capturing key creation dates, last-used timestamps, and associated principal permissions, establishing which credentials were active during the suspected APT41 activity window and what data they could have accessed | Third-party OAuth grant exports from all four cloud platforms (AWS SSO application assignments, GCP Workspace token grants, Azure AD OAuth2 permission grants, Alibaba RAM role trust policies): these exports document any persistent delegated access paths APT41 may have established via T1550.001 that would survive credential rotation if not explicitly revoked | Cloud workload outbound HTTPS connection logs (AWS VPC Flow Logs, GCP VPC Flow Logs, Azure NSG Flow Logs): filter on destination port 443 to IP addresses that resolved from typosquatted domains, capturing the actual C2 traffic sessions — these logs establish the duration, volume, and timing of any data exfiltration from compromised workloads to APT41 infrastructure

Per-Action IR Details

Containment — Audit all active cloud IAM credentials, service account keys, and OAuth tokens across AWS, GCP, Azure, and Alibaba Cloud. Rotate any credentials that have accessed unusual endpoints or that were issued more than 90 days ago. Flag and disable accounts with anomalous API call patterns.

NIST Phase: Containment

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

Controls: NIST IR-4 (Incident Handling), NIST AC-2 (Account Management), NIST IA-5 (Authenticator Management), CIS 5.1 (Establish and Maintain an Inventory of Accounts), CIS 5.3 (Disable Dormant Accounts), CIS 6.2 (Establish an Access Revoking Process)

Compensating: For teams without CSPM tooling: run AWS CLI 'aws iam generate-credential-report && aws iam get-credential-report' to export all IAM user key ages and last-used timestamps; for GCP, use 'gcloud iam service-accounts list' combined with 'gcloud iam service-accounts keys list --iam-account=' to enumerate keys and their creation dates; for Azure, use 'az ad sp list --all --query "[].{displayName:displayName, appId:appId}"' piped to review service principal credentials. Cross-reference each key's last-used timestamp against your CloudTrail/Audit Log baseline to flag keys that called endpoints matching typosquatted domain patterns.

Evidence: Before rotating any credential, capture the full AWS CloudTrail event history for that principal (filter on 'userIdentity.principalId' or 'userIdentity.accessKeyId') to preserve a record of every API call made with the potentially compromised key — particularly AssumeRole, GetCallerIdentity, and any S3 ListBuckets/GetObject events that may represent APT41 data staging (MITRE T1530). For GCP, export Cloud Audit Logs Admin Activity and Data Access logs for the implicated service account before disabling it. For Azure, pull the AAD Sign-In Logs and Service Principal Sign-In Logs for the flagged principal from Azure Monitor before credential rotation destroys the correlation chain.

Detection — Query DNS resolver logs and cloud-native DNS services (Route 53 Resolver, Cloud DNS, Azure DNS) for lookups matching typosquatted patterns of known cloud service domains (e.g., variations on amazonaws.com, googleapis.com, azure.com, aliyuncs.com). Review CloudTrail (AWS), Cloud Audit Logs (GCP), Azure Monitor Activity Logs, and Alibaba ActionTrail for T1528 indicators: unexpected token issuances, metadata API calls from non-standard source IPs, and ListBuckets/ListObjects calls outside expected service accounts.

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), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: Without a SIEM, use AWS CloudTrail Lake or Athena with the query: 'SELECT eventTime, userIdentity.principalId, eventName, sourceIPAddress FROM cloudtrail_logs WHERE eventName IN ("AssumeRoleWithWebIdentity","GetFederationToken","ListBuckets") AND sourceIPAddress NOT LIKE "%amazonaws.com%" ORDER BY eventTime DESC'. For DNS typosquatting detection without enterprise tooling, pull Route 53 Resolver query logs to S3 and run a Python script using editdistance library to flag any resolved domain with Levenshtein distance ≤ 2 from 'amazonaws.com', 'googleapis.com', 'azure.com', or 'aliyuncs.com'. A free Sigma rule targeting CloudTrail JSON (rule category: cloud, logsource: aws.cloudtrail) can be run against exported logs using sigma-cli with an Elasticsearch or plain Python backend.

Evidence: Capture Route 53 Resolver Query Logs (enable via Route 53 console → Resolver → Query logging) and export to S3 before any network-layer blocking — these logs contain the exact FQDN queried, the source VPC, and the resolver response, which will show whether APT41 C2 typosquatted domains were successfully resolved from your environment. Simultaneously, export raw CloudTrail JSON for the 90-day window prior to detection, specifically filtering for eventSource 'sts.amazonaws.com' (token issuance), 'ec2.amazonaws.com' with eventName 'DescribeInstances' from unexpected source IPs (consistent with APT41 cloud enumeration TTPs mapped to MITRE T1526 — Cloud Infrastructure Discovery), and any IMDS/metadata service calls (sourceIPAddress '169.254.169.254') from workloads not in your approved compute baseline.

Eradication — Remove any unrecognized IAM roles, service principals, or API keys. Enforce conditional access policies requiring MFA for all cloud console and API access. Block outbound DNS and HTTPS to typosquatted domains identified through log review. No vendor patch is applicable — this is an identity and network control gap, not a software vulnerability with a patch ID.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST IR-4 (Incident Handling), NIST AC-2 (Account Management), NIST AC-3 (Access Enforcement), NIST SC-7 (Boundary Protection), NIST SI-2 (Flaw Remediation), CIS 4.4 (Implement and Manage a Firewall on Servers), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 6.5 (Require MFA for Administrative Access)

Compensating: To enumerate and remove unrecognized IAM entities without CSPM: run 'aws iam list-roles --query "Roles[?CreateDate' for any not in inventory. For blocking typosquatted C2 domains without an enterprise DNS firewall, configure AWS Route 53 Resolver DNS Firewall (free tier available) with a rule group blocking identified typosquatted FQDNs; on GCP, use Cloud DNS Response Policy Zones to return NXDOMAIN for flagged domains; on-premises or in workloads, deploy Pi-hole or use /etc/hosts entries as a zero-cost DNS sinkhole for identified APT41 C2 domains.

Evidence: Before deleting any IAM role or service principal, run 'aws iam simulate-principal-policy' and 'aws iam list-attached-role-policies' to document the full permission scope of the suspicious entity — this establishes the blast radius of what APT41 could have accessed if the credential was compromised (required for breach scoping under NIST 800-61r3 §3.4). For Azure service principals flagged for removal, export the full audit log entry from AAD showing the principal's creation event, consent grants, and all OAuth token issuances (available via Azure AD Audit Logs, category 'ApplicationManagement') before deletion, as these records are required for post-incident forensic reconstruction and potential regulatory notification.

Recovery — Validate that all rotated credentials have propagated and that no legacy keys remain active. Re-baseline cloud access patterns using SIEM or cloud-native CSPM tooling. Monitor for re-authentication attempts using revoked tokens (T1550.001). Confirm no persistent backdoor access exists via reviewing all federated identity configurations and third-party OAuth grants.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST IR-4 (Incident Handling), NIST IA-5 (Authenticator Management), NIST SI-7 (Software, Firmware, and Information Integrity), NIST AU-6 (Audit Record Review, Analysis, and Reporting), CIS 5.1 (Establish and Maintain an Inventory of Accounts), CIS 6.1 (Establish an Access Granting Process), CIS 6.2 (Establish an Access Revoking Process)

Compensating: To verify no legacy AWS keys remain active after rotation, re-run 'aws iam generate-credential-report' 15 minutes post-rotation and confirm all previously flagged access_key_last_used timestamps predate the rotation event; any key showing post-rotation activity indicates either propagation failure or an undetected duplicate credential. To detect T1550.001 (Pass-the-Token) re-authentication attempts without a SIEM, configure an AWS CloudWatch Metric Filter on CloudTrail for eventName 'AssumeRoleWithWebIdentity' where the WebIdentityToken subject matches a revoked federated identity, and set a CloudWatch Alarm to page on any match — this is a zero-cost detection for APT41 attempting to reuse harvested tokens against your environment.

Evidence: Before declaring recovery complete, export the full list of third-party OAuth application grants from each platform (AWS: 'aws sso-admin list-application-assignments'; GCP: Google Workspace Admin SDK → tokens.list; Azure: 'Get-AzureADOAuth2PermissionGrant' via Microsoft Graph) and manually review for any grant created or last-used during the suspected APT41 activity window — OAuth grants are a known APT41 persistence mechanism (MITRE T1550.001) that survive credential rotation if not explicitly revoked. Additionally, audit all SAML federation configurations and external identity provider trust relationships in each cloud platform's IAM federation settings, as APT41 has been observed establishing persistent federated access paths that are not invalidated by key rotation alone.

Post-Incident — Conduct a cloud identity hygiene review against CIS Benchmarks for AWS, Azure, and GCP. Implement DNS-layer monitoring with anomaly alerting for newly registered or typosquatted domains. Map gaps to NIST CSF PR.AC-4 (access permissions) and DE.CM-7 (monitoring for unauthorized connections). If this campaign is confirmed against your environment, treat as a potential espionage-motivated intrusion and engage IR counsel before disclosing findings externally.

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-6 (Audit Record Review, Analysis, and Reporting), NIST SI-5 (Security Alerts, Advisories, and Directives), NIST RA-3 (Risk Assessment), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process), CIS 8.2 (Collect Audit Logs)

Compensating: For teams without a dedicated CSPM tool, run Prowler (open-source, free: github.com/prowler-cloud/prowler) against AWS, Azure, and GCP post-incident — it maps findings directly to CIS Benchmarks and NIST controls, providing a structured gap report without licensing cost. For DNS-layer typosquatting monitoring going forward, configure a free dnstwist (github.com/elceef/dnstwist) cron job running weekly against your organization's cloud domain list and alert on any newly registered permutations — this directly addresses APT41's reported C2 technique of typosquatted cloud service domains and costs nothing beyond a Linux VM.

Evidence: For the lessons-learned record required by NIST 800-61r3 §4, preserve a complete timeline artifact package including: all CloudTrail/Audit Log exports covering the full suspected compromise window (minimum 90 days prior to detection), the DNS resolver query log exports showing any successful resolution of typosquatted domains, the pre- and post-rotation IAM credential reports, and the list of all OAuth grants and federation configurations reviewed during recovery — this package constitutes the forensic record of the APT41 campaign's footprint in your environment and is required both for internal lessons-learned and for any regulatory notification analysis if PII or sensitive data was accessible to the compromised credentials.

Detection Guidance

Primary detection surface is DNS and cloud audit logs. Query for outbound DNS lookups to domains with high visual similarity to canonical cloud endpoints, character substitutions (rn for m, 0 for o), added hyphens, or TLD variations on amazonaws.com, googleapis.com, azure.com, login.microsoftonline.com, and aliyuncs.com. In AWS CloudTrail, alert on GetCallerIdentity calls from unexpected source IPs, AssumeRoleWithWebIdentity from external identity providers not in your approved list, and IMDSv1 metadata API calls (T1552.005). In GCP Cloud Audit Logs, flag service account key creation events and storage.objects.list calls from service accounts with no prior storage access history. In Azure Monitor, alert on OAuth token grants to unrecognized applications and Sign-In logs showing access from atypical geolocations using valid credentials (T1078.004). Behavioral indicator: cloud API calls to metadata endpoints followed immediately by outbound HTTPS to a newly registered domain; this pattern is consistent with T1552.005 into T1071.001 C2 callback. No confirmed IOC hashes, IPs, or specific domains are available from verified sources at this time.

Indicators of Compromise

Type	Value	Context	Confidence
DOMAIN	`[not confirmed — no verified IOCs available from authoritative sources]`	Typosquatted domains mimicking cloud service endpoints reported as C2 infrastructure; specific domains not confirmed in available verified source data	LOW

Framework Mappings

MITRE-ATTACK

- **T1078** — Valid Accounts
- **T1078.004** — Cloud Accounts
- **T1071.001** — Web Protocols
- **T1528** — Steal Application Access Token

- **T1530** — Data from Cloud Storage
- **T1583.001** — Domains
- **T1556** — Modify Authentication Process
- **T1552.005** — Cloud Instance Metadata API
- **T1550.001** — Application Access Token
- **T1568** — Dynamic Resolution

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **SI-4** — System Monitoring
- **SI-7** — Software, Firmware, and Information Integrity
- **IA-8** — Identification and Authentication (Non-Organizational Users)

OWASP-TOP10-2021

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

CIS-V8

- **5.2** — Use Unique Passwords
- **6.3** — Require MFA for Externally-Exposed Applications
- **6.4** — Require MFA for Remote Network Access
- **6.5** — Require MFA for Administrative Access
- **8.2** — Collect Audit Logs

HIPAA-SECURITY

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

SOC2-TSC

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

NIST-CSF-2

- **DE.CM-01** — Networks and network services are monitored

ISO-27001-2022

- **A.5.23** — Information security for use of cloud services

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1078	Valid Accounts	Defense-Evasion
T1078.004	Cloud Accounts	Defense-Evasion
T1071.001	Web Protocols	Command-And-Control
T1528	Steal Application Access Token	Credential-Access
T1530	Data from Cloud Storage	Collection
T1583.001	Domains	Resource-Development
T1556	Modify Authentication Process	Credential-Access
T1552.005	Cloud Instance Metadata API	Credential-Access
T1550.001	Application Access Token	Defense-Evasion
T1568	Dynamic Resolution	Command-And-Control

Sources

Source	URL	Tier
Security News	https://www.darkreading.com/cloud-security/apt41-zero-detection-bac...	T3
GitHub - hashishrajan/cloud-security-vulnerabilities: List of all the ...	https://github.com/hashishrajan/cloud-security-vulnerabilities	T3
LeakyCLI: AWS & Google Cloud Command Line Tools - Orca Security	https://orca.security/resources/blog/leakycli-aws-google-cloud-comm...	T3
60K Servers Wormed Across AWS, Azure & GCP + AI Governance ...	https://www.cloudsecuritynewsletter.com/p/60k-cloud-servers-wormed-...	T3
HIPAA Vault Show Episode 52	https://www.hipaavault.com/podcast/episode-52-aws-vs-google-cloud-v...	T3

Generated 2026-04-13 18:25 UTC by TJS Security Command Center