INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-30 14:46 UTC

Cloud Breach Rates Near Universal: 94% of Enterprises Report Intrusions Amid Visibility and Detection Failures

SECURITY ANALYSIS | HIGH | CVSS 7.5

SCC Item ID	SCC-STY-2026-0306
Type	Security Analysis
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Multi-cloud enterprise environments (AWS, Azure, Google Cloud); AI/ML cloud workloads; organizations using CrowdStrike Falcon Cloud Security
Discovery Source	Rss:T1 Threatintel

Executive Summary

According to a CrowdStrike survey of enterprise security teams, 94% of organizations reported cloud intrusions resulting in data exposure or exfiltration, with CrowdStrike's own telemetry indicating cloud-conscious intrusions rose 37% year-over-year in 2025 and state-nexus actor targeting of cloud environments increased 266% over the same period. These figures originate from a single commercially interested vendor source and have not been independently corroborated, but the structural failure modes they describe, incomplete workload visibility, inability to distinguish malicious from legitimate cloud activity, and detection lag, are consistent with broader industry observations about multi-cloud security maturity gaps. For CISOs and boards, the signal is directional: adversaries, including state-nexus actors, have made cloud environments a primary attack surface, and most enterprises are not detecting intrusions before data exits.

Technical Analysis

The CrowdStrike State of Cloud Detection and Response survey, supported by telemetry from CrowdStrike's own platform, identifies three interconnected control failures driving cloud breach rates to what the vendor reports as near-universal levels among enterprises. First, incomplete visibility across cloud workloads and control planes leaves defenders blind to lateral movement, privilege escalation, and data staging activity that occurs outside traditional perimeter monitoring. Second, the inability to distinguish malicious from legitimate cloud API activity, a structural challenge in environments where cloud-native services generate high-volume, policy-conformant traffic, allows adversaries to operate inside the noise floor of standard alerting. Third,

detection and response timelines trail adversary operational tempo, a problem compounded in cloud environments where compute resources can be provisioned, exploited, and deprovisioned faster than alert triage cycles.

The MITRE ATT&CK techniques associated with this campaign profile map closely to these failure modes. Cloud service discovery (T1580) and storage object discovery (T1619) exploit visibility gaps to allow adversaries to enumerate cloud infrastructure before acting. Cloud instance metadata API abuse (T1552.005) and valid cloud account abuse (T1078.004, T1078) exploit the legitimate-versus-malicious distinction problem, as these techniques use authorized API calls that are structurally indistinguishable from normal operations without behavioral baselining. Impair defenses via disable cloud logs (T1562.008) directly targets the logging and detection layer, suppressing the audit trail needed for detection and forensic reconstruction. Data from cloud storage (T1530) and transfer data to cloud account (T1537) represent the exfiltration endpoint of this kill chain. Abuse elevation control mechanism (T1548) and exploit public-facing application (T1190) round out the initial access and privilege escalation phases.

CrowdStrike's telemetry separately attributes a 266% year-over-year increase in state-nexus cloud targeting to unnamed actors catalogued in its 2026 Global Threat Report. The structural weakness enabling these intrusions maps to CWE-284 (improper access control), CWE-200 (exposure of sensitive information to unauthorized actors), CWE-732 (incorrect permission assignment for critical resources), and CWE-778 (insufficient logging). The AI/ML workload dimension noted in affected systems is consistent with CrowdStrike's published guidance on protecting AI development pipelines, where model training data, inference endpoints, and pipeline credentials represent high-value targets with immature security tooling coverage. All statistical claims in this story derive from CrowdStrike sources and should be evaluated with awareness of the vendor's commercial interest in reporting elevated threat severity.

Action Checklist

1. Step 1: Assess exposure, inventory all cloud environments (AWS, Azure, Google Cloud) and determine whether cloud workload protection, CSPM, and control plane logging are deployed consistently across all accounts and regions, not just primary production environments.
2. Step 2: Review controls, verify that cloud audit logging is enabled and protected from tampering (NIST AU-2 Event Logging, NIST AU-9 Protection of Audit Information, CIS 8.2 Collect Audit Logs); audit IAM roles and service account permissions for least-privilege compliance (NIST AC-6 Least Privilege, CIS 5.4 Restrict Administrator Privileges); require MFA on all cloud management console and API access (CIS 6.3 Require MFA for Externally-Exposed Applications, CIS 6.4 Require MFA for Remote Network Access, CIS 6.5 Require MFA for Administrative Access); and validate that default or unused accounts in cloud environments are disabled (CIS 4.7 Manage Default Accounts, CIS 5.3 Disable Dormant Accounts).
3. Step 3: Update threat model, incorporate T1562.008 (impair defenses: disable cloud logs), T1552.005 (cloud instance metadata API), T1078.004 (valid cloud accounts), and T1537 (transfer data to cloud account) as active TTP hypotheses in your cloud threat register; flag state-nexus actors as in-scope adversaries for cloud environments given the reported targeting trend.
4. Step 4: Communicate findings, brief leadership on the specific gap between adversary speed in cloud environments and current detection and response cycle times; frame as a detection engineering and visibility investment decision, not a compliance checkbox; reference the 37% and 266% YoY figures as vendor-reported directional indicators, not verified industry benchmarks.

5. Step 5: Monitor developments, track CrowdStrike's 2026 Global Threat Report release for expanded attribution and TTP detail; watch for independent corroboration of cloud intrusion rate statistics from sources such as CISA, Verizon DBIR, or Mandiant M-Trends (now Google Cloud Threat Intelligence) before treating the 94% figure as an established industry baseline.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to CISO and legal counsel if CloudTrail, Azure Monitor, or GCP Cloud Audit Logs show evidence of T1562.008 activity (StopLogging, DeleteTrail, or PutEventSelectors reducing scope) in any production account, or if IAM credential report reveals active sessions or API key usage from unrecognized principals — either condition indicates a probable active cloud-conscious intrusion requiring breach notification assessment under applicable data protection regulations (GDPR Art. 33, HIPAA §164.412, or applicable state breach statutes) if sensitive data is hosted in the affected environments.
Recovery Notes	After remediating IAM over-privilege and re-enabling suppressed logging across all cloud accounts and regions, verify log integrity by confirming CloudTrail log file validation hashes are intact and that no gaps exist in the CloudTrail event history for the 90 days preceding remediation — gaps are themselves evidence of T1562.008 activity. Monitor for re-enablement of any disabled accounts, new cross-account AssumeRole relationships, and outbound replication or snapshot-sharing events for a minimum of 30 days post-remediation, as state-nexus actors with established footholds via valid cloud accounts frequently re-enter through dormant credentials or pre-staged persistence mechanisms that survive initial IAM cleanup. Validate that all cloud workloads in non-production accounts (dev, staging, sandbox) are subject to the same CSPM and logging controls as production, as these environments are the most common blind-spot entry point identified in cloud-conscious intrusion campaigns.
Forensic Artifacts	AWS CloudTrail management event logs filtered for StopLogging, DeleteTrail, PutEventSelectors, AssumeRole, AssumeRoleWithWebIdentity, and ConsoleLogin events — these are the primary forensic record of T1562.008 (log suppression) and T1078.004 (valid cloud account abuse) in AWS environments | AWS VPC Flow Logs and host-level HTTP access logs showing outbound connections to the instance metadata service endpoint (169.254.169.254) from non-orchestration workloads — anomalous IMDS queries are the forensic signature of T1552.005 credential harvesting from cloud instance metadata APIs | AWS S3 server access logs and CloudTrail data events for S3, filtered for PutBucketReplication, PutObject with cross-account principal, and cross-account ModifySnapshotAttribute or CreateSnapshot API calls — these are the artifact trail left by T1537 (transfer data to cloud account) in AWS-hosted data exfiltration scenarios | Azure Entra ID (Azure AD) sign-in logs and Conditional Access logs, specifically entries with authentication from unfamiliar ASNs, legacy authentication protocols, or service principal logins outside normal working hours — state-nexus actors targeting Azure environments preferentially abuse service principal credentials and OAuth tokens that bypass MFA policies | GCP Cloud Audit Logs (Admin Activity and Data Access log types) filtered for SetIamPolicy, serviceAccounts.keys.create, and storage.buckets.setIamPolicy events — these represent the control-plane manipulation footprint consistent with cloud-conscious intrusions targeting GCP environments via IAM privilege escalation and data staging

Per-Action IR Details

Step 1: Assess exposure — inventory all cloud environments (AWS, Azure, Google Cloud) and determine whether cloud workload protection, CSPM, and control plane logging are deployed consistently across all accounts and regions, not just primary production environments.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: establishing visibility and tooling coverage before an incident occurs

Controls: NIST AU-2 (Event Logging), NIST AU-12 (Audit Record Generation), CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 2.1 (Establish and Maintain a Software Inventory)

Compensating: For teams without a CSPM tool: run AWS CLI `aws ec2 describe-regions --all-regions` and `aws cloudtrail describe-trails --include-shadow-trails` across every account to enumerate regions and trail status; for Azure, use `az monitor diagnostic-settings list` per subscription; for GCP, use `gcloud logging sinks list --project=PROJECT_ID`. Pipe results to a CSV and manually diff against your asset register. Two analysts can cover a mid-sized multi-cloud estate in a half-day using these commands in parallel.

Evidence: This is a preparation step and does not alter live system state, so no volatile capture is required prior to execution. However, before executing any remediation that follows, document the current CloudTrail/Azure Monitor/GCP Cloud Audit Logs configuration state (enabled regions, S3/storage bucket ARNs, log integrity validation status) as a forensic baseline — state-nexus actors targeting cloud control planes (up 266% YoY per CrowdStrike telemetry) frequently disable or redirect logging as a first action, and this snapshot will establish what was and was not visible at the time of any potential intrusion.

Step 2: Review controls — verify that cloud audit logging is enabled and protected from tampering (NIST AU-2 Event Logging, NIST AU-9 Protection of Audit Information, CIS 8.2 Collect Audit Logs); audit IAM roles and service account permissions for least-privilege compliance (NIST AC-6 Least Privilege, CIS 5.4 Restrict Administrator Privileges); require MFA on all cloud management console and API access (CIS 6.3 Require MFA for Externally-Exposed Applications, CIS 6.4 Require MFA for Remote Network Access, CIS 6.5 Require MFA for Administrative Access); and validate that default or unused accounts in cloud environments are disabled (CIS 4.7 Manage Default Accounts, CIS 5.3 Disable Dormant Accounts).

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: hardening controls and closing structural gaps that enable cloud-conscious intrusions before detection is required

Controls: NIST AU-2 (Event Logging), NIST AU-9 (Protection Of Audit Information), NIST AC-6 (Least Privilege), NIST AC-2 (Account Management), NIST AC-3 (Access Enforcement), CIS 8.2 (Collect Audit Logs), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 6.4 (Require MFA for Remote Network Access), CIS 6.5 (Require MFA for Administrative Access), CIS 4.7 (Manage Default Accounts on Enterprise Assets and Software), CIS 5.3 (Disable Dormant Accounts)

Compensating: IAM audit without enterprise tooling: in AWS, run `aws iam generate-credential-report && aws iam get-credential-report` to identify unused credentials and accounts lacking MFA; use `aws iam list-attached-role-policies` and `aws iam simulate-principal-policy` to spot over-privileged roles. For Azure, use `az ad user list --query '[?accountEnabled==false]'` and the free Microsoft Entra ID Access Review feature. For GCP, use `gcloud projects get-iam-policy PROJECT_ID --format=json` and review for roles/owner or roles/editor bindings on service accounts. For log integrity in AWS, enable CloudTrail log file validation (`--enable-log-file-validation`) and store to an S3 bucket with Object Lock (WORM) enabled — both are free-tier compatible.

Evidence: Credential modification and MFA changes in cloud environments alter live IAM state; before revoking sessions or disabling accounts, export the full IAM credential report and capture active session tokens via `aws sts get-caller-identity` and active console sessions from CloudTrail `ConsoleLogin` events and `AssumeRole` events within the last 90 days. In Azure, export the sign-in logs and Conditional Access evaluation logs from Entra ID before any account disablement. These exports document which identities were active at assessment time and are critical for forensic reconstruction if a valid cloud account (aligned with T1078.004 TTP noted in this threat) was already in use by an adversary.

Step 3: Update threat model — incorporate T1562.008 (impair defenses: disable cloud logs), T1552.005 (cloud instance metadata API), T1078.004 (valid cloud accounts), and T1537 (transfer data to cloud account) as active TTP hypotheses in your cloud threat register; flag state-nexus actors as in-scope adversaries for cloud environments given the reported targeting trend.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: developing and maintaining threat models and hunting hypotheses specific to the organization's threat profile

Compensating: Document the threat model update in a shared wiki or incident tracker entry (even a dated Confluence page or GitHub markdown file). For each TTP hypothesis, write one Sigma rule stub targeting the relevant log source: for T1562.008, a Sigma rule against CloudTrail `StopLogging` or `DeleteTrail` events; for T1552.005, alert on unusual outbound HTTP to 169.254.169.254 from non-orchestration hosts using host-based firewall logs or VPC Flow Logs; for T1078.004, a Sigma rule on CloudTrail `ConsoleLogin` from new ASNs or geolocations not in baseline; for T1537, alert on S3 `PutBucketReplication` or `CreateSnapshot` followed by cross-account `ModifySnapshotAttribute` events. Free Sigma rule templates for cloud log sources are available in the SigmaHQ repository.

Evidence: This step does not alter live system state and requires no volatile capture prior to execution. However, the threat model update should be informed by a retrospective pull of the last 90 days of CloudTrail management events filtered for `StopLogging`, `DeleteTrail`, `PutEventSelectors` (reducing log scope), `AssumeRoleWithWebIdentity`, and cross-account `AssumeRole` calls — these are the forensic fingerprints that T1562.008 and T1078.004 would leave in a cloud-conscious intrusion campaign targeting the environments named in this threat advisory.

Step 4: Communicate findings — brief leadership on the specific gap between adversary speed in cloud environments and current detection and response cycle times; frame as a detection engineering and visibility investment decision, not a compliance checkbox; reference the 37% and 266% YoY figures as vendor-reported directional indicators, not verified industry benchmarks.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: lessons learned, reporting, and communicating gaps to leadership to drive program improvement

Controls: NIST AC-1 (Policy And Procedures), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: For teams without a formal reporting infrastructure, a one-page briefing built around three columns — current visibility gap, adversary capability (cloud-conscious intrusion dwell time), and proposed investment — is sufficient. Anchor the 37% and 266% figures with the explicit caveat that they are single-vendor, commercially interested statistics not independently corroborated by CISA, Verizon DBIR, or Mandiant M-Trends as of this advisory date, and present your own organization's CloudTrail coverage gaps (from Step 1) as the primary evidence base. This grounds the conversation in your own observable data rather than contested external statistics.

Evidence: No live system state is altered by this step; no volatile capture is required. To support the briefing with first-party evidence, pull your organization's CloudTrail event volume by region for the last 60 days and identify any regions, accounts, or services with zero or near-zero log ingestion — these coverage voids are the structural failure mode this threat advisory describes and are more compelling to leadership than vendor survey statistics.

Step 5: Monitor developments — track CrowdStrike's 2026 Global Threat Report release for expanded attribution and TTP detail; watch for independent corroboration of cloud intrusion rate statistics from sources such as CISA, Verizon DBIR, or Mandiant M-Trends before treating the 94% figure as an established industry baseline.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: using external threat intelligence and updated reporting to refine detection capabilities and improve program posture over time

Controls: NIST AU-6 (Audit Record Review, Analysis, And Reporting)

Compensating: Assign one analyst to maintain a dated log entry (a simple markdown file in your threat intel repository is sufficient) tracking publication dates and key findings from CISA cloud security advisories, the Verizon DBIR cloud chapter, and Mandiant M-Trends. Set a calendar reminder for Q1 2026 to pull the CrowdStrike Global Threat Report

and extract cloud-specific TTP updates. For state-nexus actor tracking without a commercial threat intel subscription, monitor CISA's Known Exploited Vulnerabilities catalog and Joint Advisories filtered for cloud service providers, and subscribe to the CISA mailing list for free. This keeps the threat model current without recurring cost.

Evidence: No live system state is altered by this step; no volatile capture is required. Intelligence collected under this step should be used to retroactively query your CloudTrail, Azure Monitor, and GCP Cloud Audit Logs for any newly published IOCs or TTPs associated with the state-nexus actor campaigns described in updated reports — specifically any new indicators tied to cloud instance metadata API abuse (T1552.005) or cross-account data staging (T1537) that may match artifacts already present in retained logs.

Detection Guidance

Detection focus for this story should center on four behavioral clusters aligned to the reported failure modes.

1. Log tampering and visibility suppression: Hunt for CloudTrail, Azure Monitor, or GCP Cloud Audit Logs being disabled, modified, or having retention policies shortened (maps to T1562.008). Alert on IAM policy changes that remove logging permissions from service accounts. Reference NIST AU-9 and CIS 8.2, if logging is not centralized and write-protected, these detections will not fire. D3-SFA (MITRE D3FEND: System File Analysis) countermeasure applies: monitor audit configuration state continuously, not just at deployment.

2. Abnormal API and metadata service access: Baseline cloud API call patterns per service account and alert on deviations, particularly Instance Metadata Service (IMDS) queries from workloads that do not normally make them (T1552.005), and enumeration patterns across storage buckets and compute inventory (T1580, T1619). D3-LAM (MITRE D3FEND: Local Account Monitoring) and D3-UAP (MITRE D3FEND: User Account Permissions) apply to cloud IAM context: flag accounts accessing resources outside their documented function.

3. Privilege escalation and permission boundary violations: Alert on role assumption chains, policy attachment to previously unprivileged accounts, and use of administrative APIs from non-standard source IPs or new user agents (T1548, T1078.004). NIST AC-6 (Least Privilege) and NIST AC-3 (Access Enforcement) define the control baseline; detection engineering should verify that violations of these controls generate alerts, not just log entries.

4. Unusual outbound data movement: Alert on large-scale reads from cloud storage combined with cross-account or cross-region replication events (T1530, T1537). Establish data egress baselines per account and alert on volume anomalies. NIST AC-4 (Information Flow Enforcement) and CIS 3.3 (Configure Data Access Control Lists) are the relevant controls to audit for gap coverage.

For AI/ML workload environments specifically: review access controls on model artifact storage, training data buckets, and CI/CD pipeline credentials, these are high-value targets with frequently weaker access governance than production application infrastructure.

No specific IOCs (hashes, IPs, domains) are present in the provided source material. The cited CrowdStrike sources may contain platform-specific indicators, consult CrowdStrike's Falcon Cloud Security telemetry and the 2026 Global Threat Report directly for indicator values.

Framework Mappings

MITRE-ATTACK

- **T1580** — Cloud Infrastructure Discovery
- **T1552.005** — Cloud Instance Metadata API

- **T1530** — Data from Cloud Storage
- **T1619** — Cloud Storage Object Discovery
- **T1078.004** — Cloud Accounts
- **T1548** — Abuse Elevation Control Mechanism
- **T1562.008** — Disable or Modify Cloud Logs
- **T1537** — Transfer Data to Cloud Account
- **T1078** — Valid Accounts
- **T1190** — Exploit Public-Facing Application

NIST-800-53R5

- **AC-6** — Least Privilege
- **CM-6** — Configuration Settings
- **AC-2** — Account Management
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **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
- **AC-3** — Access Enforcement
- **SC-28** — Protection of Information at Rest
- **SI-4** — System Monitoring

OWASP-TOP10-2021

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

CIS-V8

- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **3.3** — Configure Data Access Control Lists
- **8.2** — Collect Audit Logs

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(a)(1)** — Access Control

ISO-27001-2022

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

NIST-CSF-2

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

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1580	Cloud Infrastructure Discovery	Discovery
T1552.005	Cloud Instance Metadata API	Credential-Access
T1530	Data from Cloud Storage	Collection
T1619	Cloud Storage Object Discovery	Discovery
T1078.004	Cloud Accounts	Defense-Evasion
T1548	Abuse Elevation Control Mechanism	Privilege-Escalation
T1562.008	Disable or Modify Cloud Logs	Defense-Evasion
T1537	Transfer Data to Cloud Account	Exfiltration
T1078	Valid Accounts	Defense-Evasion
T1190	Exploit Public-Facing Application	Initial-Access

Sources

Source	URL	Tier
Blog	https://www.crowdstrike.com/en-us/blog/crowdstrike-state-of-cdr-sur...	T1
CrowdStrike	https://www.crowdstrike.com/en-us/blog/new-in-falcon-cloud-security...	T1
CrowdStrike	https://www.crowdstrike.com/en-us/blog/crowdstrike-expands-real-tim...	T1
CrowdStrike	https://www.crowdstrike.com/en-us/blog/crowdstrike-advances-cnapp-w...	T1
Protect AI Development with Falcon Cloud Security | CrowdStrike	https://www.crowdstrike.com/en-us/blog/protect-ai-development-with-...	T1

DISCLAIMER

This intelligence report is produced by Tech Jacks Solutions Security Command Center (SCC) for informational purposes only. It does not constitute professional security advice, legal counsel, or an incident response engagement. The information herein is derived from publicly available sources and AI-assisted analysis; while every effort is made to ensure accuracy, Tech Jacks Solutions makes no warranties regarding completeness or timeliness.

Organizations should conduct their own validation and consult qualified security professionals before taking action based on this report. Tech Jacks Solutions is not liable for any damages resulting from the use of this information.

Generated 2026-06-30 14:46 UTC by TJS Security Command Center