INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-23 13:39 UTC

Zealot AI-Driven Cloud Attack Framework Demonstrates Autonomous Full-Chain Exploitation Outpacing Human Response

SECURITY ANALYSIS | HIGH | CVSS 5.0

SCC Item ID	SCC-STY-2026-0078
Type	Security Analysis
Severity	HIGH
CVSS Base Score	5.0
Affected Products	Cloud environments (multi-vendor; no specific platform scoped in available research summary)
Published	2026-04-23T06:00:00
Discovery Source	Rss

Executive Summary

Researchers have demonstrated 'Zealot,' a proof-of-concept AI-driven attack framework capable of executing complete cloud attack chains, from initial access through privilege escalation and lateral movement, autonomously and at speeds that research reporting suggests may outpace conventional SOC triage workflows. The significance is not a specific vulnerability but a capability threshold: AI-augmented adversaries may, according to early secondary reporting, complete attack objectives before defenders finish classifying the first alert. Note: Primary research publication has not yet been located; this summary is based on secondary coverage and should be validated against primary sources before determining organizational response scope. This signals a potential structural shift in the offense-defense timeline that organizations relying on human-paced incident response should assess in their threat models.

Technical Analysis

Based on currently available secondary reporting (see Sources section), which has not yet located primary research publication, the Zealot framework is reported as a proof-of-concept demonstration that AI-driven systems can chain cloud attack phases autonomously without the human deliberation that currently constrains adversary speed. The attack chain maps cleanly to established MITRE ATT&CK techniques: initial access via Valid Accounts (T1078) or exploitation of public-facing applications (T1190), followed by cloud resource and permission enumeration (T1069, T1087, T1580), privilege escalation through abuse of elevation control mechanisms (T1548), lateral movement across cloud environments, and mission completion through data

exfiltration from cloud storage (T1530) or resource hijacking for compute abuse (T1496). Defenders were also noted to contend with potential impairment of logging and monitoring capabilities (T1562), a standard technique to blind SOC visibility. The underlying weaknesses the framework exploits reflect systemic cloud hygiene gaps rather than novel vulnerabilities: improper access control (CWE-284), improper privilege management (CWE-269), and incorrect permission assignment for critical resources (CWE-732). These are not new weaknesses; they are persistent failures in cloud identity and permission hygiene that AI tooling can now exploit faster than humans can respond. Critically, according to available reporting, the system exhibited adaptive behavior during execution, suggesting the framework did not follow a rigid playbook but made autonomous decisions mid-chain. No specific cloud vendor, platform, or threat actor is attributed in the available source material. The source quality score of 0.44 reflects that primary sourcing is currently limited to T3 reporting; organizations should monitor for primary research publication before finalizing threat model updates. The CVSS base score of 5.0 provided in source data appears to be a placeholder pending formal CVE assignment (no CVE currently anchors this framework); the qualitative rating of 'high' reflects the capability-threat implications rather than CVSS metrics.

Action Checklist

1. Step 0 (Priority): Locate and review primary research. Current sourcing is limited to secondary reporting. Obtain the primary research paper, conference presentation, or vendor advisory before full implementation of Steps 1-5. These interim steps assume the capability is substantiated by primary research.
2. Step 1: Assess exposure, audit your cloud environments (AWS, Azure, GCP, or multi-cloud) for the permission and access control weaknesses described in Zealot research: over-permissioned roles, unused privileged accounts, and publicly accessible resources. These are the structural conditions autonomous attack frameworks would target.
3. Step 2: Review controls, verify that cloud-native CSPM tooling is actively enforcing least-privilege across IAM roles and service accounts; confirm that MFA is enforced on all cloud console and API access paths, including service principals and federated identities.
4. Step 3: Stress-test your response timeline. The core finding from research is that AI-driven attacks can complete before human triage concludes. Run a tabletop or purple team exercise measuring time-to-detect and time-to-contain for a privilege escalation chain in your cloud environment. Compare that timeline to the attack chain described in the research.
5. Step 4: Evaluate automated detection and response coverage. Identify where in the attack chain (T1078, T1548, T1580, T1530) your current detections fire and whether automated containment (not just alerting) is in place to act within the adversarial window.
6. Step 5: Update threat model. Incorporate AI-augmented, autonomous cloud attack chains as an explicit threat scenario in your threat register. Brief cloud operations, security engineering, and incident response teams on the capability-gap finding, not just the specific framework.
7. Step 6: Monitor for primary research publication. Current sourcing is limited to secondary reporting. Watch for the primary research paper or conference presentation for technical indicators, detection signatures, and tool-specific details.

IR / Forensic Enrichment

Triage Priority	URGENT

Escalation Criteria	Escalate immediately to CISO and cloud platform owners if CloudTrail, Azure Activity Log, or GCP Admin Activity logs show any principal executing T1078 (role assumption from anomalous source), T1548 (IAM policy modification), T1580 (broad resource enumeration), and T1530 (bulk storage object access) within a compressed timeframe (under 10 minutes), as this sequence matches the Zealot autonomous chain and may represent an active or completed compromise before human triage has concluded.
Recovery Notes	Post-containment, do not restore affected cloud IAM roles or service accounts to their prior permission state — rebuild them from a least-privilege baseline derived from the IAM Access Analyzer or GCP IAM Recommender output captured in Step 1. Monitor all reconstituted principals with CloudTrail data event logging enabled at the object level for a minimum of 30 days post-recovery, with automated alerting on any API call sequence matching the T1078 → T1548 → T1580 → T1530 chain. Because Zealot's autonomous operation means lateral movement and data access may have completed before containment, conduct a full data exposure assessment of any cloud storage objects accessible to compromised principals before declaring recovery complete.
Forensic Artifacts	AWS CloudTrail Management Events — filter for AssumeRole, AttachRolePolicy, PutRolePolicy, CreatePolicyVersion, and GetCallerIdentity API calls from the same principal ARN within compressed timeframes; the sequential, low-latency pattern of these calls (milliseconds to seconds between steps) is a distinguishing artifact of autonomous tooling versus human operator activity | AWS CloudTrail Data Events on S3 — enable and review GetObject and ListObjects calls on all buckets accessible to workload identities or service accounts; bulk enumeration across multiple buckets in a single session is consistent with T1530 execution by Zealot's discovery module | Azure Entra ID Sign-In Logs and Azure Activity Log — query for Microsoft.Authorization/roleAssignments/write and Microsoft.Authorization/roleDefinitions/write events correlated with service principal authentication events lacking MFA (authenticationRequirement: singleFactorAuthentication); automated tooling exploiting missing MFA on service principals will produce this pattern | GCP Cloud Audit Logs — Admin Activity log entries for SetIamPolicy, CreateServiceAccountKey, and storage.objects.list operations attributed to the same service account identity within a single session; cross-reference with GCP VPC Flow Logs for any outbound data transfer to external IPs following the storage enumeration sequence | Cloud provider IAM credential last-used metadata — AWS IAM credential report fields `access_key_1_last_used_date` and `password_last_used`; Azure Entra ID `signInActivity` per service principal; GCP service account `last_authenticated_time` — anomalous first-use of long-dormant credentials or access keys immediately preceding the attack chain is consistent with Zealot's valid account abuse entry technique (T1078) and should be treated as a high-confidence precursor artifact

Per-Action IR Details

Step 1: Assess exposure — audit your cloud environments (AWS, Azure, GCP, or multi-cloud) for the permission and access control weaknesses Zealot targets: over-permissioned roles, unused privileged accounts, and publicly accessible resources. These are the structural conditions the framework exploits.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establish IR capability and reduce attack surface prior to incident

Controls: NIST AC-2 (Account Management) — enumerate and right-size cloud IAM roles and service accounts, NIST AC-6 (Least Privilege) — enforce minimum necessary permissions on all cloud principals, NIST RA-3 (Risk Assessment) — document over-permissioned role findings as risk register entries, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — extend to cloud IAM principals and service accounts, CIS 5.1 (Establish and Maintain an Inventory of Accounts) — include cloud console, API keys, service principals, and federated identities, CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts)

Compensating: For teams without CSPM licensing: run AWS IAM Access Analyzer via CLI (`aws accessanalyzer list-findings --analyzer-name `) to surface externally accessible resources and unused permissions at no cost. On Azure, use `az role assignment list --all` and cross-reference against Azure AD sign-in logs exported to a local file. On GCP, use the free Policy Analyzer (`gcloud asset analyze-iam-policy`) to enumerate effective permissions. Pipe each output to a CSV and flag any role with wildcard actions (`*`) or any service account inactive for 45+ days — these match the structural preconditions Zealot requires for autonomous escalation.

Evidence: Before remediating, capture a point-in-time snapshot of the current IAM state as forensic baseline: export AWS IAM credential report (`aws iam generate-credential-report && aws iam get-credential-report`), Azure AD audit logs (sign-in and role assignment logs from Entra ID), and GCP Cloud Audit Logs — Admin Activity log (`cloudaudit.googleapis.com/activity`) for all `SetIamPolicy` and `CreateServiceAccountKey` events from the past 90 days. This baseline establishes whether any over-permissioned roles were created or modified during a potential reconnaissance or staging period prior to discovery.

Step 2: Review controls — verify that cloud-native CSPM tooling is actively enforcing least-privilege across IAM roles and service accounts; confirm that MFA is enforced on all cloud console and API access paths, including service principals and federated identities.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Verify defensive tools are operational before adversarial window opens

Controls: NIST AC-6 (Least Privilege) — validate enforcement is active, not merely configured, NIST IA-5 (Authenticator Management) — confirm MFA is enforced on all cloud API and console authentication paths, not just human interactive logins, NIST SI-4 (System Monitoring) — verify CSPM policies are triggering alerts, not silently suppressed, NIST CM-6 (Configuration Settings) — validate CSPM baseline configuration matches current cloud resource state, CIS 6.3 (Require MFA for Externally-Exposed Applications) — include cloud console endpoints, CIS 6.4 (Require MFA for Remote Network Access), CIS 6.5 (Require MFA for Administrative Access) — explicitly include cloud service principals and federated SAML/OIDC identities

Compensating: Without paid CSPM: use AWS Config with the free managed rules `iam-user-mfa-enabled`, `iam-root-access-key-check`, and `access-keys-rotated` — enable via `aws configservice put-config-rule`. On Azure, enable Microsoft Defender for Cloud free tier and review the 'Secure Score' recommendations for identity controls. On GCP, enable the Security Command Center standard tier and run the IAM recommender: `gcloud recommender recommendations list --recommender=google.iam.policy.Recommender`. For federated identities, manually audit SAML trust relationships and OAuth 2.0 app grants — Zealot's autonomous chain depends on finding at least one credential path without MFA enforcement.

Evidence: Before any control changes, export current MFA enforcement state as evidence: AWS — `aws iam list-virtual-mfa-devices` and `aws iam list-users` cross-referenced for gaps; Azure — Entra ID Conditional Access policy export and sign-in logs filtered for `authenticationRequirement: singleFactorAuthentication` on privileged roles; GCP — `gcloud iam service-accounts list` with `gcloud iam service-accounts get-iam-policy` per account. Retain these exports as pre-remediation artifacts to establish whether MFA gaps existed during the adversarial window of interest.

Step 3: Stress-test your response timeline — the core finding is that AI-driven attacks can complete before human triage concludes. Run a tabletop or purple team exercise measuring time-to-detect and time-to-contain for a privilege escalation chain in your cloud environment. Compare that timeline to the attack chain described in the research.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Test IR capability effectiveness; validate detection and response timelines against realistic threat scenarios

Controls: NIST IR-3 (Incident Response Testing) — conduct exercises specifically simulating autonomous, multi-stage cloud privilege escalation chains at machine speed, NIST IR-4 (Incident Handling) — measure preparation, detection, containment phase durations against Zealot's demonstrated full-chain execution timeline, NIST IR-2 (Incident Response Training) — train SOC analysts to recognize that initial alert classification may already represent a completed attack chain, not an in-progress one, NIST CA-8 (Penetration Testing) — use purple team to simulate T1078, T1548, T1580, T1530 in sequence against cloud environment, CIS 7.1 (Establish and Maintain a Vulnerability

Management Process) — incorporate AI-speed adversarial timelines into exercise design criteria

Compensating: For a 2-person team without a red team budget: design a tabletop using the MITRE ATT&CK Cloud matrix as the scenario script — walk through T1078 (Valid Accounts) → T1548 (Abuse Elevation Control Mechanism) → T1580 (Cloud Infrastructure Discovery) → T1530 (Data from Cloud Storage Object) as a linear chain. Record the elapsed wall-clock time at each detection decision point using only your current tooling. Compare to the Zealot research timeline once the primary paper is published (Step 6). Document gaps where human decision latency exceeds the inter-step adversarial interval — these are your automated response requirements.

Evidence: During the exercise, instrument the following to capture baseline detection latency: timestamp of first alert generated in your cloud provider's native alerting (AWS GuardDuty finding, Azure Defender alert, or GCP Security Command Center finding); timestamp of first human acknowledgment in your ticketing or SIEM system; timestamp of first containment action taken. Capture these three timestamps per attack chain step. The delta between cloud-provider alert timestamp and human containment action timestamp is the adversarial dwell window Zealot is designed to operate within — document this gap as a formal finding.

Step 4: Evaluate automated detection and response coverage — identify where in the Zealot attack chain (T1078, T1548, T1580, T1530) your current detections fire and whether automated containment (not just alerting) is in place to act within the adversarial window.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Identify detection gaps against known adversarial technique sequences; validate that detection coverage maps to complete attack chain, not isolated events

Controls: NIST SI-4 (System Monitoring) — validate monitoring covers all four ATT&CK techniques in Zealot's documented chain: T1078, T1548, T1580, T1530, NIST IR-4 (Incident Handling) — confirm automated response playbooks exist and are active, not just alerting pipelines, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — verify cloud audit log analysis is occurring at machine speed, not human-reviewed batches, NIST AU-2 (Event Logging) — confirm logging is enabled for IAM role assumption, privilege escalation API calls, storage object enumeration, and cross-account access attempts, CIS 8.2 (Collect Audit Logs) — validate cloud audit logs are collected and retained centrally, not only in provider-native consoles where retention windows may be short

Compensating: For each ATT&CK technique in the Zealot chain, deploy the following free detection controls: T1078 — enable AWS CloudTrail `ConsoleLogin` and `AssumeRole` event monitoring; alert on `AssumeRoleWithWebIdentity` from unusual source IPs using a CloudWatch Metric Filter. T1548 — alert on AWS `AttachRolePolicy`, `PutRolePolicy`, and `CreatePolicyVersion` API calls via CloudTrail + CloudWatch Events rule; Azure equivalent: monitor `Microsoft.Authorization/roleAssignments/write` in Activity Log. T1580 — alert on `DescribeInstances`, `ListBuckets`, `GetCallerIdentity` in sequence from the same principal within a 5-minute window — indicative of automated enumeration. T1530 — alert on `GetObject` calls on S3 buckets not accessed by that principal in the prior 30 days via CloudTrail data events. Use free Sigma rules from the SigmaHQ cloud-aws ruleset (github.com/SigmaHQ/sigma) converted to your SIEM query language via sigma-cli.

Evidence: Before modifying detection rules, export current alerting state as a coverage map: pull the last 30 days of GuardDuty findings (AWS), Defender for Cloud alerts (Azure), or Security Command Center findings (GCP) and classify each finding against the four Zealot ATT&CK technique IDs — T1078, T1548, T1580, T1530. Identify which techniques produced zero alerts in that period; absence of findings for T1580 (Cloud Infrastructure Discovery) or T1530 (Data from Cloud Storage) is a critical gap given Zealot's reconnaissance dependency on these techniques. Retain this coverage gap analysis as a pre-remediation artifact.

Step 5: Update threat model — incorporate AI-augmented, autonomous cloud attack chains as an explicit threat scenario in your threat register. Brief cloud operations, security engineering, and incident response teams on the capability-gap finding, not just the specific framework.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Update policies and threat models based on lessons learned and emerging capability intelligence; this step applies proactively given the capability-gap nature of the threat

Controls: NIST IR-8 (Incident Response Plan) — update IR plan to explicitly address adversarial timelines that may complete before human triage concludes, requiring pre-authorized automated containment authorities, NIST RA-3

(Risk Assessment) — add AI-augmented autonomous cloud attack chains as a named threat scenario with likelihood and impact ratings, NIST PM-16 (Threat Awareness Program) — formally brief cloud ops, security engineering, and IR on the Zealot capability-gap finding as a program-level threat awareness action, NIST SI-5 (Security Alerts, Advisories, and Directives) — track the Zealot research publication as a pending advisory requiring IR plan update upon release, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — update process documentation to include AI-speed adversarial tooling as an explicit risk category, CIS 7.2 (Establish and Maintain a Remediation Process) — add automated response authority decisions to the remediation process for cloud IAM abuse scenarios

Compensating: For a 2-person team: create a one-page threat scenario card for AI-augmented autonomous cloud attack chains covering: (1) threat actor capability — autonomous full-chain from initial access to data access without human operator interaction; (2) target conditions — over-permissioned IAM roles, no MFA on service accounts, cloud storage accessible to workload identities; (3) detection challenge — attack may complete within a single human triage cycle; (4) required IR response change — containment authority must be pre-delegated to automated controls, not require human approval. Add this card to your threat register using a plain-text file under version control (git) if no GRC platform is available. Share via internal wiki or email to cloud ops and engineering leads — briefing does not require tooling.

Evidence: As input to the threat model update, collect and retain: the secondary research reporting on Zealot that triggered this assessment (URLs, publication dates, author attributions); the results of the coverage gap analysis from Step 4; the response timeline deltas documented in Step 3's tabletop exercise; and the IAM permission and MFA gap findings from Steps 1 and 2. These artifacts constitute the threat intelligence basis for the model update and support audit defensibility under NIST RA-3 documentation requirements.

Step 6: Monitor for primary research publication — current sourcing is limited to secondary reporting. Watch for the primary research paper or conference presentation for technical indicators, detection signatures, and tool-specific details.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Integrate cyber threat intelligence and contextual information into adverse event analysis; update detection baselines as authoritative technical detail becomes available

Controls: NIST SI-5 (Security Alerts, Advisories, and Directives) — establish a formal watch process for the primary Zealot research publication from the authoring research team or conference (e.g., DEF CON, Black Hat, Usenix Security, academic preprint servers), NIST IR-5 (Incident Monitoring) — update incident tracking documentation with a pending intelligence item for Zealot primary research; assign an owner and review cadence, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — plan to update cloud audit log query filters and correlation rules immediately upon publication of technical indicators of compromise or tool-specific signatures, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — treat the primary research publication as a triggering event for a vulnerability management process review cycle

Compensating: For a 2-person team: set up a free Google Alert or RSS feed monitoring for terms 'Zealot AI attack framework,' 'autonomous cloud exploitation,' and the names of any attributed research authors identified in secondary reporting. Monitor arXiv.org (cs.CR category), USENIX Security proceedings, Black Hat Briefings archive, and DEF CON CFP publications on a weekly cadence. When the primary paper publishes, immediately extract: (1) any YARA or Sigma rules released by the authors; (2) specific API call sequences or cloud SDK invocations used in the attack chain; (3) any IOCs tied to the proof-of-concept tooling. Convert Sigma rules to your query language via sigma-cli and deploy to CloudWatch Logs Insights or Azure Monitor Logs within 24 hours of publication.

Evidence: Upon primary research publication, immediately extract and preserve: the full list of cloud API calls invoked by Zealot in sequence (these become the basis for CloudTrail / Activity Log detection queries); any hardcoded strings, user-agent patterns, or SDK fingerprints present in the proof-of-concept tool; and any persistence mechanisms or credential staging artifacts described in the paper. Retroactively query your cloud audit logs for these indicators across the 90-day window preceding the publication date to determine whether any reconnaissance activity matching the Zealot technique chain occurred in your environment before detection capability was established.

Detection Guidance

Contingent on primary research publication, detection strategies should focus on the behavioral sequences likely to characterize AI-augmented autonomous cloud attack chains, whether or not Zealot-specific instances are observed. These are defensive improvements that apply to the broader threat class, independent of this specific framework. Focus detection engineering on the behavioral sequence rather than static indicators, since no verified IOCs are available from current sourcing. Key detection hypotheses to build or validate: (1) Rapid enumeration chains, high-velocity IAM and permissions queries (T1069, T1087) from a single identity or service account within a compressed time window, particularly following any new authentication event. (2) Privilege escalation following enumeration, any attempt to assume a higher-privileged role (T1548) within minutes of permissions discovery activity; cloud audit logs (CloudTrail, Azure Activity Log, GCP Cloud Audit Logs) should surface this sequence. (3) Logging impairment (T1562), any modification or disabling of cloud logging configurations, especially GuardDuty, Defender for Cloud, or GCP Cloud Audit Logs; treat this as a high-confidence attack signal requiring immediate response. (4) Cross-service lateral movement, access to cloud storage buckets (T1530) or compute resources (T1496) by identities that have no prior access history to those services. (5) AI-paced execution cadence, traditional human-operated attack chains include pause intervals; autonomous chains may show suspiciously uniform or near-instantaneous timing between steps. Behavioral analytics tuned to inter-event timing may surface this pattern. Log sources to prioritize: cloud control plane audit logs, IAM activity logs, network flow logs for east-west cloud traffic, and CSPM alert queues. Audit for gaps in logging coverage; T1562 exploitation means defenders may face incomplete logs after the fact.

Indicators of Compromise

Type	Value	Context	Confidence
TOOL	`Pending — refer to primary Zealot research publication for published indicators`	The Zealot framework is described in secondary reporting as an AI-driven cloud attack chain tool. Technical indicators including tool signatures, behavioral patterns, and any associated artifacts have not been extracted in available source material. Monitor for the primary research paper or conference presentation.	LOW

Framework Mappings

MITRE-ATTACK

- **T1078** — Valid Accounts
- **T1190** — Exploit Public-Facing Application
- **T1069** — Permission Groups Discovery
- **T1087** — Account Discovery
- **T1548** — Abuse Elevation Control Mechanism
- **T1580** — Cloud Infrastructure Discovery
- **T1530** — Data from Cloud Storage
- **T1562** — Impair Defenses

- **T1496** — Resource Hijacking

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **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
- **CM-6** — Configuration Settings
- **AU-9** — Protection of Audit Information
- **SI-4** — System Monitoring
- **AC-3** — Access Enforcement

OWASP-TOP10-2021

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

CIS-V8

- **3.3** — Configure Data Access Control Lists
- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **5.4** — Restrict Administrator Privileges to Dedicated Administrator Accounts
- **6.8** — Define and Maintain Role-Based Access Control
- **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
- **CC6.3** — Authorizes, modifies, or removes access

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

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1078	Valid Accounts	Defense-Evasion
T1190	Exploit Public-Facing Application	Initial-Access
T1069	Permission Groups Discovery	Discovery
T1087	Account Discovery	Discovery
T1548	Abuse Elevation Control Mechanism	Privilege-Escalation
T1580	Cloud Infrastructure Discovery	Discovery
T1530	Data from Cloud Storage	Collection
T1562	Impair Defenses	Defense-Evasion
T1496	Resource Hijacking	Impact

Sources

Source	URL	Tier
Security News	https://www.darkreading.com/cyber-risk/zealot-shows-ai-execute-full...	T3
Vendor refuses CVEs for third-party findings. Anything you can do?	https://www.reddit.com/r/cybersecurity/comments/1spov5s/vendor_refu...	T3
Lack of vendor visibility is number one pain point for cloud customers	https://www.cloudcomputing-news.net/news/lack-vendor-visibility-num...	T3
Top 11 Cloud Security Vulnerabilities and How to Fix Them - Wiz	https://www.wiz.io/academy/cloud-security/common-cloud-vulnerabilities	T3
Top 15 Cloud Security Vulnerabilities - SentinelOne	https://www.sentinelone.com/cybersecurity-101/cloud-security/cloud-...	T3

Generated 2026-04-23 13:39 UTC by TJS Security Command Center