INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-15 06:52 UTC

Bring out your dead: How agentic AI for cybersecurity helps you rid your cloud of forgotten, risky assets

SECURITY ANALYSIS | MEDIUM

SCC Item ID	SCC-STY-2026-0130
Type	Security Analysis
Severity	MEDIUM
Affected Products	Cloud environments broadly (multi-cloud and hybrid cloud infrastructure)
Published	2026-05-15
Discovery Source	Gemini

Executive Summary

Cloud environments accumulate forgotten, unmanaged assets at scale, abandoned virtual machines, orphaned storage buckets, stale IAM roles, idle cloud functions, and these 'zombie assets' quietly expand the attack surface without triggering conventional security controls. Attackers can discover and exploit these resources through cloud enumeration techniques, often accessing sensitive data or establishing lateral movement footholds before defenders are aware the asset exists. Agentic AI is demonstrating practical value as an answer to the visibility gap, autonomously discovering and classifying these assets at a speed and scale that manual audit cycles cannot match, a signal that cloud hygiene is becoming an AI-augmented discipline.

Technical Analysis

The core risk is structural. Cloud provisioning is fast and distributed; decommissioning is slow, inconsistent, and often skipped entirely. The result is a category of assets that NIST SP 800-53 revision 5 addresses under CM-8 (Information System Component Inventory) and RA-5 (Vulnerability Monitoring and Scanning), controls that assume an accurate, maintained asset inventory exists. In practice, that assumption fails routinely in multi-cloud and hybrid environments.

The attack surface created by zombie assets maps directly to four MITRE ATT&CK techniques: T1526 (Cloud Service Discovery), T1538 (Cloud Service Dashboard), T1530 (Data from Cloud Storage), and T1078 (Valid Accounts). An attacker who discovers an unmanaged VM may find it running an unpatched OS image from the last time it was actively managed, months or years ago. An orphaned storage bucket may retain a permissive bucket policy applied during development that was never tightened for production and never cleaned up after

the project ended. A stale IAM role may carry permissions that were broad during initial integration work and were never scoped down.

These conditions align with three CWEs flagged in the item data: CWE-284 (Improper Access Control), CWE-732 (Incorrect Permission Assignment for Critical Resource), and CWE-1059 (Insufficient Technical Documentation), the last of which is an underappreciated enabler. Zombie assets persist partly because they were never properly documented; no ticket records their purpose, owner, or intended decommission date.

Agentic AI approaches the problem differently than scheduled scans or manual audit cycles. Rather than producing a point-in-time inventory report, an agentic system maintains continuous discovery across cloud APIs, classifies assets by risk posture (age, privilege level, network exposure, data sensitivity signals), and can execute or queue remediation actions without waiting for a human to process a report. The capability shift enables continuous rather than point-in-time asset visibility: manual cloud hygiene programs typically run quarterly or annually and review findings over weeks; agentic systems can compress that cycle to near-real-time.

The tradeoff is governance. Agentic AI with remediation authority introduces a new risk class, automated deletion or access revocation of an asset that is not actually abandoned but simply undocumented. Security teams implementing agentic cloud hygiene tools need clear policy boundaries defining which remediation actions require human approval versus which can execute autonomously, and they need logging sufficient to reconstruct any automated action for audit purposes. This maps back to NIST SP 800-53 AU-2 and AU-12 (Audit Events, Audit Record Generation) and CM-3 (Configuration Change Control).

The vendor sources cited (CrowdStrike, SentinelOne, Fidelis, Pentera) are tier-3 commercial sources and should be read as marketing-adjacent perspectives on a real underlying problem. The core risk model they describe is consistent with CISA cloud security guidance and the CSA Cloud Controls Matrix, which provide more authoritative framing.

Action Checklist

1. Step 1: Assess exposure, run a full cloud asset discovery across all accounts, subscriptions, and projects in every cloud environment your organization uses; include regions that may have been provisioned during past projects or by shadow IT
2. Step 2: Review controls, verify that CM-8 (asset inventory) and RA-5 (vulnerability scanning) controls in your NIST SP 800-53 implementation actually cover cloud resources, not just on-premises infrastructure; confirm IAM role and policy reviews are included in your cloud hygiene scope
3. Step 3: Update threat model, add cloud enumeration (T1526, T1538) and stale credential abuse (T1078) to your cloud threat register; model the scenario where an attacker discovers an unmanaged asset before your team does
4. Step 4: Communicate findings, brief leadership on the gap between cloud provisioning speed and decommission discipline; quantify the number of unmanaged assets discovered and map at least a sample to potential data exposure or privilege risk
5. Step 5: Monitor developments, track CISA cloud security guidance, CSA Cloud Controls Matrix updates, and emerging NIST guidance on agentic AI governance; agentic remediation tooling is maturing quickly and governance frameworks are still forming

IR / Forensic Enrichment

Triage Priority	STANDARD
Escalation Criteria	Escalate immediately to incident response if CloudTrail, Azure Activity Log, or GCP Audit Logs reveal that any discovered unmanaged asset has already received external enumeration API calls (T1526, T1538) or if a stale IAM role (T1078) shows an AssumeRole event from an unrecognized principal or external IP within the prior 90 days, indicating the attacker-first scenario modeled in Step 3 may have already occurred; additionally escalate if any orphaned storage asset is confirmed to contain PII, PHI, or credentials, triggering potential regulatory breach notification obligations.
Recovery Notes	After decommissioning or securing zombie assets, re-run the full cloud asset discovery (Step 1 tooling) within 72 hours to confirm no new unmanaged assets were provisioned during the remediation window and that no orphaned resources were missed. Monitor CloudTrail and equivalent logs for 30 days post-remediation for any resumed enumeration activity targeting the now-removed assets — an attacker who bookmarked those endpoints may retry and reveal themselves through failed API calls against resources that no longer exist. Establish a recurring monthly cloud hygiene review using the same CLI-based discovery approach as a lightweight compensating control until a formal CSPM solution is budgeted and deployed.
Forensic Artifacts	AWS CloudTrail logs (all regions, 90-day lookback): filter for ListBuckets, DescribeInstances, ListRoles, ListFunctions, and AssumeRole events from IAM principals not present in current active IAM inventory — these API calls are the primary forensic signal for T1526 (Cloud Account Discovery) and T1538 (Cloud Service Dashboard) enumeration activity against zombie assets | AWS IAM credential report and GetAccountAuthorizationDetails export: documents stale access keys and roles with no RoleLastUsed timestamp — establishes which T1078 (Valid Accounts) vectors were exposed and for how long, and provides the pre-remediation evidence chain if a stale credential is later found to have been used maliciously | S3 bucket ACL and bucket policy exports for all buckets lacking owner tags or last-modified activity: captures the public-access exposure state at the time of discovery, essential for establishing breach notification timelines if PII or credentials are found in orphaned buckets | Cloud resource tag audit export (all asset types, all regions): the absence of owner, project, or environment tags on a resource is itself forensic evidence of the decommission discipline gap described in the threat summary — documents which assets fell outside any governance process and for how long | AWS Config historical snapshots or equivalent Azure Resource Graph query history: reveals the timeline of when unmanaged assets were created versus when they fell out of active management, and identifies gaps in Config rule coverage (regions or resource types not tracked) that allowed zombie assets to persist outside CM-8 and RA-5 control scope

Per-Action IR Details

Step 1: Assess exposure — run a full cloud asset discovery across all accounts, subscriptions, and projects in every cloud environment your organization uses; include regions that may have been provisioned during past projects or by shadow IT

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR Capability and Asset Visibility

Controls: NIST CM-8 (System Component Inventory), NIST RA-5 (Vulnerability Monitoring and Scanning), NIST CA-7 (Continuous Monitoring), CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 1.2 (Address Unauthorized Assets)

Compensating: For AWS: run 'aws ec2 describe-instances --region ' and 'aws s3api list-buckets' across every region programmatically using a bash loop over the output of 'aws ec2 describe-regions --query Regions[].RegionName'. For Azure: use 'az resource list --output table' across all subscriptions via 'az account list'. For GCP: use 'gcloud asset search-all-resources --scope=organizations/ORG_ID'. Aggregate results into a CSV; flag any asset with no owner tag, no last-modified activity in 90+ days, or a creation date tied to a terminated project. ScoutSuite (free, open-source) can run multi-cloud enumeration and export findings without requiring a SIEM.

Evidence: Before running discovery, capture a point-in-time snapshot of cloud provider audit logs to establish baseline: AWS CloudTrail (all regions, S3-backed trail), Azure Activity Log export, GCP Cloud Audit Logs (Admin Activity and Data Access). Preserve these logs before any remediation so you can later determine whether an attacker already performed cloud enumeration (T1526 — Cloud Account Discovery, T1538 — Cloud Service Dashboard enumeration) against the same forgotten assets you are about to find. Look specifically for ListBuckets, DescribeInstances, ListFunctions, ListRoles API calls from unexpected IAM principals or external IPs in the 30–90 days prior to this assessment.

Step 2: Review controls — verify that CM-8 (asset inventory) and RA-5 (vulnerability scanning) controls in your NIST SP 800-53 implementation actually cover cloud resources, not just on-premises infrastructure; confirm IAM role and policy reviews are included in your cloud hygiene scope

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Policy, Control Coverage, and Tool Readiness

Controls: NIST CM-8 (System Component Inventory), NIST RA-5 (Vulnerability Monitoring and Scanning), NIST IR-8 (Incident Response Plan), NIST AC-2 (Account Management), NIST AC-6 (Least Privilege), CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 5.1 (Establish and Maintain an Inventory of Accounts)

Compensating: Export all IAM roles and attached policies using 'aws iam get-account-authorization-details --output json > iam_snapshot.json' and parse for roles with no last-used date (field: RoleLastUsed) or last-used older than 90 days using a simple Python script with the boto3 library. For Azure: 'az role assignment list --all --output json' filtered for service principals assigned to deleted or non-existent resources. For GCP: 'gcloud iam service-accounts list' cross-referenced against active projects. Flag any IAM role with AdministratorAccess or wildcard (*) actions attached to a compute or storage resource that is itself orphaned. PMapper (free, open-source) can visualize IAM privilege escalation paths from stale roles without requiring a commercial CSPM.

Evidence: Collect AWS IAM credential report ('aws iam generate-credential-report; aws iam get-credential-report') to identify access keys with no recent use and roles with no RoleLastUsed timestamp — these are the exact stale credentials (T1078 — Valid Accounts) an attacker would target if they discovered them via cloud enumeration before your team did. Also pull AWS Config historical snapshots if enabled; gaps in Config coverage (regions or resource types not tracked) directly reveal which assets were outside your CM-8 scope. Document this gap as forensic evidence of the control failure, not just a remediation task.

Step 3: Update threat model — add cloud enumeration (T1526, T1538) and stale credential abuse (T1078) to your cloud threat register; model the scenario where an attacker discovers an unmanaged asset before your team does

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Threat Modeling and Attack Scenario Development

Controls: NIST RA-3 (Risk Assessment), NIST SI-4 (System Monitoring), NIST IR-4 (Incident Handling), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: Model the attacker-first scenario concretely: document that T1526 (Cloud Account Discovery) involves API calls such as ListAccounts, DescribeOrganization, ListBuckets, and ListInstances that will appear in CloudTrail/Activity Logs before any exploitation occurs — these are your earliest detectable signals. Write a Sigma rule targeting CloudTrail events where userIdentity.type is 'IAMUser' or 'AssumedRole' and eventName matches ListBuckets, DescribeInstances, or ListRoles with a sourceIPAddress outside known corporate CIDR ranges. Publish this Sigma rule to your detection backlog and test it against the historical CloudTrail logs captured in Step 1. Threat

model document can be maintained as a simple markdown file in version control; no SIEM required to document the model, only to operationalize the detections.

Evidence: To validate whether the threat scenario has already materialized, query CloudTrail for the 90 days prior to this assessment: filter on eventSource = 's3.amazonaws.com' with eventName = 'ListBuckets' OR eventSource = 'ec2.amazonaws.com' with eventName = 'DescribeInstances' from principals that do not appear in your current active IAM inventory (i.e., principals tied to roles or users that are themselves orphaned or stale). A principal you cannot identify in your current IAM state performing enumeration calls is a confirmed T1526 indicator. Preserve these CloudTrail records as forensic evidence before any IAM cleanup removes the ability to trace them.

Step 4: Communicate findings — brief leadership on the gap between cloud provisioning speed and decommission discipline; quantify the number of unmanaged assets discovered and map at least a sample to potential data exposure or privilege risk

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Lessons Learned and Risk Communication

Controls: NIST IR-6 (Incident Reporting), NIST IR-8 (Incident Response Plan), NIST RA-3 (Risk Assessment), NIST PM-9 (Risk Management Strategy), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: Structure the leadership brief around concrete risk figures derived from the Step 1 discovery output: number of unmanaged assets by type (VMs, S3/Blob buckets, IAM roles, Lambda/Cloud Functions), number of those with public-facing exposure (e.g., S3 buckets with public ACLs — query with 'aws s3api get-bucket-acl --bucket ' or use AWS Trusted Advisor free tier), and number of stale IAM roles with high-privilege policies still attached. Map at least 3–5 sample assets to a concrete impact scenario (e.g., 'This orphaned S3 bucket contains backup files tagged with PII and has no bucket policy — it is accessible to any authenticated AWS principal in the account'). Use a simple risk matrix (likelihood x impact) with no special tooling required.

Evidence: Before briefing leadership, preserve point-in-time evidence of the exposure state: export the S3 bucket ACL and policy configurations for any publicly accessible buckets identified, screenshot or export the IAM role last-used data, and capture the resource tags (or absence thereof) showing which assets lack owner or project attribution. This evidence documents the pre-remediation exposure window and is essential if a later investigation reveals an attacker accessed one of these assets during the gap period — it establishes the timeline for potential breach notification obligations under regulations such as GDPR, HIPAA, or state breach notification laws. Worth noting this touches regulatory notification thresholds — you may want to verify with legal counsel whether any discovered exposed assets containing PII or PHI trigger mandatory reporting obligations.

Step 5: Monitor developments — track CISA cloud security guidance and CSA research on agentic AI governance for cloud environments; agentic remediation tooling is maturing quickly and governance frameworks are still forming

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Policy Updates and Continuous Improvement

Controls: NIST SI-5 (Security Alerts, Advisories, and Directives), NIST IR-8 (Incident Response Plan), NIST CA-7 (Continuous Monitoring), NIST PM-9 (Risk Management Strategy), CIS 7.1 (Establish and Maintain a Vulnerability Management Process)

Compensating: Subscribe to CISA's free email alerts at cisa.gov/uscert/ncas/alerts and to the CSA Cloud Controls Matrix working group updates. Assign one team member to review these on a biweekly cadence and log any new cloud enumeration TTPs or agentic AI governance guidance into the threat register updated in Step 3. Until formal agentic AI governance frameworks mature, apply a conservative control principle: any automated remediation agent (whether commercial or open-source) must operate in read-only/audit mode by default and require explicit human approval before modifying or deleting cloud assets — document this as a standing policy in your IR plan under NIST 800-61r3 §2 preparation requirements. Track MITRE ATT&CK cloud matrix updates specifically for additions to the Discovery tactic (TA0007) that affect T1526 and T1538.

Evidence: Establish a monitoring baseline now so future changes are detectable: enable AWS Config with a conformance pack covering S3 public access, IAM password policy, and EC2 security group rules; for Azure, enable Microsoft Defender for Cloud free tier recommendations as a detective control. Document the current count of

unmanaged assets, stale IAM roles, and publicly exposed buckets as the remediation baseline. If agentic tooling is later introduced, its audit logs (what it enumerated, what it modified, under whose authorization) become primary forensic artifacts and must be retained under NIST AU-11 (Audit Record Retention) policy from day one of deployment.

Detection Guidance

Detection for zombie asset abuse centers on cloud API activity rather than endpoint telemetry. Review CloudTrail (AWS), Unified Audit Log (Azure), and Cloud Audit Logs (GCP) for enumeration patterns: repeated DescribeInstances, ListBuckets, ListRoles, or equivalent calls from principals with no recent operational history. Flag IAM role assumption events where the assuming principal has not used that role in 90 or more days. Alert on any access to S3 buckets, Azure Blob containers, or GCP Cloud Storage buckets that have had no legitimate access in 60 or more days, data access on a forgotten bucket is a high-signal anomaly. For stale VM access, monitor for authentication events against instances that have not received patch updates in 90 or more days; these are both hygiene indicators and potential compromise signals. On the hunting side, query your CSPM or cloud inventory tool for IAM roles with AdministratorAccess or equivalent that are not attached to any active principal or service, these represent standing privilege that exists purely due to oversight. Audit cloud functions and serverless workloads for execution events; abandoned functions that suddenly execute are worth investigating. Log retention for cloud API activity should meet at minimum the 90-day availability threshold recommended in NIST SP 800-92.

Framework Mappings

MITRE-ATTACK

- **T1526** — Cloud Service Discovery
- **T1538** — Cloud Service Dashboard
- **T1530** — Data from Cloud Storage
- **T1078** — Valid Accounts

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **AC-3** — Access Enforcement

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
- **7.3** — Perform Automated Operating System Patch Management

- **7.4** — Perform Automated Application Patch Management

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

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.23** — Information security for use of cloud services

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1526	Cloud Service Discovery	Discovery
T1538	Cloud Service Dashboard	Discovery
T1530	Data from Cloud Storage	Collection
T1078	Valid Accounts	Defense-Evasion

Sources

Source	URL	Tier
Top 8 Cloud Vulnerabilities | CrowdStrike	https://www.crowdstrike.com/en-us/cybersecurity-101/cloud-security/...	T3
17 Security Risks of Cloud Computing in 2026 - SentinelOne	https://www.sentinelone.com/cybersecurity-101/cloud-security/securi...	T3
Vulnerability Scanning from IT Assets to Cloud Environments	https://fidelissecurity.com/threatgeek/threats-and-vulnerabilities/...	T3
Top 10 Cloud Security Vulnerabilities (And How to Fix Them)	https://www.secure.com/blog/infrastructure-security/cloud-security-...	T3
What Is Cloud Vulnerability Identification? - Pentera	https://pentera.io/glossary/cloud-vulnerability-identification/	T3

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 2026-05-15 06:52 UTC by TJS Security Command Center