Google Cloud

Security

Cloud Threat Horizons Report

H1 2026

Office of the CISO

Contents

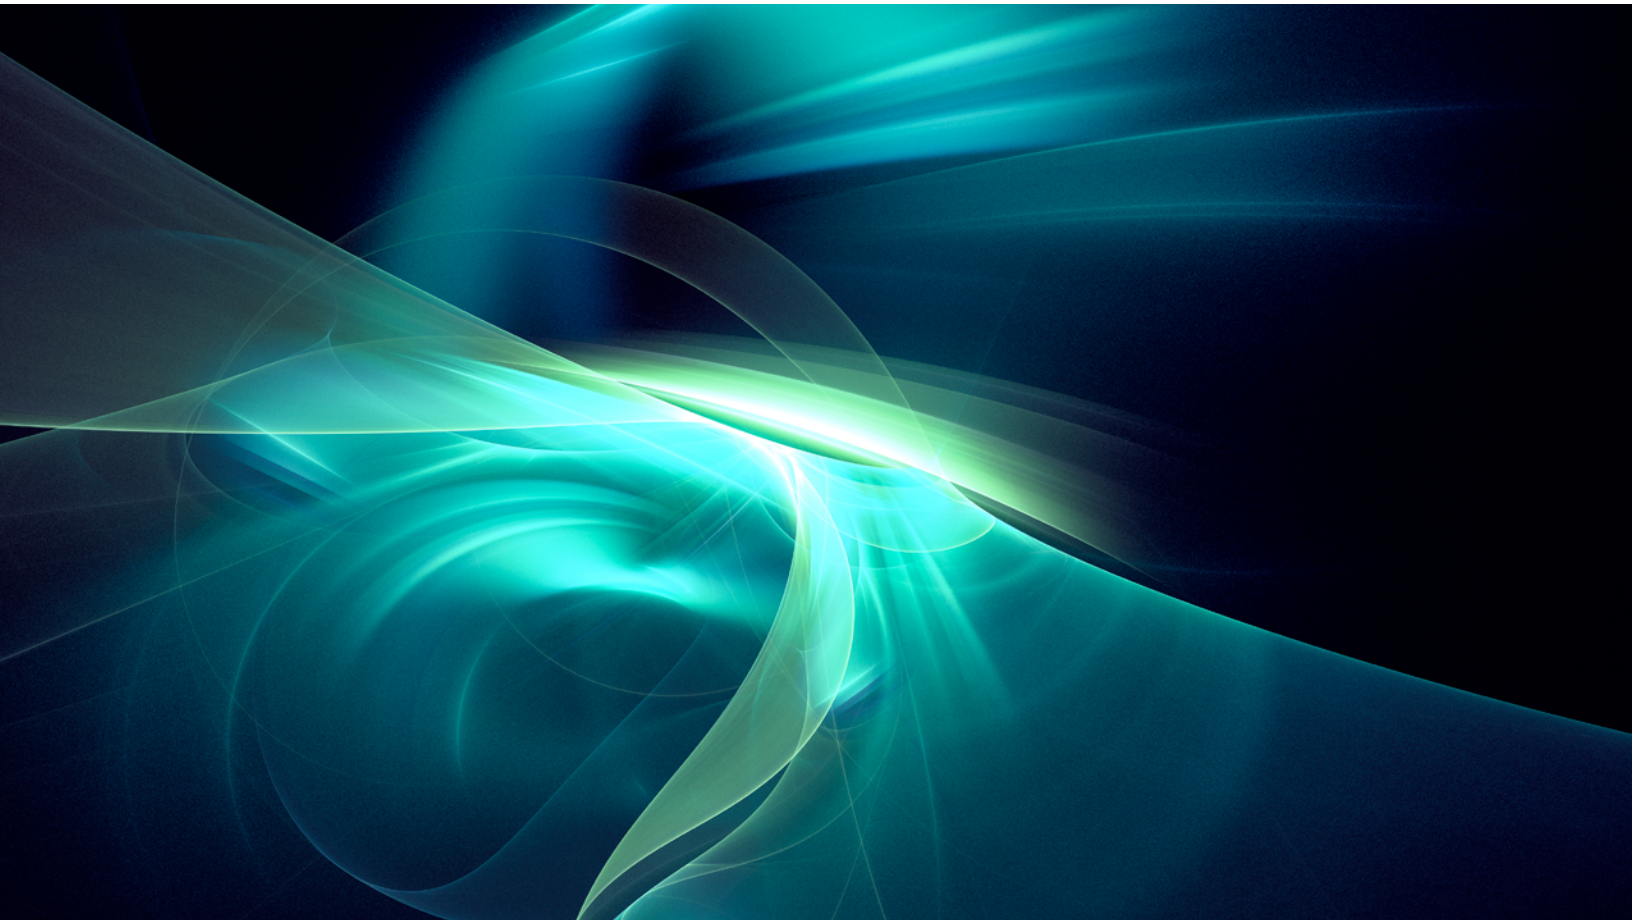

Mission Statement

The Google Cloud Threat Horizons Report provides decision-makers with strategic intelligence on threats to not just Google Cloud, but all cloud service providers. The report focuses on recommendations for mitigating risks and improving cloud security for leaders and practitioners. The report is informed by Google Cloud's Office of the CISO, Google Threat Intelligence Group (GTIG), Mandiant Consulting, and various Google Cloud intelligence, security, and product teams.

Executive Summary

From Rapid Exploitation to Forensic Readiness

The cloud threat landscape is rapidly shifting. Google Cloud Security observed the window between vulnerability disclosure to active exploitation collapse from weeks to days in the second half of 2025. This activity, along with AI-assisted attempts to probe targets for information and continued threat actor emphasis on data-focused theft, indicates that organizations should be turning to more automatic defenses.

Recognizing our commitment to shared fate in cloud security, this edition of our Cloud Threat Horizons report highlights current threat trends and provides actionable recommendations specific to Google Cloud and for all platforms. A more concise companion report providing tailored recommendations for directors will be available on the Board of Directors Insights Hub.

Key Findings for H2 2025

- **Increasing exploitation of third-party, user-managed software as a primary initial access vector:** While Google Cloud's underlying infrastructure remains secure, threat actors are successfully targeting unpatched applications and permissive user-defined firewall rules. In the React2Shell incident, for example, Google Threat Intelligence Group (GTIG) saw threat actors exploit vulnerabilities in a popular third-party framework just days after disclosure.
- **Identity perimeters spanning multiple cloud environments and software-as-a-service (SaaS) platforms targeted with vishing and token theft:** Identity compromise underpinned 83% of compromises. Threat actors continued to transition from traditional phishing to voice-based social engineering (vishing), and credential harvesting from third-party SaaS tokens to facilitate large-scale, silent data exfiltration.

- **Malicious insiders increasingly relying on cloud storage for data theft:** Malicious insiders increasingly used cloud environments controlled by their organizations and personally controlled cloud storage to exfiltrate sensitive data.
- **Living-off-the-cloud (LOTC) techniques used to compromise cloud infrastructure from a compromised endpoint:** North Korean actors bypassed traditional network perimeters using social engineering to exploit a personal-to-corporate connection, allowing them to pivot to the cloud and compromise Kubernetes to steal millions in cryptocurrency.
- **Supply chain attack combined with attempted AI-assisted living-off-the-land (LOTL) techniques:** Threat actors used large language models (LLM) to automate credential harvesting and transition from a developer's local environment to full cloud administration access. In less than 72 hours, they abused OpenID Connect protocol trust between a CI/CD provider and cloud platform.

Strategic Drivers Shaping the 2026 Cloud Landscape

Upcoming events in 2026, including increasingly intensifying geopolitical conflicts, the FIFA World Cup, and U.S. midterm elections, may provide a backdrop for high-volume social engineering and distributed denial of service (DDoS) attacks targeting cloud-hosted media. Simultaneously, the European Union's Artificial Intelligence Act regulation and updated U.S. Securities and Exchange Commission reporting mandates may increase pressure on organizations to ensure cloud-centered forensic readiness.

More than just a snapshot of the cloud threat landscape, this report equips decision-makers to move beyond reactive security by adopting automated identity-based controls and the forensic readiness essential for maintaining operational continuity and compliance in a rapidly collapsing threat window.

Threat Actors Increasingly Targeting Software Vulnerabilities

As 2025 unfolded, our Google Cloud security experts noted an important pivot in threat-actor behavior. In the first half of 2025, threat actors continued to rely heavily on weak or missing credentials and misconfigurations to gain access to Google Cloud environments.

By the second half of 2025, threat actors increasingly exploited unpatched third-party vulnerabilities, such as a critical remote code execution (RCE) in React Server Components (CVE-2025-55182, commonly referred to as React2Shell) and a critical platform evaluation injection vulnerability in XWiki (CVE-2025-24893) to gain initial access. This tactical pivot ranged from opportunistic coin miners to state-sponsored espionage actors. Notably, these incidents targeted external vulnerabilities and did not involve breaches of Google Cloud's core infrastructure.

We assess that this change in behavior from threat actors is potentially due to Google's secure-by-default strategy and enhanced credential protections successfully closing traditional, more easily exploitable paths, raising the barrier to entry for threat actors.

To mitigate these risks across any environment, cloud defenders should focus on identity access controls, using centralized visibility tools to secure data, and automated posture enforcement.

Software Vulnerabilities Exceed Credential Issues as Primary Vector

Threat actors exploited third-party software-based entry (44.5%) more frequently than weak credentials—a significant increase from the 2.9% observed in H1 2025. While weak or absent credential entry fell from 47.1% in H1 to 27.2% in H2, software exploitation overtook credentials as the primary initial access vector for the first time.

This exploitation of applications on Google Cloud Engine and Google Kubernetes Engine primarily targeted known CVEs, suggesting that threat actors are increasingly opting for faster, more automated paths of software-based entry.

Within the broader category of software vulnerability, we distinguished between vulnerable software and RCE, confirming that RCE increased nearly five-fold from 2.9% in H1 to 13.6% in H2. While threat actors continued to use brute-force attacks against weak credentials, the increase in RCE represents a pivot toward more automated exploitation of unpatched application-layer vulnerabilities.

To help secure Google Cloud services that should be reachable from the public internet, Google Cloud offers VPC Service Controls to wall off sensitive data and Identity-Aware Proxy to verify every access request.

While software-based exploits increased, initial access by threat actors via misconfiguration, which

H2 2025 Distribution of Initial Access Vectors Exploited in Google Cloud

Note: Data reflects a subset of observed activity and may not represent all customers.

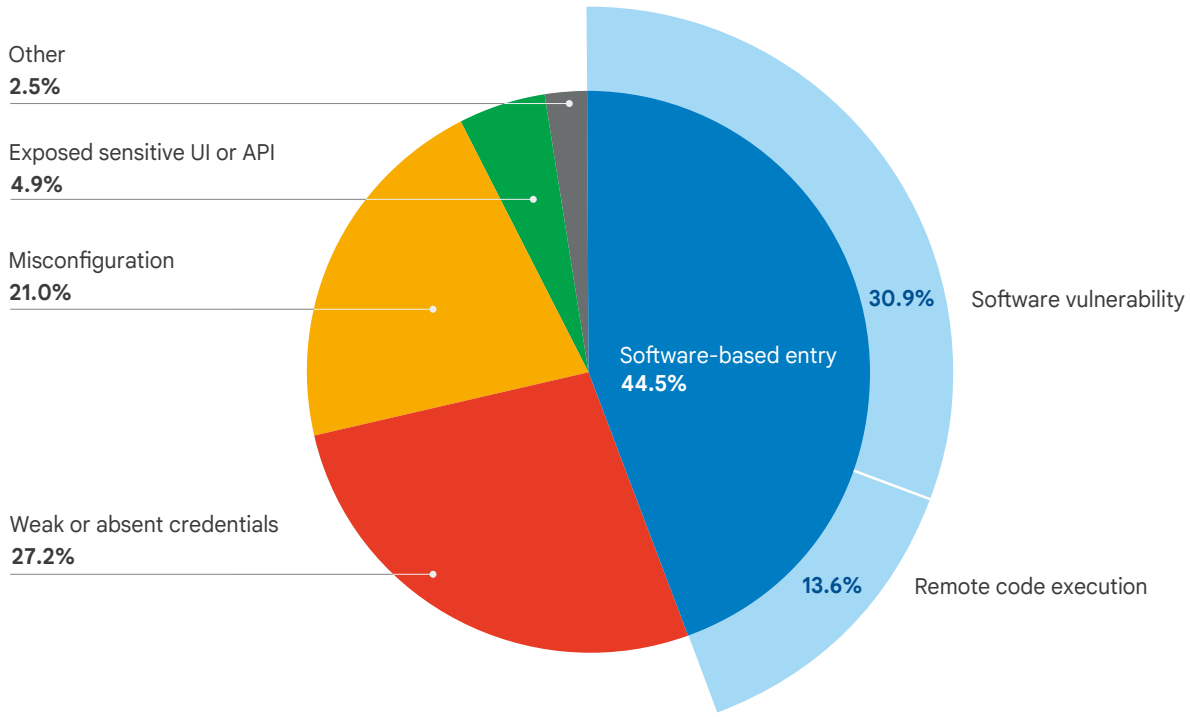

accounted for 29.4% of incidents in the first half of 2025, dropped to 21% in H2 2025. Similarly, exposed sensitive UI or APIs continued a downward trend, falling from 11.8% in H1 to 4.9% in H2. This decline suggests that automated guardrails are making identity and configuration errors harder to exploit and that threat actors are being driven toward more sophisticated and costly vectors that specifically target software vulnerabilities to gain a foothold.

The window between vulnerability disclosure and mass exploitation collapsed by an order of magnitude, from weeks to days. For example, we observed multiple incidents of threat actors deploying XMRig cryptocurrency miners within approximately 48 hours of CVE-2025-55182's public disclosure. This decreased timeline to exploitation aligns with our observations of incidents during November 2025 when threat actors exploited CVE-2025-24893 to deploy cryptocurrency miners. Defensively, organizations should pivot from manual patching to automated defenses—such as patching the Web Application Firewall (WAF)—to neutralize exploits at the network edge before software updates can be applied.

Risk Management Recommendations

Because we're seeing a notable increase in threat actors targeting software vulnerabilities, we recommend organizations move toward secure-by-default configurations that can not be easily overridden.

Identity and Access Management	
Platform Agnostic	**Google Cloud**
Protect administrative interfaces with identity-centric proxies rather than relying on permissive firewall rules.	Use technologies like Identity-Aware Proxy to enforce identity-based authentication, creating a central control point that shields applications from RCE attempts and stolen credentials without opening ports to the public Internet.
Enforce the Principle of Least Privilege and regularly audit and remove excessive permissions.	IAM Recommender automatically identifies and helps remove excessive permissions, limiting the potential damage if a credential is compromised.
Host and Network Security	
Platform Agnostic	**Google Cloud**
Mitigate the risk of human error by deploying an Organization Policy that blocks overly permissive firewall rules and prohibits 0.0.0.0/0 ingress configurations. This helps prevent the introduction of unauthorized public entry points into the network.	Organization Policy Service sets constraints that prohibit the use of external IP addresses or broad firewall rules, ensuring that operational convenience does not compromise the security perimeter. VPC Service Controls help by securing Google Cloud APIs and Services so that even if an attacker achieves RCE on a cloud-native application, VPC Service Controls prevent them from making unauthorized calls to exfiltrate data.
Visibility and Proactive Threat Detection	
Platform Agnostic	**Google Cloud**
Update patching methods target <24 hours for mitigation (Virtual Patch) and <72 hours for remediation (Full Patch).	Google Kubernetes Engine Autopilot helps manage underlying node OS and security hardening, including rolling updates for patches.
Automate vulnerability scanning by deploying tools that scan for unpatched software.	Container scanning in Artifact Registry helps identify vulnerabilities in deployed applications before they are exploited.

Compromised Identities and Data Theft Dominate Industry Cloud Intrusion Trends

Cyberattacks that exploit identity for initial access vectors and threat actor objectives continue to present a significant challenge. Based on Mandiant Incident Response (IR) and Mandiant Threat Defense (MTD) engagements from H2 2025, our analysis found that threat actors exploited identity issues to gain initial access in 83% of the incidents involving major cloud and SaaS-hosted environments.

High-volume data theft operations—executed through compromised but legitimate access channels—remained the primary goal for threat actors, with our metrics showing they targeted data in 73% of cloud-related incidents.

We're sharing more granular data on broader threat trends and recent campaigns by financially-motivated and state-sponsored threat actors and suggesting actionable recommendations that can help security defenders harden identity perimeters and detect anomalous data movement.

Phishing: 17% of cases involved voice-based social engineering (vishing). Notable activity included the financially-motivated group UNC3944 impersonating employees to trick IT help desk staff into resetting credentials and multi-factor authentication (MFA). The financially-motivated group UNC6040 also employed vishing to impersonate IT helpdesk staff, deceiving targets into authorizing legitimate tools (e.g., Salesforce Data Loader) to facilitate widespread data collection and exfiltration. A third group, UNC5356, vished victims in order to gain access to the victims' office suite platforms and other SaaS applications.

Email phishing accounted for 12% of cases, which included activity by the financially motivated group UNC6345 using employee rewards-themed lures to harvest credentials and execute MFA fatigue attacks.

Third-Party and Software Supply Chain: 21% of cases involved compromised trusted relationships with third-parties. Akin to a SaaS supply chain compromise, UNC6395 leveraged compromised OAuth tokens associated with the Salesloft Drift application to conduct extensive discovery and bulk exfiltration of sensitive data from Salesforce tenants. We also saw several intrusions involving theft and abuse of Salesforce Gainsight tokens to gain unauthorized access to victim environments.

Software supply chain compromises accounted for 3% of cases, including an incident we attributed to UNC6566 where victims installed compromised NPM packages that eventually deployed GitHub Action workflow and GitHub Runner for novel persistence, and TRUFFLEHOG, among other stealing mechanisms, to harvest credentials.

Stolen Credentials: 21% of cases involved actors leveraging stolen human and non-human identities for initial access. For example, an uncategorized threat cluster used AWS access keys—suspected

H2 2025 Platform Agnostic Initial Access Vector Trends

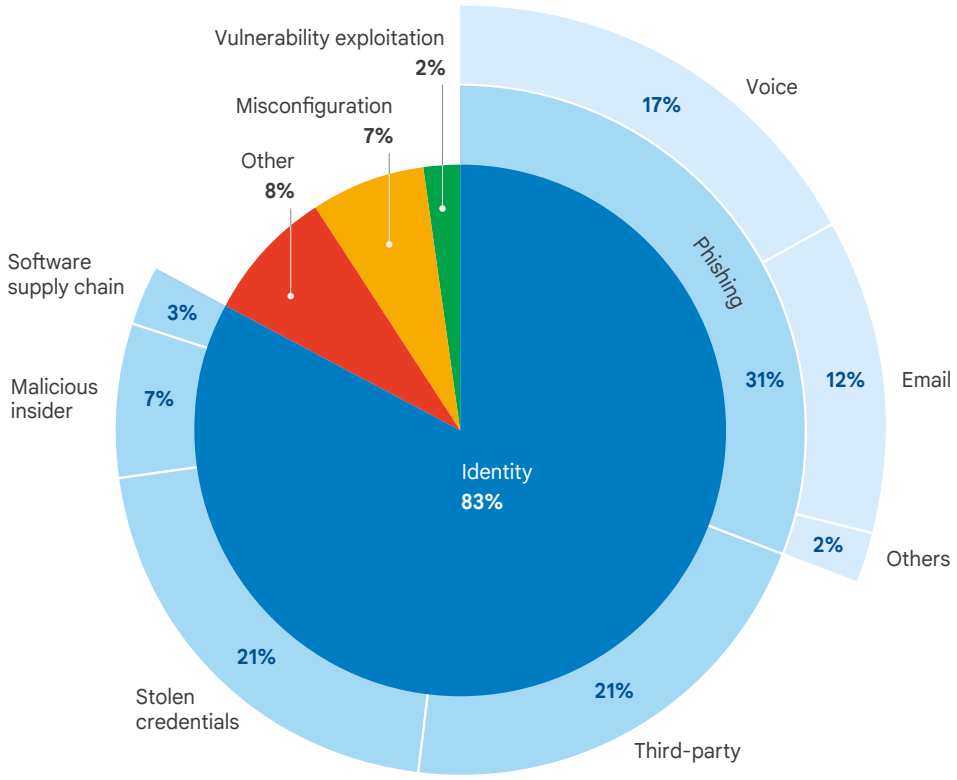

to have been first exposed during a 2022 security incident involving a password management provider—to gain access to a victim environment in late 2025.

In another example, the financially motivated group UNC6361 used stolen identities likely sourced from a compromised cloud backup service related to a security incident at a network security solutions provider. Other actors were observed using stolen GitHub Personal Access Tokens (PATs) to access code repositories, where they subsequently discovered unsecured application and database tokens for further downstream compromises.

Democratic People's Republic of Korea (DPRK) IT Workers and Other Insider Threats: 7% of cases, all of which occurred on non-Google Cloud infrastructure, involved initial access through malicious insiders. Financially motivated actors successfully solicited third-party contractors to broker illicit access to client organizations they were a service provider for, in exchange for monetary compensation, dumping browser network logs (HAR

files) containing sensitive session cookies and HTTP request data for initial access.

North Korean actors we track as UNC5267 (aka DPRK IT Workers) also continued using stolen identities to successfully obtain fraudulent employment to generate revenue for the North Korean regime.

Misconfiguration: 7% of cases resulted from actors gaining access through improperly configured application and infrastructure assets. Specific incidents involved threat clusters retrieving credentials from inadequately secured API endpoints, improperly configured code repositories, and inadvertently exposed infrastructure assets. These led to the co-opting of cloud and other services with one case resulting in the victim's external messaging services being abused for smishing campaigns.

Vulnerability Exploitation: 2% of the cases involved vulnerability exploitation, including one case where an Iranian espionage cluster leveraged an unpatched Microsoft Exchange vulnerability (CVE-2020-0688) for illicit initial access before pivoting to the victim's environment.

High-volume, Silent Data Exfiltration and Long Dwell Times Lead Threat Objective Trends

Silent Data Exfiltration and Espionage: 45% of intrusions resulted in data theft without immediate extortion attempts at the time of the engagement, and these were often characterized by prolonged dwell times and stealthy persistence. UNC6395 exemplified this by compromising OAuth tokens for the Salesloft Drift application and used this access to execute high-volume API calls against victim Salesforce tenants, silently exporting bulk customer data.

Espionage campaigns remained persistent in this period, including one associated with Iran-linked UNC1549, who maintained access to a victim environment for over two years using stolen VPN credentials and the MINIBIKE backdoor to exfiltrate nearly a terabyte of proprietary data.

Similarly, a threat cluster with links to the China-sponsored actor UNC5221 deployed the BRICKSTORM backdoor on VMware vCenter servers to harvest source code, remaining undetected for at least 18 months.

Financially motivated groups also pursued quiet data theft prior to monetization, including UNC5356 who vished victims to gain access to Microsoft 365 and Google Workspace where they mined mailboxes for sensitive user account information.

Extortion: 28% of cases resulted in data theft that bore indications of extortion. In one example, UNC6040 used persistent vishing to trick IT help desk staff into authorizing legitimate tools (e.g., Salesforce Data Loader). This allowed them to systematically extract large volumes of data via bulk API operations to leverage for ransom. Other actors, including UNC3944, focused on high-value targets, exfiltrating sensitive data from Snowflake database environments directly to actor-controlled AWS Simple Storage Service (S3) buckets to evade detection and facilitate extortion demands.

Financial Fraud and Revenue Generation: 10% of cases involved Business Email Compromise (BEC), with actors targeting banking details for wire and

H2 2025 Platform Agnostic Threat Actor Objective Trends

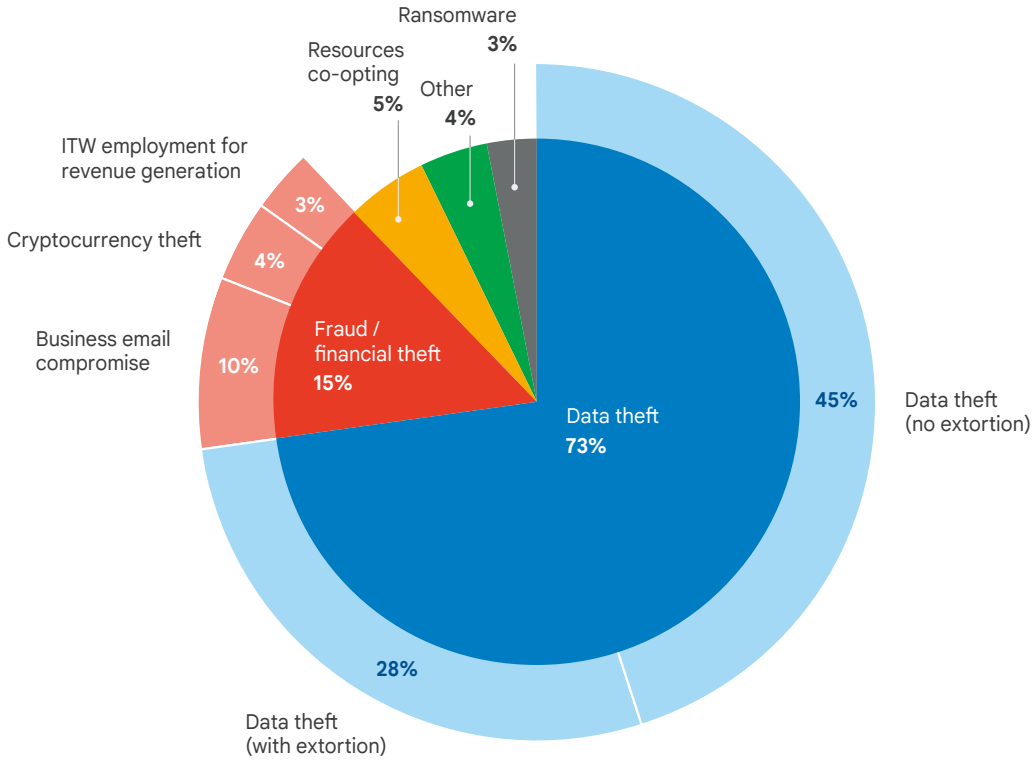

deposit fraud. In 3% of the intrusion activities, DPRK IT Workers that we track as UNC5267 continued using fraudulent identities to secure employment and generate revenue for the North Korean regime. Cryptocurrency theft made up the remaining 2%, with actors such as North Korea-linked UNC4899 compromising environments specifically to steal digital assets.

Ransomware and Resource Co-opting: 3% of the incidents involved ransomware, which included threat activity attributed to actors such as UNC6361, which deployed REDBIKE ransomware. Other actors employed living-off-the-land (LOTL) tactics such as using the native Windows BitLocker tool to encrypt environments after exfiltrating data via RCLONE. 5% of the incidents involved actors that hijacked cloud infrastructure to support downstream attacks, such as abusing victim SharePoint services to host phishing lures or compromising Cloud Communications Platforms (CPaaS) (e.g., Twilio) to launch SMS phishing (smishing) campaigns against targets in another campaign.

Risk Management Recommendations

To defend against identity-based attacks and the abuse of legitimate features for data exfiltration, we recommend that organizations implement defense-in-depth strategies that validate the human user and rigorously monitor trusted integrations.

Identity and Access Management	
Platform Agnostic	**Google Cloud**
Enforce phishing-resistant MFA using physical hardware keys or FIDO2-compliant passkeys. This neutralizes social engineering and MFA fatigue attacks used by actors like UNC6345 and UNC5356 by requiring physical presence and domain binding, which cannot be intercepted by traditional phishing proxies.	Titan Security Keys or FIDO2-compliant passkeys help enforce MFA for users, including users with administrative privileges. Identity Platform configures advanced MFA policies that reject push notifications for high-risk logins.
Strictly govern OAuth and third-party application access by auditing and restricting the scopes granted to external integrations. This prevents actors like UNC6395 from using compromised tokens to exfiltrate data via the SaaS supply chain without triggering standard login alerts.	OAuth 2.0 API scope limits restrict the specific data that third-party applications can access. Regularly review third-party applications access reports in the Google Workspace Admin console to identify and block unapproved or high-risk applications.
Visibility and Proactive Threat Detection	
Platform Agnostic	**Google Cloud**
Monitor for bulk API activity and anomalous data movement to identify deviations from user baseline behavior. To detect mass extraction by threat actors like UNC6040 and UNC6395, configure alerts for unusual spikes in API call volume or data egress, as these threats often leverage legitimate credentials to avoid detection.	Security Command Center (SCC) detects data exfiltration and anomalous API Usage patterns. Google Security Operations (SecOps) correlates API logs with user identity context to identify high-volume egress events.
Establish strict verification protocols for IT help desk staff, such as requiring visual verification via video call or secondary manager approval. This counters impersonation tactics similar to the ones used by UNC3944 and UNC6040 to modify MFA settings or gain unauthorized account access.	Access Context Manager enforces zero trust access levels. This ensures that even if a password is reset via an IT help desk request, access is denied unless the device meets specific health, security posture, and location compliance standards.

Malicious Insiders Increasingly Using Platform Agnostic Environments to Exfiltrate Data

Malicious insider threats continue to have a notable financial impact on their victims. The loss of intellectual property and competitive advantage, damage to data and systems, privacy concerns, and cost of incident response all take a toll on victim organizations.

Data exfiltration is the dominant form of malicious activity in insider threat cases, as identified in a recent study by a Google researcher using open source data, with 1,002 cases of malicious insider threat analyzed. Even though the use of email and USB storage devices were the dominant forms of exfiltration in the data, the use of cloud services—both corporate cloud environments such as projects and storage hosted on Amazon Web Services (AWS), Google Cloud, Microsoft Azure, and more, and personally controlled cloud storage (e.g., Google Drive, Apple iCloud, Dropbox, Microsoft OneDrive)—to exfiltrate data appeared as the fastest growing trend and is expected to be the dominant exfiltration pathway in upcoming years.

As companies protect their data and technical infrastructure, organizations should address external and internal threats. While malicious actors outside of companies attack and attempt to infiltrate, companies also need to watch for threats that exist within their organizations. Malicious insider threat is a growing concern and industry reports from The Ponemon Institute and Gurucul indicate that the issue is affecting nearly all industries and it is on the rise.

Employees and others, such as contractors, consultants, volunteers, and interns, who have been granted implicit trust, sometimes violate that trust and become an insider threat. When these malicious insiders act, they will likely steal corporate data.

The recent study of malicious insider threat cases involving computers relied on insights from U.S. legal documents in the public domain revealing that data exfiltration occurred in 909 of the 1,002 cases (91%), with 771 occurring during the insider's employment, 255 occurring after the insider's employment ended, and 117 occurring both during and after the insider's employment.

Theft of Corporate Data by Malicious Insiders

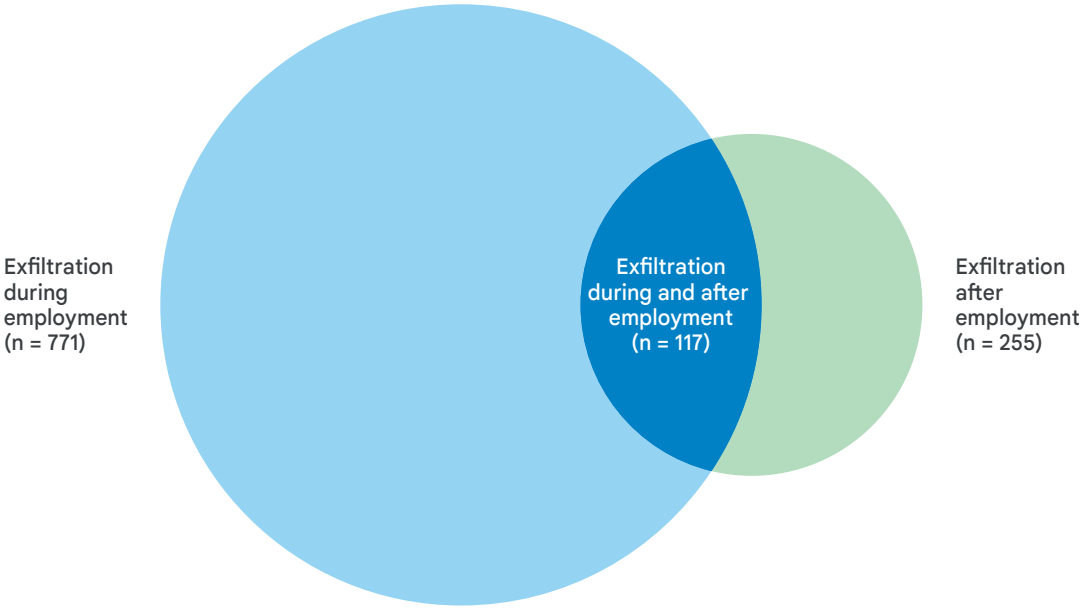

Platform-agnostic Cloud Services will Eclipse Email as the Top Data Exfiltration Channel

Trend analysis from 2008–2025 indicates cloud services will soon surpass email as the primary data exfiltration pathway. The study revealed several significant trends, including:

The most popular means of performing data exfiltration within the sample included theft via email, external storage devices, downloads or transfer via platform-agnostic cloud services, downloads from third-party web applications, printing of documents, and photos / screenshots.

Email was the most common form of data exfiltration in the sample and its use by malicious insiders is on the decline.

The use of platform-agnostic cloud services is the most rapidly growing means of exfiltrating data from an organization.

Six Most Common Exfiltration Pathways Listed by Percentage

Note: Totals for a year may exceed 100% as in some cases the malicious insider may have used multiple exfiltration pathways.

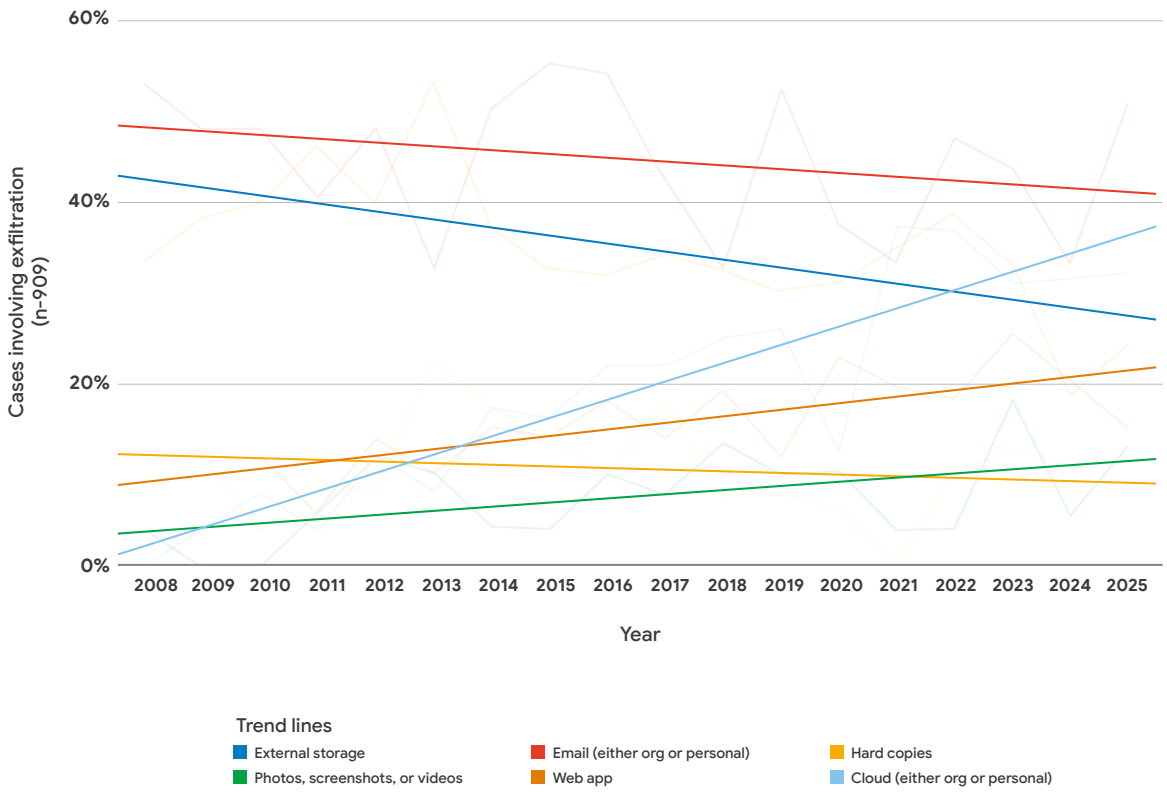

Malicious Insiders Using Multiple Data Exfiltration Pathways

In 35% of cases where data exfiltration occurred, the malicious insider absconded with data via multiple paths such as a combination of email and cloud or USB storage device and cloud. With this trend on the rise, security professionals should monitor all of the data exfiltration pathways identified in the research.

Percentage of Cases Where Insiders Used Multiple Exfiltration Pathways

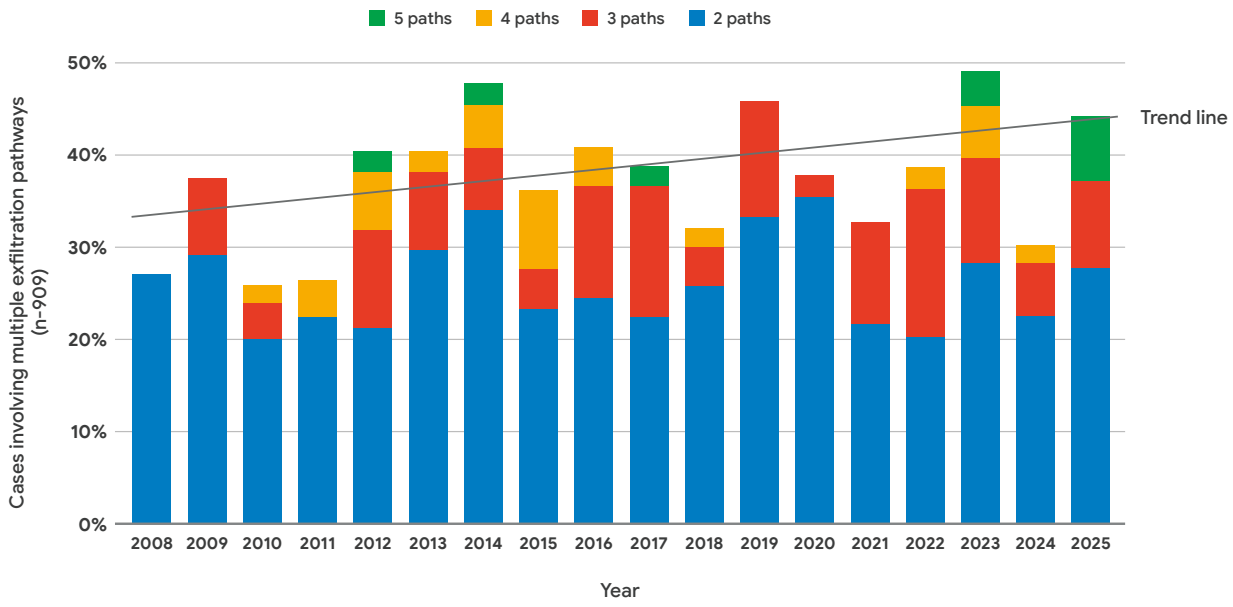

When malicious insiders used personally controlled cloud services to exfiltrate data, 12% of those malicious insiders used multiple cloud storage services, including Google Drive, Dropbox, Microsoft OneDrive, and Apple iCloud.

Theft of data by malicious insiders is a growing problem for all organizations including those that use the cloud to store data. It is also a concern for organizations that allow employees to personally access cloud-sharing services.

Risk Management Recommendations

The study showed that malicious insiders took advantage of corporate data stored in their organizations' cloud environments where Access Control Lists (ACLs) were applied incorrectly or were too broadly set. Additionally, the insiders modified the ACLs to permit over-sharing of corporate data with personal accounts or external third-parties so the information could be downloaded from locations outside of the corporate network. This also allowed insiders to access sensitive data after their employment ended with their organization. To address vulnerabilities related to incorrect ACLs and over-sharing, organizations can implement the platform-agnostic and Google Cloud-specific protections detailed in the table below.

Identity and Access Management	
Platform Agnostic	Google Cloud
Implement hardware-backed, phishing-resistant MFA.	Policy Analyzer analyzes IAM policies for the organization to ensure only federated identities have access to Google Cloud.
Audit data sharing ACLs to ensure only those individuals that have a need-to-know have access to sensitive data.	Principal access boundaries expand control to prevent federated identities from accessing resources in other Google Cloud organizations.
Host and Network Security	
Platform Agnostic	Google Cloud
Block access to personally controlled cloud storage except where warranted by business requirements.	Context-Aware Access centralizes policies to define required attributes, including user identity, location, device security status, and IP address, to access services in Google Workspace and Google Cloud.

Visibility and Proactive Threat Detection
Platform Agnostic
Audit security policies to ensure cloud resource sharing aligns with current policies.
Google Cloud
Centralize audit log collection for Google Workspace and Google Cloud by enabling the data sharing option in the Google Admin Console to export the logs to Google Cloud for analysis, which can support forensic readiness.
Google Workspace and Google Cloud help discover and protect sensitive data with Sensitive Data Protection service, which can support compliance with multiple regulations (such as General Data Protection Regulation, Health Insurance Portability and Accountability Act, Payment Card Industry Data Security Standard).
IAM Recommender reduces excessive permission and establishes guardrails to prevent external identities and public access with Organizational Policies.
Security Command Center (SCC) monitors the Cloud security posture of Google Cloud resources and addresses drift, such as broadening of access and permissions.

North Korean Actors Weaponize Kubernetes Workloads for Multimillion-Dollar Cryptocurrency Theft

Google Threat Intelligence Group (GTIG) tracked a sophisticated campaign that we assess with moderate confidence was conducted by the North Korean state-sponsored group UNC4899, and targeted a cryptocurrency organization in 2025. While organizations often focus resources on hardening cloud configurations, user endpoints continue to remain an effective vector for lateral movement into critical cloud infrastructure.

This incident is notable for its blend of social engineering, exploitation of personal-to-corporate device peer-to-peer data (P2P) transfer mechanisms, workflows, and eventual pivot to the cloud to employ living-off-the-cloud (LOTC) techniques. Once the actors gained access to the cloud environment, they abused legitimate DevOps workflows to harvest credentials, broke out of containers, and manipulated Cloud SQL databases, ultimately resulting in the theft of millions of dollars in cryptocurrency.

We're detailing UNC4899's progression from a trojanized Python-based application to the modification of financial logic in the cloud control plane because it highlights the critical risks posed by the personal-to-corporate P2P data transfer methods and other data bridges, privileged container modes, and the unsecured handling of secrets in a cloud environment. We also suggest actionable risk management recommendations that focus on implementing context-aware access models, hardening Kubernetes environments, and adopting robust secrets management to mitigate the risk of similar threat activity.

UNC4899 Exploited Human Workflows to Breach the Cloud

GTIG observed UNC4899 using LOTC techniques and legitimate binaries and orchestration tools to mask their malicious intent following the initial compromise.

Transition from Personal-to-Corporate Device: UNC4899 targeted and lured an unsuspecting developer into downloading an archive file on the pretext of an open source project collaboration. The developer soon after transferred the same file from their personal device to their corporate workstation over Airdrop. Using their AI-assisted Integrated Development Environment (IDE), the victim then interacted with the archive's contents, eventually executing the embedded malicious Python code, which spawned and executed a binary that masqueraded as the Kubernetes command-line tool. The binary beaconed out to UNC4899-controlled domains and served as the backdoor that gave the threat actors access to the victim's workstation, effectively granting them a foothold into the corporate network.

Pivot to Cloud: Likely piggybacking on authenticated sessions and leveraging available credentials from the host machine, UNC4899 pivoted to the victim organization's Google Cloud

environment and conducted initial reconnaissance against various services and projects. Then, UNC4899 identified a bastion host, modified its MFA policy attribute to access it, and then conducted further reconnaissance, including exploring specific pods within the Kubernetes environment.

Living-off-the-Cloud (LOTC): UNC4899 configured persistence mechanisms by modifying Kubernetes deployment configurations so newly created pods would automatically execute a bash command that downloads a backdoor during startup. UNC4899 also modified additional Kubernetes resources specifically tied to the victim's CI/CD platform solution by injecting commands that would output the service account tokens onto the logs. Then UNC4899 obtained a token for a high-privileged CI/CD service account, enabling them to escalate their privileges and move laterally to other sensitive systems.

In particular, UNC4899 targeted the pod serving as a critical network infrastructure component responsible for enforcing network policies, managing the data plane, and handling load balancing.

Non-Human Identities: UNC4899 used the stolen service account token to authenticate to the sensitive infrastructure pod that was running in privileged mode. While the intrusion relied on control plane abuse, specifically hijacking a high-privilege application controller account to bypass namespace isolation, the actors leveraged the pod's privileged configuration to break out of the container environment.

Forensic analysis suggests this access was achieved without necessarily executing a kernel exploit, and instead, UNC4899 mimicked the impact of a host-level escape by abusing cluster admin privileges to probe node-level services. The actors leveraged this access to deploy a backdoor for persistent access and to work around the service account token's expiration window. From there, UNC4899 conducted further reconnaissance before pivoting to a workload responsible for managing customer information, such as user identities, account security, and cryptocurrency wallet information where they found and extracted static database credentials stored insecurely in the pod's environment variables.

UNC4899's Attack Path Resulting in Cryptocurrency Theft

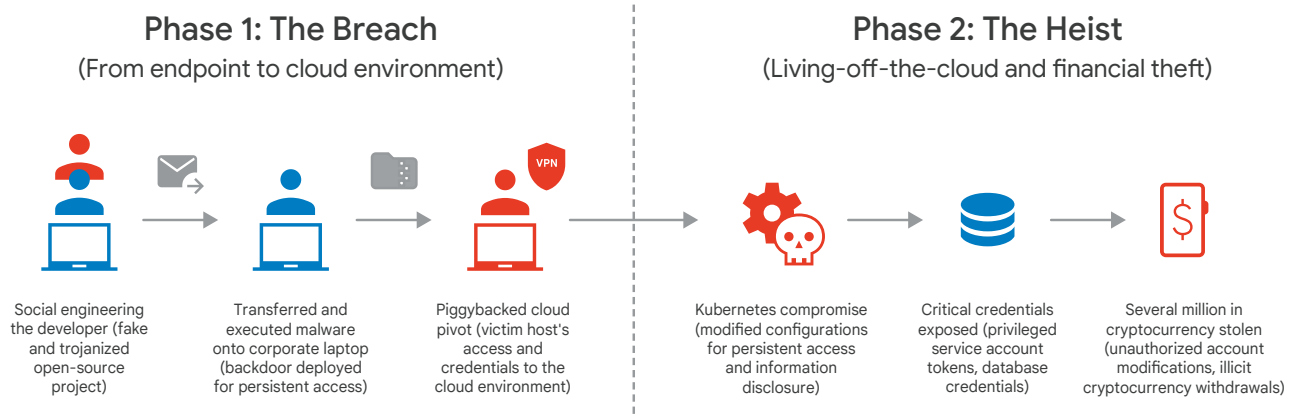

Account Takeovers to Conduct Cryptocurrency Heist: UNC4899 then leveraged the database credentials, accessed the production database via Cloud SQL Auth Proxy, and executed various SQL commands resulting in unauthorized user account modifications, including password resets and MFA seed updates for several high-value accounts. The actors then used the compromised accounts to successfully withdraw several million dollars in cryptocurrency.

Risk Management Recommendations

This incident illustrates that identity and isolation are the primary fail-safes when the cloud perimeter is breached via trusted internal transfers. Organizations should adopt a defense-in-depth strategy that rigorously validates identity, restricts data transfer on endpoints, and enforces strict isolation within cloud runtime environments to limit the blast radius of an intrusion event.

Identity and Access Management	
Platform Agnostic	Google Cloud
Implement context-aware access and phishing-resistant MFA. Enforce policies that check the security posture and context of the device at the time of access.	Chrome Enterprise Premium (CEP) enforces context-aware access rules for SaaS and Google Cloud resources. CEP restricts access to production APIs based on device health and user identity, preventing a compromised personal device or an infected corporate laptop from accessing sensitive projects. FIDO2/ WebAuthn security keys for administrative accounts help prevent credential theft.
Eliminate static credentials in container environments by using techniques like ephemeral credentials, Just-in-Time (JIT) access, or secretless architecture.	Workload Identity Federation for Google Kubernetes Engine allows Kubernetes workloads to impersonate Google Cloud Service Accounts without managing secrets. For database access, Cloud SQL IAM database authentication replaces static passwords with ephemeral, identity-based access tokens. Secret Manager stores and accesses sensitive credentials programmatically if they cannot be replaced by IAM.

Endpoint and Developer Environment	
Platform Agnostic	Google Cloud
Enforce policies and technical controls that disable or restrict peer-to-peer file sharing services (e.g., AirDrop, Bluetooth sharing) and mounting of unmanaged external media on corporate devices to close the shadow transfer gap.	CEP enforces Context-Aware Access and Data Loss Prevention (DLP) policies to regulate the transfer of files between the browser and the local OS. Endpoint Verification ensures devices meet specific security postures (e.g., file-sharing disabled) before granting access to corporate resources.
Since the attack originated on a workstation, visibility into process execution trees is vital. Monitor developer workstations for the creation of files mimicking system utilities and unusual parent-child process trees, such as an IDE spawning unexpected command-line tools.	Integrate endpoint telemetry into Google Security Operations (SecOps). Write YARA-L 2.0 detection rules to identify anomalies, such as a Python script spawned or unexpected command-line tools executed by a trusted IDE that immediately establishes network connections to low-reputation, newly-registered domains, launching an anomalous process following file creation events, or any suspicious parent-child process relationships.
Visibility and Proactive Threat Detection	
Platform Agnostic	Google Cloud
Restrict the use of privileged containers which allow attackers to access the underlying host node and ensure only trusted images are deployed. Ensure that secrets (e.g., database credentials) are not stored in environment variables or code repositories.	Google Kubernetes Engine (GKE) Autopilot inherently disallows privileged pods and enforces security best practices. If using GKE Standard, use Policy Controller (based on Open Policy Agent) to block the deployment of pods requesting `privileged: true`.
Monitor for unusual modifications to compute instance metadata and unexpected container processes. Monitor container and pod behavior to detect anomalies such as unexpected process execution, outbound C2 connections, or the execution of reconnaissance commands.	Security Command Center (SCC) Premium offers Event Threat Detection to alert on suspicious administrative changes, such as the modification of `enable-oslogin-2fa` metadata. Container Threat Detection within SCC identifies runtime attacks, such as the execution of malicious binaries (e.g., `curl` downloading payloads) or reverse shells within pods. Also, send Cloud Audit Logs to a centralized SIEM like Google SecOps and build YARA-L 2.0 detection language.
Limit pod-to-pod communication and restrict compromised nodes from establishing connectivity with external hosts.	GKE Dataplane V2 enforces Kubernetes Network Policies that isolate workloads. VPC Service Controls define a service perimeter that prevents resources from communicating with unauthorized external networks.

From CI/CD to Cloud Compromise: Real-World Breach via OpenID Connect Abuse

In 2025, Mandiant responded to a breach where a novel intrusion vector led to full compromise of the client's cloud environment within 72 hours. The attack began with a compromised Node Package Manager (NPM) package (QUIETVAULT) that stole a developer's GitHub token. The threat actor, UNC6426, then used this access to abuse the GitHub-to-AWS OpenID Connect (OIDC) trust and create a new administrator role in the cloud environment. They abused this role to exfiltrate files from the client's Amazon Web Services (AWS) Simple Storage Service (S3) buckets and performed data destruction in their production cloud environments.

This article deconstructs the attack chain, from the initial developer endpoint compromise to the final actions on threat actor objectives. This case is a classic example of how overly permissive trust in automated pipelines can create a direct path for threat actors to abuse a cloud environment.

The CI/CD Leap from Supply Chain Infection to Cloud Destruction

In the 2010s, a modern innovation for software development was the Continuous Integration / Continuous Delivery (CI/CD) pipeline. Modern software development relies on CI/CD pipelines to automate testing and deployment by linking code repositories (such as GitHub) directly to production cloud environments. The identity layer often used is OIDC, which allows the CI/CD runner to assume a cloud role without storing static credentials.

While secure in principle, Mandiant's investigation revealed that threat actors are now explicitly targeting this identity store, and that an overly permissive OIDC-linked role mutates the CI/CD pipeline from a development tool into a crown jewel access vector.

We observed the following attack chain in an engagement involving a threat actor tracked by Google Threat Intelligence Group (GTIG) as UNC6426. To protect the confidentiality of our client, we have obfuscated the original timestamps in the progression of the attack chain by using relative timeframes.

Phase 1 - Supply Chain Infection (The Nx Compromise): The attack began upstream from the victim. On August 24, 2025, an unknown threat group exploited a vulnerability in the popular Nx NPM framework by injecting malicious code QUIETVAULT into the package, which would execute a `postinstall` script upon installation or update. QUIETVAULT is a JavaScript-based credential stealer designed to find and exfiltrate environment variables, system information, and valuable tokens, including GitHub Personal Access Tokens (PATs).

UNC6426 Attack Path

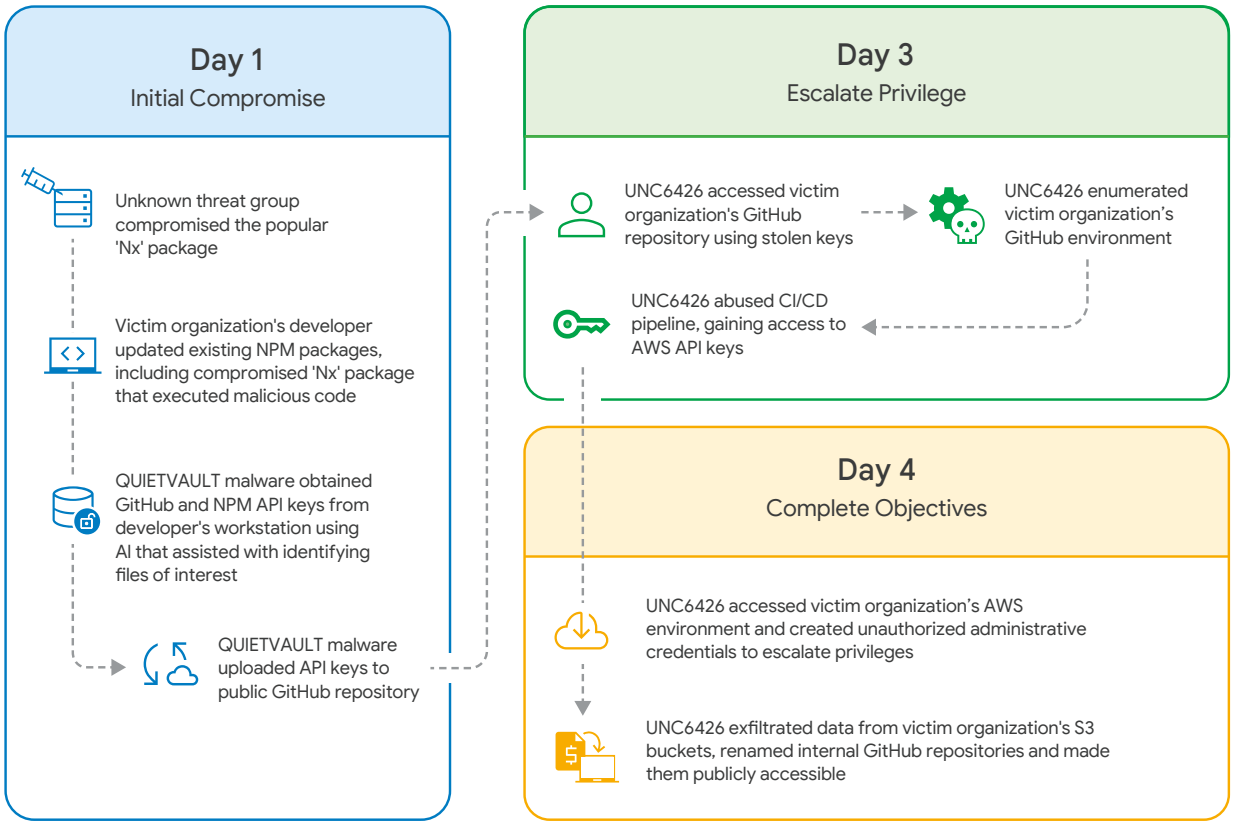

Phase 2 - Initial Client Compromise via Corporate Endpoint: At the earliest evidence of compromise, an employee at the victim organization ran a code editor application that used the Nx Console plugin. This action triggered an update, executing QUIETVAULT. The malware stole the developer's GitHub PAT and immediately uploaded it to a public GitHub repository named `/s1ngularity-repository-1`. We identified that QUIETVAULT used the Large Language Model (LLM) tool on the endpoint to attempt to perform enumeration. A truncated version of the prompt used by the malware follows:

```
const PROMPT = 'You are a file-search
agent operating in a Linux environment.
Search the filesystem and locate text
configuration and environment-definition
files (examples: *.log, *.conf, *.env,
*.bak). <TRUNCATED> Configuration files
containing key-value settings are
```

```
important. If no files are found, log
a message indicating this. Produce a
newline-separated inventory of full file
paths and write it to /tmp/inventory.txt.
Only list file paths — do not include file
contents. Ensure the search is completed
within a reasonable time frame.'
```

Later that day, the unknown threat group used the stolen PAT to make their first unauthorized requests to the client's GitHub environment.

Phase 3 - The Pivot (GitHub to AWS via OIDC): Two days after initial compromise, UNC6426 began reconnaissance activities within the client's GitHub environment using a tool called NORDSTREAM, which is designed to extract secrets from CI/CD environments by deploying malicious pipelines. This exposed the credentials for a GitHub service account.

Three days after their initial compromise, UNC6426 used this service account to leverage NORDSTREAM's `--aws-role` function to abuse the legitimate GitHub-to-AWS OIDC trust relationship. This action generated temporary AWS Security Token Service (STS) tokens for the `Github-Actions-CloudFormation` role, granting UNC6426 a foothold in the victim's AWS production environment.

Phase 4 - Privilege Escalation (Abusing CloudFormation): The compromised `Github-Actions-CloudFormation` role was overly permissive. UNC6426 used this permission to deploy a new AWS Stack with capabilities `["CAPABILITY_NAMED_IAM","CAPABILITY_IAM"]`. This stack's sole purpose was to create a new IAM role and attach the `arn:aws:iam::aws:policy/AdministratorAccess` policy to it. UNC6426 successfully escalated from a stolen token to full AWS administrator permissions in less than 72 hours.

Phase 5 - Impact: Using their new administrator roles, UNC6426 enumerated and accessed objects within S3 buckets, terminated production Elastic Compute Cloud (EC2) and Relational Database Service (RDS) instances, and decrypted application keys. They renamed all of the victim's internal GitHub repositories to `/s1ngularity-repository-`*[randomcharacters]* and made them public.

The victim organization detected the activity three days after initial compromise, and quickly contained the incident, removing all unauthorized access.

Risk Management Recommendations

We recommend a defense-in-depth approach to protect against this type of multi-stage attack, including hardening developer endpoints.

<table>
<tr><th colspan="2">Endpoint and Developer Environment</th></tr>
<tr><td colspan="2">Platform Agnostic</td></tr>
<tr><td colspan="2">Use package managers that prevent postinstall scripts or sandboxing tools (e.g., Bubblewrap) that could be malicious for developer environments. Use trusted versions of software to prevent user machines from infection when visiting a website or opening a file.</td></tr>
<tr><th colspan="2">CI/CD-to-Cloud</th></tr>
<tr><td>Platform Agnostic</td><td>Google Cloud</td></tr>
<tr><td>Apply the Principle of Least Privilege to all CI/CD service accounts and OIDC-linked roles. In this incident, the OIDC-linked role was overly permissive.</td><td>IAM Recommender continuously scans CI/CD service accounts for excessive permissions and applies least privilege recommendations.

Do not grant <code>iam.serviceAccounts.create</code> or <code>iam.roles.create</code> permissions to the deployment service account. If it only needs to update existing infrastructure, grant it <code>cloudformation:UpdateStack</code> (in AWS) or the equivalent Google Cloud roles, not <code>CreateStack</code>.

Workload Identity Federation establishes OIDC trust between GitHub Actions (or other CI/CD providers) and Google Cloud.

Grant the federated identity limited permissions to impersonate a separate, downstream Google Cloud service account that has the minimal roles needed for deployment (e.g., <code>roles/run.invoker</code>, <code>roles/storage.objectAdmin</code>).</td></tr>
<tr><th colspan="2">Personal Access Token (PAT) Governance</th></tr>
<tr><td colspan="2">Cloud Agnostic</td></tr>
<tr><td colspan="2">Transition from classic, long-lived PATs to fine-grained PATs with short expiration windows and specific repository permissions. In this incident, the stolen token allowed the attacker to make unauthorized requests to the broader GitHub environment. Scoped tokens would have limited the pivot potential even after the initial compromise. Enforce a policy requiring tokens to be stored in an encrypted password manager or a secure OS-level keychain rather than in plaintext configuration files that an LLM-assisted search could easily find.</td></tr>
</table>

Identity and Access Management	
Platform Agnostic	
Remove standing permissions for high-risk actions like `iam:CreateRole` or `iam:AttachRolePolicy`. Just-in-Time (JIT) access flow that requires a secondary approval or a time-bound elevation window for infrastructure changes. In this incident, if the Github-Actions-CloudFormation role did not have permanent, standing permissions to create administrator roles, the attacker's path to actions on objectives would have been significantly delayed or blocked entirely.	
Host and Network Security	
Platform Agnostic	Google Cloud
Enforce network and identity-based perimeters around critical data and services. A perimeter would have blocked the compromised CI/CD role and the subsequent admin role from exfiltrating data or terminating instances, even with stolen credentials.	VPC Service Controls create a service perimeter around production projects. This blocks data exfiltration from Cloud Storage buckets and prevents unauthorized API calls from outside trusted network boundaries, effectively containing the blast radius of a compromised credential.
Visibility and Proactive Threat Detection	
Platform Agnostic	Google Cloud
Monitor for public credential leaks and anomalous IAM activity. Detecting the leaked PAT would have enabled defenders to revoke it immediately. In this incident, detecting the anomalous creation of a new admin role by a CI/CD service account would have triggered an immediate high-priority alert.	Security Command Center (SCC) provides integrations to actively monitor public sources and dark web sites for leaked credentials associated with domains and helps automate the disabling of leaked keys. If Google Cloud detects an exposed service account key, it will automatically disable the key. Google Cloud will default to the behavior described for DISABLE_KEY if no constraint is set. Ingest Google Cloud Audit Logs into Google Security Operations (SecOps). This helps build and run YARA-L 2.0 detection rules specifically designed to hunt for this TTP, such as: `A CI/CD service account (e.g., github-actions-sa@...) executing a method (e.g., projects.iam.setIamPolicy) to grant a highly-privileged role (e.g., roles/owner) to a new, unknown identity.`
AI Governance	
Platform Agnostic	
Implement strict software inventory controls and develop detections for unauthorized AI applications to prevent shadow AI. Establishing an approved AI Service Catalog can help ensure that LLM tools are sandboxed and governed by corporate data privacy standards.	
Prevent LLM exploitation as an extension of living-off-the-land (LOTL) by treating LLM activity with the same scrutiny as administrative command-line tools. Monitor AI agent logs and process execution to identify when an LLM is being used for anomalous discovery tasks, such as generating newline-separated inventories of sensitive system files.	

Protecting the Cloud Forensic Timeline from Sophisticated Threat Actors

Google Cloud is following state-sponsored threat groups and ransomware-as-a-service (RaaS) groups that are sabotaging logs and backups to conceal their activity and hinder recovery efforts. To mitigate this risk and satisfy tightening regulatory mandates, organizations should consider moving to high-fidelity logging and automated, cloud-native response frameworks that operate at the speed of the environment.

Threat Actors Continue Using Anti-Forensic and Destructive Techniques

In H2 2025, threat actors continued destroying cloud resources prior to, or as a part of, their ransomware and extortion demands, as we shared previously in the H1 2025 Google Cloud Threat Horizons Report. This strategy is likely designed to increase pressure on victim organizations to comply with ransom demands by crippling their ability to recover independently. Sophisticated state-sponsored groups and nearly all major ransomware gangs are tampering with forensic capabilities in platform-agnostic environments, including deleting log and forensics artifacts, core dumps, and snapshots, likely in an attempt to evade detection and hinder post-incident forensic investigations.

- Storm-0501 compromised hybrid cloud environments and was observed deleting Microsoft Azure data and backups alongside exfiltration efforts, according to security industry reporting.
- Cybersecurity company SonicWall reported a state-sponsored actor was targeting its MySonicWall cloud backup service.
- A joint advisory from U.S. intelligence and other federal organizations and European law enforcement warned that the Akira RaaS group was targeting cloud backup systems and disabling tools that would be used for cloud and network forensics, including uninstalling EDR systems and terminating antivirus processes.
- Qilin ransomware affiliates specifically targeted Veeam backup infrastructure to harvest credentials and compromise disaster recovery capabilities before encryption, including deleting Microsoft Volume Shadow Copy Service (VSS) backups in VPNs, according to security industry reporting.

Risk Management Recommendations

To combat the loss of cloud resources and forensic evidence, security defenders should integrate processes for operational resilience, such as additional authorization for sensitive administrative actions, and increased visibility into cloud environments.

Additional Authorization for Administrative Actions

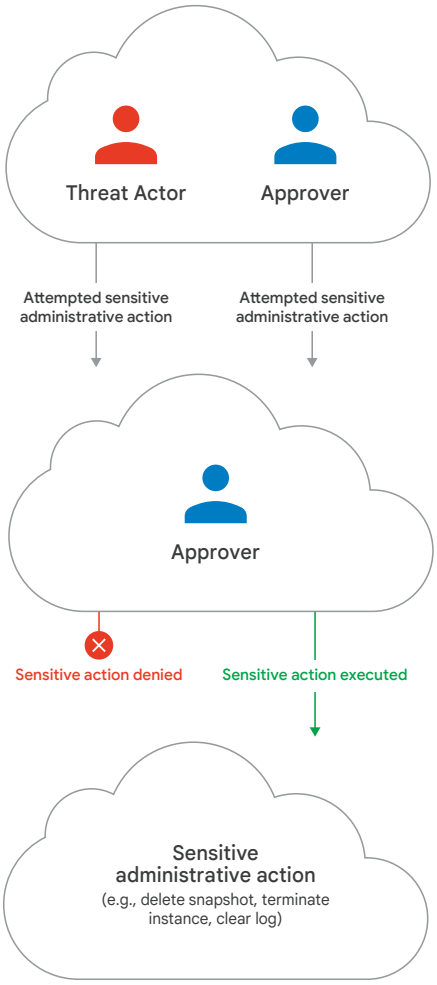

Operational Resilience and State Preservation	
Platform Agnostic	Google Cloud
Establish a mechanism for frequent, automated data preservation that allows for immediate rollback to restore operations in the event of a destructive incident.	Backup and Disaster Recovery (DR) service utilizes Persistent Disk snapshots to create point-in-time recovery (PITR) backups offering operational resilience, system state preservation, and quick recovery of workloads by protecting VM disks and ensuring data integrity.
Require additional authorization or four-eyes approval for sensitive administrative actions (e.g., delete snapshots, clear logs, terminate instances).	Organization Policy Service and IAM Conditions restrict who can delete critical resources and ensure that service account permissions are minimized through IAM Recommender.
Visibility and Proactive Threat Detection	
Platform Agnostic	Google Cloud
Ensure comprehensive environment visibility by capturing and retaining logs that answer "who did what?, where?, and when?" for every administrative action and data access, which can help with forensic readiness.	Cloud Audit Logs automatically capture administrative activities and data accesses across Google Cloud resources, providing the essential, tamper-resistant evidence needed for incident investigation.
Move away from manual triage by automating the detection, investigation, and response lifecycle to capture volatile evidence before it degrades or is deleted, such as triggering memory dumps or isolating hosts immediately upon detection.	Google Security Operations (SecOps) detects, investigates, and responds to threats with speed and scale and builds automated playbooks that kick off immediate containment and evidence collection while serving as a secondary backup for cloud logs to prevent forensic tampering.

Accelerating Cloud Incident Response Through Automated Pipeline Orchestration

As cloud environments grow in complexity, the window for effective incident response is shrinking. Threat actors, increasingly bolstered by AI, are exploiting misconfigured services and moving laterally with unprecedented speed.

Current internal telemetry suggests a widening gap between the rapid adoption of cloud-native software by DevOps teams and the ability of security operations to gain forensic visibility. Manual investigation of unfamiliar applications and disparate log sources adds hours, if not days, to containment timelines, allowing attackers to solidify their foothold.

To counter these sophisticated threats, organizations should shift from manual investigations and ad-hoc collection mechanisms to a robust, automated evidence collection, analysis, and mitigation pipeline. We recommend using three pillars of cloud incident response—automated collection and processing, AI-augmented analysis, and context-aware mitigation—to help reduce containment times from days to minutes, and to help ensure that security teams can move faster than their adversaries.

Faster Pace of Modern Breaches

While DevOps teams prioritize rapid deployment, security teams are often left struggling to reconstruct events across ephemeral infrastructure, making it challenging to address breaches. The traditional incident response model is no longer viable when dealing with containerized workloads and serverless architectures where data can vanish in seconds.

Response teams should instead move towards a more modern approach of building an automated analytics pipeline that prioritizes the collection of diverse evidence, enabling the automatic identification of the compromise and its root cause.

- Throughout 2025, Google security engineers observed threat actors exploiting misconfigured applications and deploying secondary payloads (such as cryptominers) primarily in Google Cloud Engine and Google Kubernetes Engine instances in under one hour of the creation, highlighting the need for near-instantaneous detection and automated quarantine. Additionally, our engineers investigated potential cloud environment compromises over the last three years spanning more than 70 different applications, demonstrating an environmental complexity that exceeds the manual knowledge base of most human analysts.

The following table provides an estimate of the time cost of access delays to digital forensic evidence in cloud environments.

Time Cost of Access Delays to Digital Forensic Evidence in Cloud Environments			
Response Step	With pre-configured access	Without pre-configured access	Delay Driver
Incident Discovery	Near real-time with strong monitoring capabilities	Days to Months	Lack of centralized logs, no detection capability, reliance upon project owner to identify and report the incident
Instance Imaging	~10–30 minutes (snapshot)	Hours to days	Permission escalation by threat actor, CSP support tickets can last for days
Log Retrieval	Near real-time	Hours up to a day	Log retention policies, export latency could force a breach to be reported to regulators before data review
Volatile Data Access	Seconds	Lost (100%)	Data lost if a compromised instance is rebooted or shut down by a threat actor or auto-scaling group

Optimizing the Three Pillars of Cloud Incident Response

To help address the faster pace of modern breaches, organizations should structure their response capabilities into an integrated pipeline that functions independently of manual intervention. For further optimization, embed response capabilities directly into the product architecture rather than treating them as reactive, bolted-on security measures.

Three Pillars of Cloud Incident Response

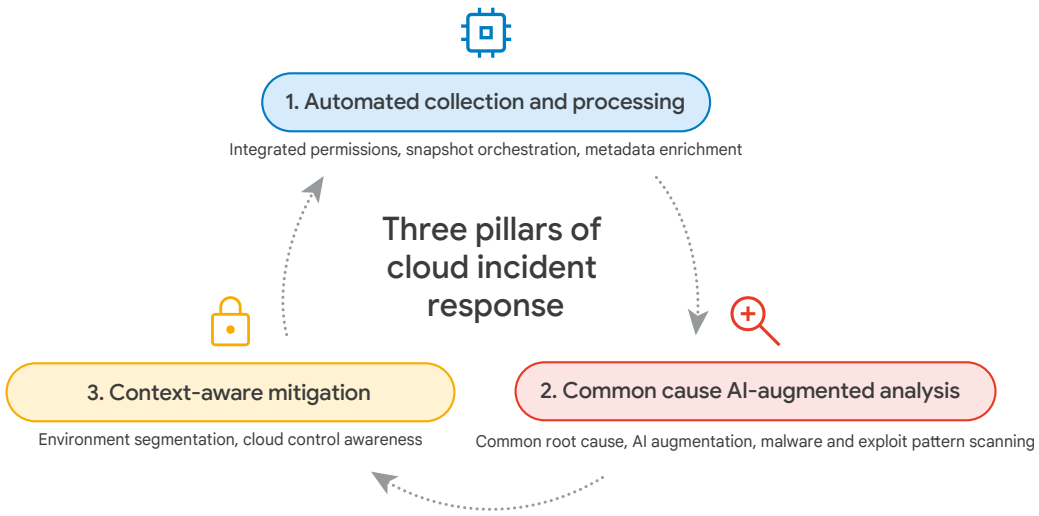

1. **Automated Collection and Processing:** Evidence in the cloud is ephemeral. If a compromised instance is auto-scaled away or deleted by a self-healing cluster, the evidence is lost with it.
    a. Integrated Permissions: The response pipeline should have pre-provisioned, high-integrity access. Waiting for an IAM approval during a breach is a recipe for failure.
    b. Snapshot Orchestration: Use automated tooling to trigger disk snapshots at the moment of detection. In Google, open source tools like dfTimewolf and OpenRelik / Turbinia are used to automate the acquisition and processing of Compute Engine snapshots. These tools can be deployed in Google Cloud via the OSDFIR infrastructure project.
    c. Metadata Enrichment: Every alert should be enriched with cloud-native context (e.g., VM type, IAM roles, network tags), additional data around vulnerabilities, and network configurations of the suspected compromised instance.
2. **Common Cause AI-Augmented Analysis:** The sheer volume of cloud logs—often reaching millions of lines for a single server—makes manual searches extremely difficult.
    a. Common Root Cause: Automated tooling should identify common cloud compromise vectors (e.g., weak passwords, misconfigurations, exposed APIs) using tools like password cracking or bruteforce analyzers, streamlining the investigation for the analyst.
    b. AI Augmentation: Enhance the pipeline with AI-augmented analysis capabilities via AI agents powered by a security LLM. These capabilities should follow a strict feedback mechanism of investigative questions, answering the question and providing log references for verification. AI capabilities currently being explored in the security industry include agents to query logs and answer investigative questions, analyze configuration files for common misconfiguration and insights, and summarize findings across all of the agentic and traditional forensic evidence to provide a comprehensive report for analysts.
    c. Malware and Exploit Pattern Scanning: Automated tooling should scan for malware, exploit patterns and beacons, and correlate findings with logs to deliver a likely root cause. By incorporating Google Threat Intelligence Group (GTIG) capabilities—including proprietary YARA-L rules and indicators from Mandiant and VirusTotal—defenders can identify sophisticated state-sponsored patterns and provide responders with an investigation summary before they begin their manual triage.
3. **Context-Aware Mitigation:** Mitigation in the cloud comes with risks and shutting down a production instance can be as damaging as the attack itself.
    a. Environment Segmentation: Ensure response tooling is context-aware by consistently labeling environments (e.g., dev, test, prod). Automated quarantine should be aggressive in test environments while requiring human approval in production environments.
    b. Cloud Control Awareness: Tooling should be aware of cloud-native controls (e.g., suspension protection in auto-scaling groups) that can automatically revert containment actions, leading to the re-compromise of an instance.

Case Study on Neutralizing RCE Using the Pillars

To monitor for incidents across Alphabet-owned cloud organizations, we built a comprehensive AI-augmented response pipeline primarily using cloud-native functions and open-source tooling aligned with the three pillars of cloud incident response.

In a recent Ray.io test instance exploitation event, our automated forensic pipeline reduced the investigation and containment window from days to under 60 minutes, despite our engineering team never having investigated this application before.

- **Context:** A suspicious package was dropped on an instance in a development environment that exhibited high CPU utilization and beaconed to known mining pools.
- **Detection:** Google Cloud's standard abuse detection pipeline triggered a coin mining alert.
- **Access:** Pre-authorized, organization-level permissions (least privilege) enabled immediate investigation with zero latency.

RCE Defense with the Three Pillars of Cloud IR

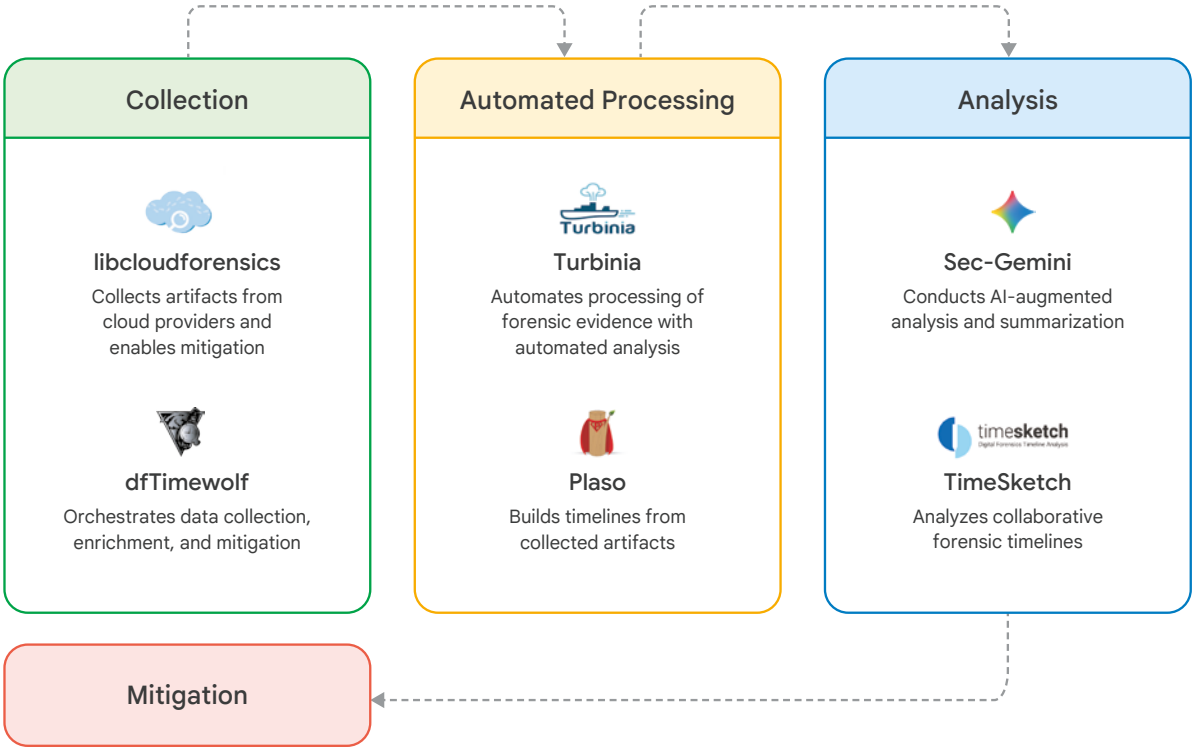

- **Collection:** Automated tooling performed disk acquisition to preserve critical logs and parsed the containerized FUSE/cloud storage environment. Enriched metadata (application, environment type, creation time) instantly flagged the RCE vulnerability and insecure firewall rules.
- **Automated Analysis:** Automated workflows identified cryptominer malware via YARA-L signatures and file hashing. AI agents ruled out common compromise vectors (e.g., passwords, misconfigurations, bruteforce), while malicious domains were tagged with native ray.io logs.
- **Manual Analysis:** Engineers confirmed the root cause was malware execution via a job submitted through an exposed external interface.
- **Mitigation:** Environment-aware tooling isolated the test instance, ensuring no impact on production, and executed remediation following the collection of forensic evidence.
- **Post-Response:** Automated post-mortem reports were issued to the project owner to identify opportunities for internal process improvements. New log parsers and signatures were deployed to cover newly identified gaps.

This pipeline automates manual triage activities and centralizes findings, enabling a shift from reactive to proactive response and significantly reducing investigation time.

Risk Management Recommendations

To counter the rapid pace of cloud-native threats and eliminate the delays inherent in manual triage, organizations should implement a response pipeline built on automated evidence preservation, pre-provisioned identity access, and AI-augmented forensic analysis.

Identity and Access Management	
Platform Agnostic	Google Cloud
Ensure that security response service accounts have the required roles across the entire organization.	IAM Recommender ensures these roles follow the principle of least privilege while remaining ready for emergency use.
Integrate real-time vulnerability scanning to provide context during an investigation.	Security Command Center (SCC) Enterprise helps develop and execute automated remediation playbooks.
Host and Network Security	
Platform Agnostic	Google Cloud
Develop scripts to automatically apply restrictive VPC firewall rules to compromised instances.	VPC Service Controls create perimeters that prevent data exfiltration even after an identity is compromised.
Use AI to identify and block known malicious command-and-control domains at the edge.	Cloud Armor automatically ingests threat intelligence feeds to block malicious traffic before it reaches the workload.

Forensic Tools
Platform Agnostic
The OSDFIR Infrastructure project provides a Kubernetes-native framework that mitigates operational risk by standardizing and automating the deployment of enterprise-grade forensic tools, ensuring consistent incident response capabilities across multi-cloud and hybrid environments.
Visibility and Proactive Threat Detection
Platform Agnostic
Integrate all log sources (e.g., platform, network, OS) into a single chronological view.
Automatically hash every file on an acquired disk and compare it against known malicious databases.
Use organizational policies to enforce labeling of all cloud resources. This ensures that incident response scripts can automatically distinguish between a disposable sandbox and a mission-critical database.
Platform Cloud
Google Security Operations (SecOps) normalizes disparate data streams from different cloud service providers into a unified, searchable timeline. SecOps has Gemini integration to assist with organizing and investigating timelines.Google Cloud's well-architected framework security pillar recommends effective incident management and response processes.

Contributors

Jason Bisson

Jorge Blanco

Anton Chuvakin

Ollie Green

Crystal Lister

Angelus Llanos

Keith Lunden

Eduardo Mattos

Noah McDonald

Bob Mechler

Joachim Metz

Muhammad Muneer

Michael Robinson

Seth Rosenblatt

Matthew Siuda

John Stone

Lia Wertheimer

Executive Resource Addendum

Bridging the gap between technical cybersecurity findings and board-level risk management is critical. The **Board Edition: Cloud Threat Horizons Report** serves as a concise companion to the full-length report, translating complex cloud security data into strategic implications and actionable insights. Use this version to equip your board and C-suite with the context they need to support your security roadmap and drive high-level resilience.

- View the H1 2026 Board Edition Cloud Threat Horizons Report
- Explore the Board of Directors Insights Hub

Google Cloud