



BGD e-GOV CIRT

TLP: CLEAR



CYBER THREAT ADVISORY

**Agentic-Driven Cloud Attack
Using Compromised Service
Principals**

TLP: CLEAR**Distribution:** Public**Advisory on:** Agentic-Driven Cloud Attack Using Compromised Service Principals**Severity:** High**Threat Type:** Cloud Account Compromise, Cloud Resource Destruction, Credential Access, Ransomware-Aligned Activity**Potentially Affected Technology:** Cloud infrastructure, workload identities/service principals, object storage, databases, Key Vault/secrets, virtual machines, application services and backup/recovery resources**Date:** 01 October 2026

Executive Summary

Cyber Threat Intelligence of BGD e-GOV CIRT is issuing this advisory regarding a recently documented **automated cloud attack campaign** in which compromised **service principal's/workload identities** were used to perform reconnaissance, resource discovery, destructive operations and credential collection across a cloud environment.

The investigated activity involved **two compromised service principals**. One primarily performed reconnaissance, while the second performed extensive discovery followed by destructive operations and credential collection. The attacker enumerated virtual machines, subscriptions, resource groups, storage accounts, application services and other cloud resources before initiating a rapid destructive sequence.

During the destructive phase, the attacker attempted **100+ storage-account deletions**, deleted a Key Vault, Function App and App Service plan, attempted database deletion, and attempted to interfere with backup and recovery protection mechanisms. Approximately 30 minutes later, the compromised identity successfully performed **30+ storage-account key retrieval operations**, potentially providing access to sensitive data.

The activity is notable because the destructive actions were performed using **legitimate cloud identities and their existing permissions**, rather than requiring conventional ransomware malware on endpoints. The observed timing and parallel token activity strongly indicate scripted or automated execution.

No ransom notes or successful data exfiltration was confirmed in the investigated activity. However, the combination of resource destruction, recovery-control targeting and credential collection is consistent with activity that could support **ransomware or extortion operations**.

Key Findings

The campaign demonstrates a significant shift from traditional endpoint-centric ransomware toward **cloud-native destructive operations**. The investigated activity demonstrates that **identity compromise can itself become the destructive mechanism**.

- Compromised **service principals** were used as the primary execution identities.
- Cloud reconnaissance preceded destructive activity.
- More than **300 successful read operations** were observed during one reconnaissance period.
- The attacker used multiple tokens for parallel operations.
- More than **150 destructive or credential-related** operations were attempted.
- More than **100 storage-account deletion** attempts were observed.
- Key Vault, Function App and App Service resources were deleted.
- Azure SQL database deletion was attempted but unsuccessful.
- Backup and recovery protection mechanisms were targeted.
- More than **30 successful storage-account key** retrievals were performed.
- Operations followed the permissions already assigned to the compromised identity.
- The activity showed characteristics of highly automated/scripted cloud operations.
- No successful data exfiltration or ransom note was confirmed in the investigated case.

Attack Chain




Figure: Agentic Cloud Attack Lifecycle

This distinction is important because the source confirmed destructive and credential-access activity but did not confirm successful exfiltration.

Initial Access — Compromised Workload Identity

The exact initial compromise mechanism was not confirmed. However, investigators found that the affected service principal's:

- Client ID
- Client secret
- Tenant ID

These had previously been exposed in plaintext through a public code-hosting issue. Although the issue was later edited, the secret remained accessible through its public edit history. The investigation could not confirm that this exposed credential was actually used for the observed attack.

Removing an exposed secret from the visible version of a public repository or issue does not invalidate the credential. If a secret has been publicly exposed, organizations should treat it as compromised until it has been revoked or rotated.

Service Principal Abuse

A service principal is a non-human identity used by applications and automation to access cloud resources. In this campaign, service principals became the attacker's operational identity.




Figure: Compromised Service Principal Architecture (Abuse of a compromised workload identity for cloud reconnaissance, destructive operations and credential collection)

The major security issue is that cloud operations performed using a valid service principal can resemble legitimate application activity unless identity, behavioral and resource context are correlated.




Figure: Two Service Principals and Division of Activity

Cloud Reconnaissance

The first compromised identity performed extensive reconnaissance. Observed enumeration included:

- Virtual Machines
- Subscriptions
- Resource Groups
- Cloud Resources
- Storage Accounts
- Web Applications
- Application configuration stores
- Recovery resources

One reconnaissance phase lasted approximately 15 hours and 30 minutes and generated more than 300 successful read operations.




Figure: Cloud Reconnaissance and Resource Enumeration

Application Configuration Discovery

The second compromised service principal later enumerated App Service configuration stores, potentially searching for exposed credentials or secrets.

It also attempted to identify cloud OpenSearch resources, although that attempt was unsuccessful. This demonstrates the importance of monitoring not only resource creation/deletion but also configuration discovery. Potentially sensitive information in application configuration may include:

- Database credentials
- API keys
- Connection strings
- Storage keys
- Application secrets
- Service credentials

Automated Attack Coordination

The observed timing strongly indicates automated or scripted execution. Five unique tokens were observed for the service principal responsible for destruction and credential collection. Four tokens supported deletion activity, while another handled storage inventory and key retrieval. Two deletion-related tokens were active during the same approximately 70-second period, allowing parallel destructive activity.




Figure: Discovery Before Destruction Timeline

This is a key detection concept as many cloud API operations in a short time window from one workload identity may indicate automation or abuse.

Destructive Operations

The destructive sequence began less than one second after an unsuccessful `ListKey` request against a non-existent storage account. The attacker subsequently attempted more than 150 destructive or credential-related operations over approximately 35 minutes.

The main destructive sequence lasted approximately seven minutes. More than 100 storage-account deletion attempts were made, with most targeted storage accounts successfully deleted.




Figure: Seven-Minute Destructive Sequence

Storage Account Destruction

Storage accounts were a primary target. Some deletion attempts were blocked by:

- Resource locks
- Storage-level deletion protection

This demonstrates the value of independent protection mechanisms that remain effective even when an identity has broad administrative privileges.

Application Resource Destruction

The attacker also deleted:

- Key Vault
- Function App
- App Service plan

These resources belonged to the same resource group and appeared to support the Function App. The attacker also attempted to delete multiple SQL databases, although those deletion attempts failed because of an unsupported API version for the database resource type.




Figure: Multi-Resource Destruction

Use different visual states:

- Deleted
- Attempted
- Blocked

This prevents successful and unsuccessful actions from being incorrectly presented as equivalent.

Backup and Recovery Inhibition




Figure: Independent Recovery Protection

The attacker also attempted to remove or interfere with recovery protections. Targets included:

- Azure Site Recovery-related locks
- Azure Backup protection locks
- Storage resources associated with backup/recovery functions

This behavior is particularly important because an attacker does not necessarily need to encrypt data if they can destroy both production resources and recovery mechanisms.

Credential and Key Collection

Approximately 30 minutes after the final destructive activity, the same service principal performed storage inventory and issued more than 30 successful `ListKeys` requests.

These operations requested storage-account access keys, including keys associated with recovery-related storage accounts.



Figure: Credential Collection Through `ListKeys` (Post-destruction storage-key collection enabling potential access to cloud data)

Permission Abuse

The attacker's actions were largely constrained and enabled by the permissions already assigned to the compromised service principal. Observed permissions included roles equivalent to:

- Storage Account Contributor
- Contributor access to application resources
- SQL Database Contributor

The destructive storage operations were authorized through existing role assignments.




Figure: Permission-to-Impact Mapping (Risks of Excessive Cloud Privilege)

Publicly Exposed Credentials

The investigation identified a potential credential-exposure path through a publicly accessible issue containing:

- Client ID
- Client Secret
- Tenant ID

The issue was later modified, but historical versions still exposed the secret. Organizations should therefore search for exposed cloud credentials:

- Public repositories
- Public issue trackers
- Source-code archives
- CI/CD logs
- Configuration files
- Developer workstations
- Build artifacts
- Documentation repositories

Public-Facing Application Probing

Separate from the confirmed service-principal activity, infrastructure associated with the campaign was observed probing public cloud-hosted applications. The probing targeted paths associated with:

- WordPress administration
- PHP-CGI
- LangFlow code-validation functionality
- Web-shell-like paths

The investigated App Service targets did not overlap with the affected subscriptions, and no App Service-to-cloud-management credential path was identified for the affected tenant. This should therefore be treated as related reconnaissance activity rather than confirmed initial access to the affected environment.

Agentic / Automated Attack Characteristics

The activity demonstrates characteristics associated with agentic or highly automated cloud operations:

1. **Continuous reconnaissance:** Long-duration resource enumeration.
2. **Parallel execution:** Multiple tokens performing different operations simultaneously.
3. **Rapid decision-making:** Transition from reconnaissance to destructive activity within seconds.
4. **API-driven operations:** Cloud resources were manipulated through legitimate management interfaces.
5. **Multi-resource targeting:** Storage, applications, databases, secrets and recovery infrastructure were targeted.
6. **Adaptive behavior:** The actor attempted different operations when particular resource actions failed.

The important defensive implication is that attack speed may exceed traditional manual SOC investigation cycles.

Token Allocation & Function

5 Unique Tokens Identified:
A single service principal generated five distinct tokens for specialized malicious activities.

4 Deletion vs. 1 Inventory Token:
Most tokens focused on destruction, while one handled storage inventory and key retrieval.



Destruction Tokens			Reconnaissance Token
Token Type	Quantity	Primary Operations	
Destruction Tokens	4	File and resource deletion	Reconnaissance Token
Reconnaissance Token	1	Storage inventory and ListKeys retrieval	

The Power of Concurrency



Figure: Automated / Parallel Token Activity

Attack Timeline



Figure: Observed Attack Timeline

This figure is particularly useful for demonstrating how **slow reconnaissance can transition into extremely rapid destruction**.

MITRE ATT&CK Mapping

Tactic	Technique	ID	Observed Activity
Initial Access	Exploit Public-Facing Application	T1190	Probing public cloud applications
Persistence/Access	Valid Accounts: Cloud Accounts	T1078.004	Compromised service principals
Discovery	Cloud Service Discovery	T1526	Enumeration of cloud resources
Impact	Data Destruction	T1485	Deletion of storage/application resources
Impact	Inhibit System Recovery	T1490	Targeting backup/recovery protections

Indicators of Compromise (IoC)

The following IP addresses were identified in connection with probing and malicious cloud-management requests:

Indicator	Type	Description
45.131.66[.]106	IPv4	App Service probing / malicious cloud-management requests
34.153.223[.]102	IPv4	App Service probing
64.20.53[.]230	IPv4	App Service probing

These IP addresses should be used as supplementary indicators. Detection should not rely exclusively on static IP blocking because the observed attack primarily **abused legitimate cloud identities and APIs**.

Cloud Detection Strategy

Traditional endpoint-focused monitoring is insufficient for this attack pattern. BGD e-GOV CIRT recommends monitoring the following telemetry:




Figure: Cloud Detection Architecture

Identity

- Service-principal authentication
- Token issuance
- Token frequency

BGD e-GOV CIRT

- Source IP
- Geographic anomalies
- New authentication locations
- Unusual authentication times

Cloud API


Monitor: (with particular attention to high-risk resources)

- List*
- Get*
- Delete*
- ListKeys
- Create*
- Update*

Resource Activity

Monitor:

- Storage deletion
- Key Vault deletion/access
- Database deletion
- Application deletion
- Recovery-resource modification
- Backup-policy modification
- Resource-lock removal



High-Risk Behavioral Correlation




Figure: Cloud Attack Detection Matrix

Recommended Actions

Organizations operating cloud environments should:

- **Inventory all service principal's/workload identities** and their permissions.
- Immediately revoke and rotate any publicly exposed client secrets.
- Review service-principal authentication and token activity for unusual source IPs, locations and timing.
- Apply **least privilege** to workload identities.
- Remove unnecessary Contributor-level permissions.
- Restrict destructive operations where operationally feasible.
- Enable independent **resource locks and deletion protections**.
- Protect backup and recovery resources using separate administrative controls.
- Monitor `Delete*`, `ListKeys`, resource-lock and recovery-control operations.
- Review cloud audit logs for unusual bulk operations.
- Hunt for the listed IP indicators across cloud, firewall, proxy and DNS telemetry.
- Review public repositories, issue trackers and CI/CD systems for exposed cloud credentials.

The investigation demonstrates that independent deletion protection can prevent some destructive actions even when an attacker controls a highly privileged identity.

Bangladesh-Specific Considerations

Organizations in Bangladesh increasingly operate workloads across:

- Government cloud environments
- Data centers
- Financial services
- Telecommunications
- Energy and utilities
- Healthcare
- Education
- E-commerce
- Software and technology services
- Critical infrastructure

Organizations should therefore assess whether service principals or workload identities have excessive permissions across production, backup, database and recovery environments.

For government and critical infrastructure environments, particular attention should be given to:

- Cloud Identity
- Application
- Database
- Storage
- Backup
- Recovery

because compromise of a single workload identity could potentially provide access across multiple interconnected cloud resources.




Figure: Attack-to-Defense Model (Securing The Cloud Lifecycle)

The investigated activity demonstrates how compromise of a cloud workload identity/service principal can enable an attacker to conduct reconnaissance, access sensitive credentials and perform destructive operations without relying on conventional ransomware deployment.

The most significant defensive lesson is that cloud identity security, API monitoring, least privilege and independent recovery controls must be treated as core components of ransomware defense.

Organizations should not assume that a cloud environment is protected simply because endpoint ransomware controls are deployed. A compromised workload identity with excessive permissions can potentially manipulate cloud resources directly through legitimate management interfaces.

References

- <https://www.microsoft.com/en-us/security/blog/2026/09/25/storm-3168-agentic-driven-cloud-attacks-using-compromised-service-principals/>
- <https://www.sysdig.com/blog/jadepuffer-agentic-ransomware-for-automated-database-extortion>



BGD e-GOV CIRT