Measuring Attack Observability in Cloud Telemetry Logs: A Cross-Platform Analysis

Mary Grace Dhooghe*, Minkyung Park*, Junghwan Rhee†,

Yung Ryn Choe‡, Chung Hwan Kim*

International Conference on Dependable Systems and Networks

June 22-25, 2026

*University of Texas at Dallas, †University of Central Oklahoma, ‡Sandia National Laboratories

Scenario

Context:

You are a cloud service customer.

Limited or No access to Machine Level Logs

Is data being stolen or destroyed?

Are cloud services being used in an unauthorized capacity?

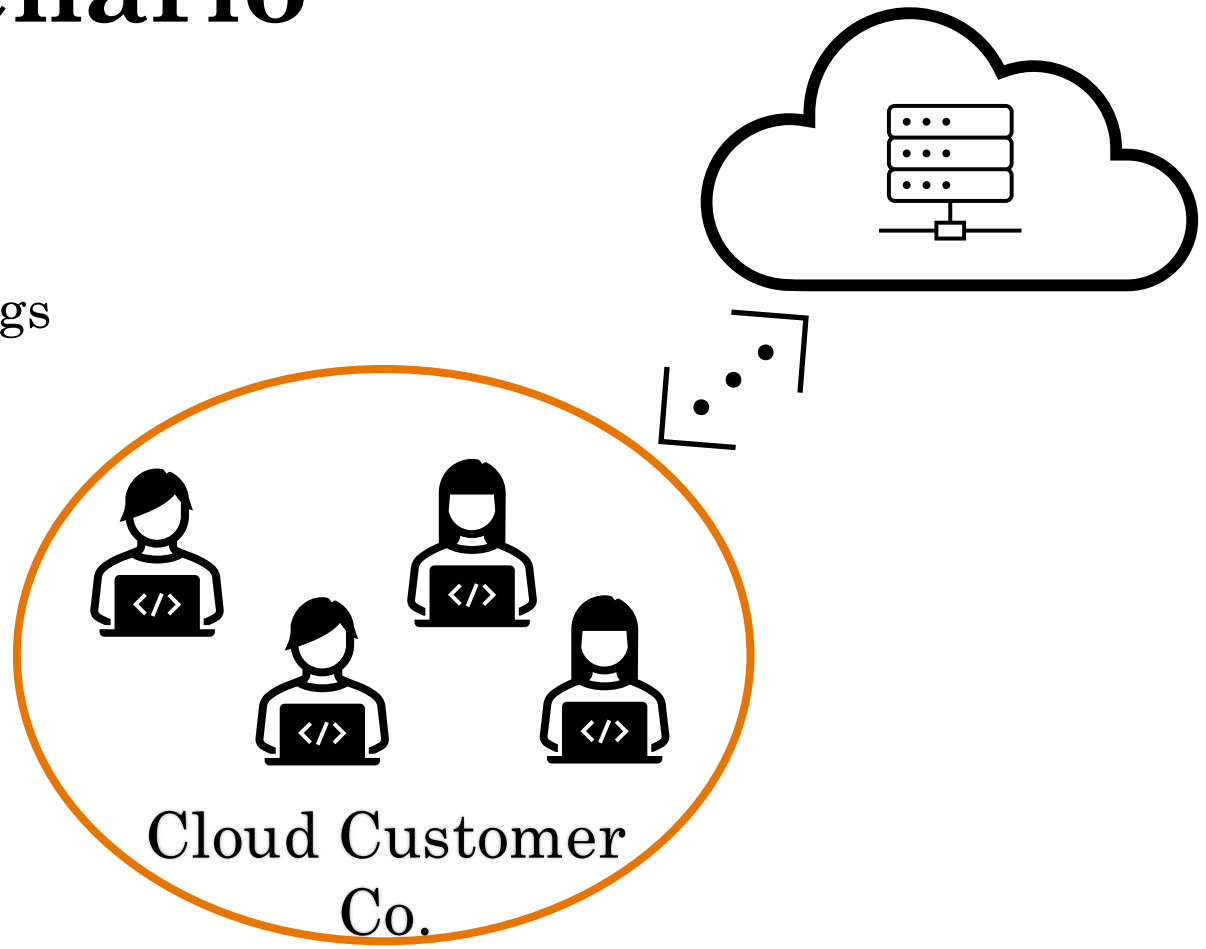

Question:

How can security-relevant activity within cloud environments be monitored?

Cloud Attacks

Traditional Attacks

On-Premise Environments

Cloud-Specific Vulnerabilities

Management Interfaces

Network Protocols

APIs

Etc.

Consequences

Service Disruptions

Unauthorized Access to Sensitive Data

Data Loss

Financial Losses

Etc.

Cloud Security

Existing Cloud Security Countermeasures

AWS GuardDuty, Google Security Operations, Microsoft Defender for Cloud

Often Require Additional Resources

Cloud Telemetry Logs

Collected by Default

No required Extra Infrastructure

Compliment Available Security Measures

Cloud Telemetry Logs

What Are Telemetry Logs?

Time Stamped Event Logs

System-level Logs

Call Stacks

System Call Arguments

Etc.

Classically used for security

Ex. Host-Based Intrusion Detection Systems (HIDS)

What Are Cloud Telemetry Logs?

Subscription Level Event Logs

Designed to Track:

Performance

Resource Usage

Accounting, etc.

Information Includes

Ex. User IDs & Resource IDs

Application & Network Activities

Cloud Telemetry Logs for Security: Overview

MITRE ATT@CK Framework

Scripted Attacks

Subscription Level

Cloud Providers

Google, Amazon, and Microsoft

Resources

Ex. Virtual Machines, Storage, Key Management, Identity Access

Log Collection

Default – Enabled without Setup

Additional – User Initialized

Cloud Telemetry Log Information

Single AWS Log

```
{
  "eventVersion": "1.11",
  "userIdentity": {
    "type": "AWSService",
    "invokedBy": "cloudtrail.amazonaws.com"
  },
  "eventTime": "2025-01-02T16:54:02Z",
  "eventSource": "kms.amazonaws.com",
  "eventName": "GenerateDataKey",
  "sourceIPAddress": "cloudtrail.amazonaws.com",
  "userAgent": "cloudtrail.amazonaws.com",
  "requestParameters": {
    "keySpec": "AES_256",
    "encryptionContext": { ... },
    "keyId": "arn:aws:kms:us-east-1:11111111:key/2222..."
  },
  "responseElements": null,
  "requestID": "33333333-3333-3333-3333-",
  "eventID": "44444444-4444-4444-4444-444444444444",
  "resources": [
    {
      "accountId": "11111111",
      "type": "AWS::KMS::Key",
      "ARN": "arn:aws:kms:us-east-1:11111111:key/2222..."
    }
    ...
  }
```

	AWS	Azure	GCP
Action	8	2	3
Success	3	2	0
Caller	2	3	3
Event ID	4	0	1
Log ID	0	1	2
Log Type	1	2	3
Resource	8	5	2
Time	2	2	1
Total	28	17	15

Threat Model

Assumptions:

1. Attacker does not have access to the cloud provider's system
2. Attacker has no ability to edit the logs at rest or in transit
3. Generated logs are accurate and retain their integrity throughout creation, transit, and storage

Attacks to the Subscription

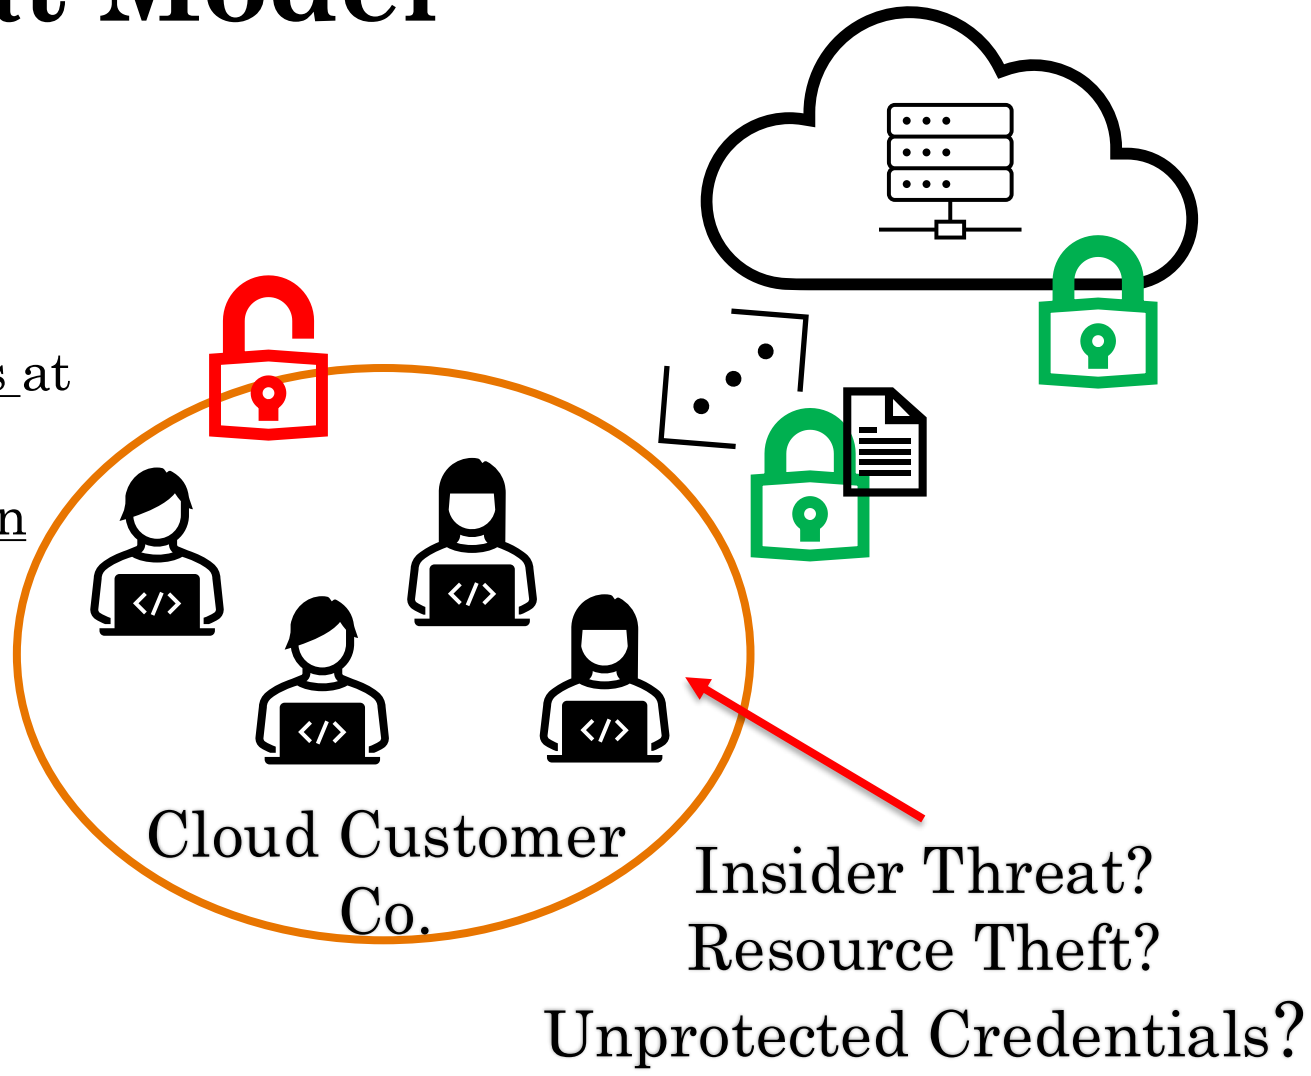

Cloud Platforms & Resources

3 Cloud Providers

7 Resource Types

Resource	Google Cloud Platform (GCP)	Amazon Web Service (AWS)	Microsoft Azure
Event Subscription	Eventarc	Lambda	Function App. Logic App.
Identity and Access Management	IAM	IAM Users	AD Service Principals
Key Management Service	KMS	KMS	Storage Encryption
Secrets Manager	Secrets Manager	Secrets Manager	Key Vault
Database (SQL)	SQL MySQL	RDS MySQL	SQL Serverless
Storage	Cloud Storage	S3	Storage
Virtual Machine	Compute Engine	EC2	Virtual Machines

MITRE ATT@CK Framework

Cloud-Executable

Performed in the victim's cloud environment

Reproducible

performed consistently on cloud resources with publicly available tools

Ethically Feasible

attacks on user-controlled instances

35 MITRE ATT@CK (11 Categories)

MITRE Category	MITRE Attacks	Resources
Collection	Archive Collected Data, Automated Collection	Storage, KMS, Virtual Machine
Credential Access	Brute Force, Credentials from Password Stores	IAM, Secrets Mgr., Storage
Command & Ctrl.	Tool Transfer, Remote Access Software	Virtual Machine
Defense Evasion	Valid Accounts, Alternative Authentication,	IAM, Storage
Execution	Cloud Administration Command,	Virtual Machine, IAM
Exfiltration	Automated Exfiltration, Scheduled Transfer	Storage, Event Subscription
Impact	Account Access Removal, Data Destruction, Disk Wipe, Denial of Service	IAM, Storage, SQL, Virtual Machine
Initial Access	Drive By Compromise, Exploit Public Application	SQL, Virtual Machine
Lateral Movement	Exploit Remote Service, Alternative Auth.	Virtual Machine, IAM, Storage
Persistence	Create Account, Implant Image, Traffic Signaling	IAM, Virtual Machine, Event Sub.
Privilege Escalation	Access Token Manipulation	IAM

Research Questions

RQ.1: Do cloud telemetry logs provide (a) consistent and (b) relevant data for detecting malicious activity?

RQ.2: Do non-default logs provide additional insights into malicious activity?

RQ.3: Can attacks evident in cloud telemetry logs be automatically detected?

Data Collection

Control & Attack Scenarios: *scripted* individually for each cloud provider using their CLI.

35 (Attacks) x 2 (Control/Attack) x 3 (Platforms) = 210 Scripts

x 2 (Default/Additional Logs) x 10 (Trials) = 4200 Datasets

Control Scenario

- Resource Creation
- Targeted resources initiated
- Attacker accounts initiated
- Resource Deletion
- Targeted resources removed
- Attacker accounts removed
- For testing consistency

Attack Scenario

- Resource Creation
- Payload Actions
- Targeted Resource Attack Actions
- Performed as a user with the corresponding access level
- Resource Deletion

Attack Visibility

$U_{control}$

(Access Token Manipulation)

$Time_{Delta}$	Log Name	Method	…	Resource
0	Activity	Create Account	…	Projects/A
65	Activity	Delete Account	…	Accounts/**

Attack Logs (Bold)

(Access Token Manipulation)

$Time_{Delta}$	Log Name	Method	…	Resource
0	Activity	Create Account	…	Projects/A
3	**Data Access**	**Get IAM Policy**	**…**	**Accounts/****
9	**Activity**	**Create IAM Key**	**…**	**Accounts/****
70	Activity	Delete Account	…	Accounts/**

Visible:

Payload execution results in log entries that differ from those generated under equivalent control conditions

$U_{control}$:

Union of all unique logs generated in any of the Control trials

Attack Logs:

Difference of $Attack_{a,i}$ - $U_{control}$

RQ.1: Do cloud telemetry logs provide consistent data for detecting malicious activity?

	Default			Additional		
Platform	$avg(\#_{pa})$	$std(\#_{pa})$	m_{ratio}	$avg(\#_{pa})$	$std(\#_{pa})$	m_{ratio}
GCP	10.3686	0.0606	1.0672	29.1714	1.0657	1.0698
AWS	8.4667	0.6450	1.4361	24.2114	0.9773	1.3942
Azure	9.2387	2.7240	1.8409	11.8743	1.9760	1.6942

$\#_{pa}$: *Number of Control Logs for Each Platform (p) and Attack (a)*

$m_{ratio} = \frac{m_{pa}}{avg(\#_{pa})}$: ratio of unique logs to average count of logs - Range:[1,10]

m_{pa} : size of $U_{control}$ for each platform and attack

All platforms are highly consistent ($m_{ratio} < 2$)

RQ.1: Do cloud telemetry logs provide relevant data for detecting malicious activity?

% visible by Resource

Visible: Any Attack Specific Logs

Def.: Default

Add.: +Additional

Resource	GCP ↑		AWS ↑		Azure ↑	
	Def.	Add.	Def.	Add.	Def.	Add.
Storage	0.56	0.94	0.6	0.93	0.6	1
IAM	0.58	0.75	0.58	0.67	0.5	0.67
VM	0.31	1	0.3	0.3	0.4	0.4
SQL	0	1	0	1	0	0
Event Sub.	1	1	1	1	1	1

Default Findings:

Storage + IAM > 50%

Additional Findings:

VM (AWS/Azure) < 50% (No Improvement)

SQL (Azure) = 0 => Exact Resource Diff.

RQ.1: Do cloud telemetry logs provide relevant data for detecting malicious activity?

Visibility Level: What amount of payload ***Action, Actor, and Affected Resources*** are identifiable from the payload-generated logs?

1. Collect Attack Logs
2. Cluster semantically similar actions

   Ex. (“create”,”insert”), (“disable”,”delete”)
3. Compare Payload Script CLI commands to Attack Logs for each Platform and Attack
4. “Full”/”Partial” if All/Some commands have a Log

Default

GCP:

20% (Full)

20% (Partial)

AWS:

20% (Full)

11% (Partial)

Azure:

3% (Full)

26% (Partial)

RQ.2: Do non-default logs provide additional insights into malicious activity?

Visibility Level: What amount of payload ***Action, Actor, and Affected Resources*** are identifiable from the payload-generated logs?

1. Collect Attack Logs
2. Cluster semantically similar actions

   Ex. (“create”,”insert”), (“disable”,”delete”)
3. Compare Payload Script CLI commands to Attack Logs for each Platform and Attack
4. “Full”/”Partial” if All/Some commands have a Log

	Default	**Additional**
GCP:	20% (Full) 20% (Partial)	49% (Full) 40% (Partial)
AWS:	20% (Full) 11% (Partial)	40% (Full) 31% (Partial)
Azure:	3% (Full) 26% (Partial)	29% (Full) 23% (Partial)

RQ.3: Can attacks evident in cloud telemetry logs be automatically detected?

Variational Auto-Encoder (VAE)

Encode → Decode

Decode Acc. < Threshold

→ Anomaly

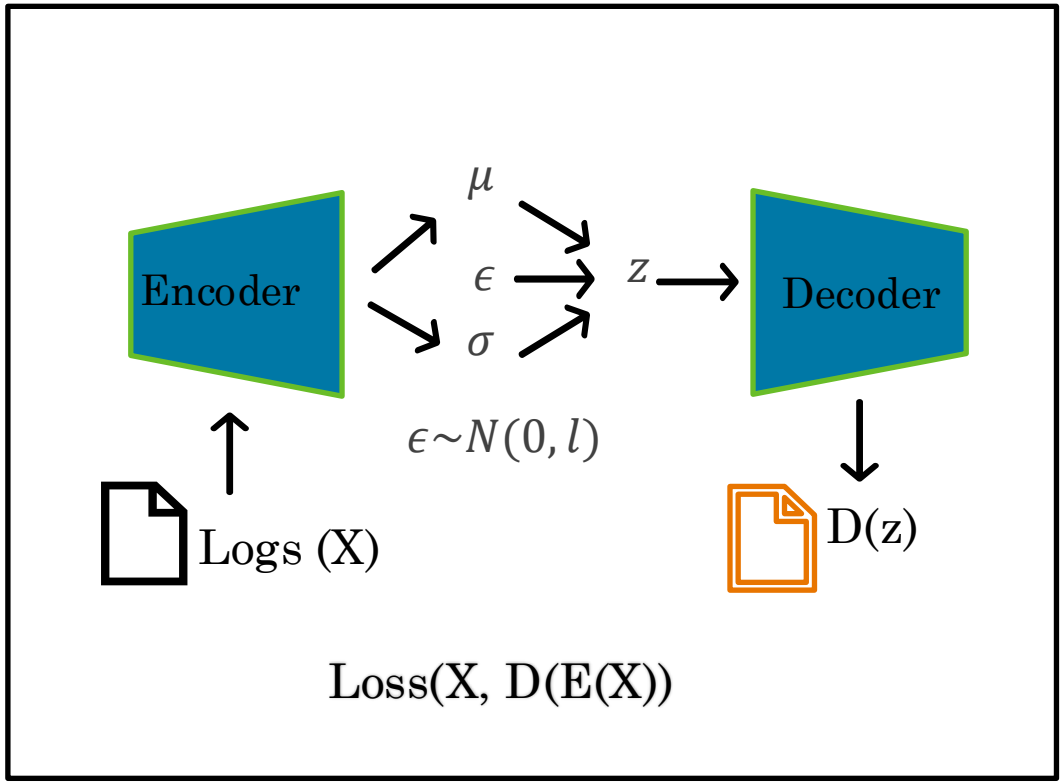

RQ.3: Can attacks evident in cloud telemetry logs be automatically detected?

Variational Auto-Encoder (VAE)

Encode → Decode

Decode Acc. < Threshold

→ Anomaly

Log Labels

T.P: True Positive

F.P: False Positive

Acc: Accuracy

Platform	Log	T. P.	F. P.	Acc.	Vis.	Inv.	Ctrl.
GCP	Def.	0.654	0.044	0.977	9/15	0/17	1/32
	Add.	0.294	0.105	0.861	16/30	0/2	4/32
AWS	Def.	0.409	0.080	0.909	8/16	0/13	2/30
	Add.	0.227	0.095	0.838	11/24	1/9	5/33
Azure	Def.	0.325	0.004	0.953	7/16	2/15	2/31
	Add.	0.408	0.103	0.868	12/23	1/11	4/34

35+% of Visible Attacks Detected in All Scenarios

Many Attacks only Produce 1 or 2 Attack related Logs

Default Scenario: This Analysis is Approx. Free

Summary

35 Attacks on **GCP/AWS/Azure** targeting **7** resource types & **2** log configs.

- Cloud Telemetry Logs security is Complimentary
- Logs are found to be consistent and relevant
- Non-Default logs increase the number of attacks detectable
- Details and Case Study included in paper & artifacts

- https://gitlab.com/s3lab-code/public/telemetry
- https://doi.org/10.5281/zenodo.19228517

Contact:
Mary.Kozuch@utdallas.edu
ChungKim@utdallas.edu