Cloudy with a Chance of DFIR

Rachel Kang

Manager, Digital Forensics and Incident Response

Aon Cyber Solutions

The Exchange Q2

June 20, 2024

Current State of Cloud Computing

% of companies using cloud (public, hybrid, private, etc.)	Global market growth in 2022[1]	Global market growth in 2023[1]	Expected market growth by 2032[1]
87%	$480 bill.	$559.2 bill.	$2,297.3 bill

*across Fortune 500 companies

- All types of cloud models (public, hybrid, private) and services (SaaS, PaaS, IaaS) are seeping into all industries, including banking, finance, education, health, etc.
- Benefits to business continuity outweigh the detriments:

✓ Cost reduction in operational expenditure

✓ Increased productivity + collaboration across remote work

✓ Flexibility + scalability of models

✓ Sustainability

✓ Overall workflow efficiency

✓ Enhanced data security control options

Current State of Cloud Computing

% of companies using cloud (public, hybrid, private, etc.)	Global market growth in 2022[1]	Global market growth in 2023[1]	Expected market growth by 2032[1]
87%	$480 bill.	$559.2 bill.	$2,297.3 bill

*across Fortune 500 companies

- All types of cloud models (public, hybrid, private) and services (SaaS, PaaS, IaaS) are seeping into all industries, including banking, finance, education, health, etc.
- Benefits to business continuity outweigh the detriments:

✓ Cost reduction in operational expenditure

✓ Increased productivity + collaboration across remote work

✓ Flexibility + scalability of models

✓ Sustainability

✓ Overall workflow efficiency

✓ Enhanced data security control options

How does this migration to cloud platforms render our DFIR (Digital Forensics + Incident Response) investigations more complex?

AON

Shift in Traditional Paradigm

20 years ago → Keep the bad guys **out** and keep the data **in** the premises

- Single ingress point is the network/firewall, with log streams flowing thru this point
- Higher confidence in visibility across information
- Data and logs resides on physical servers inside the building

Today → Keep data **in** third-party's hands, remotely accessible by (potentially) anyone at anytime from any place

- Thousands of ingress points on top of physical + virtualized networks
  - "*In the public cloud, identity is the new perimeter*"
- Unclear specifics of environment structure
- Data and logs are scattered across multiple regions and locations on different machines and storage devices, both physical and virtual

AON

Shift in Traditional Paradigm

20 years ago → Keep the bad guys **out** and keep the data **in** the premises

- Single ingress point is the network/firewall, with log streams flowing thru this point
- Higher confidence in visibility across information
- Data and logs resides on physical servers inside the building

Today → Keep data **in** third-party's hands, remotely accessible by (potentially) anyone at anytime from any place

- Thousands of ingress points on top of physical + virtualized networks
  - "*In the public cloud, identity is the new perimeter*"
- Unclear specifics of environment structure
- Data and logs are scattered across multiple regions and locations on different machines and storage devices, both physical and virtual

How do we use our traditional DFIR knowledge to tackle the complexities within Cloud IR?

AON

Understanding Cloud Log Sources

Logs are still the source of truth

DFIR in on-prem: “Let’s go get a memory dump of the system”

DFIR in the cloud: “May I please get a log of that activity, if available?

Still reviewing ***disk***, ***network*** and ***memory*** data in cloud IR to identify the root cause and scope of cloud breach.

But there are challenges....

- Not all cloud logs are enabled by default.
  - GCP Data Access Audit Logs are disabled by default (except BigQuery) for easier overhead auditing. No logging across “ADMIN_READ”, “DATA_READ”, AND “DATA_WRITE” operations for any GCP services available unless enabled. [2]
- Some logs are at the mercy of cloud providers (Shared Responsibility Model)
- Permanent loss of logs across ephemeral resources (ex. containers + serverless architecture)

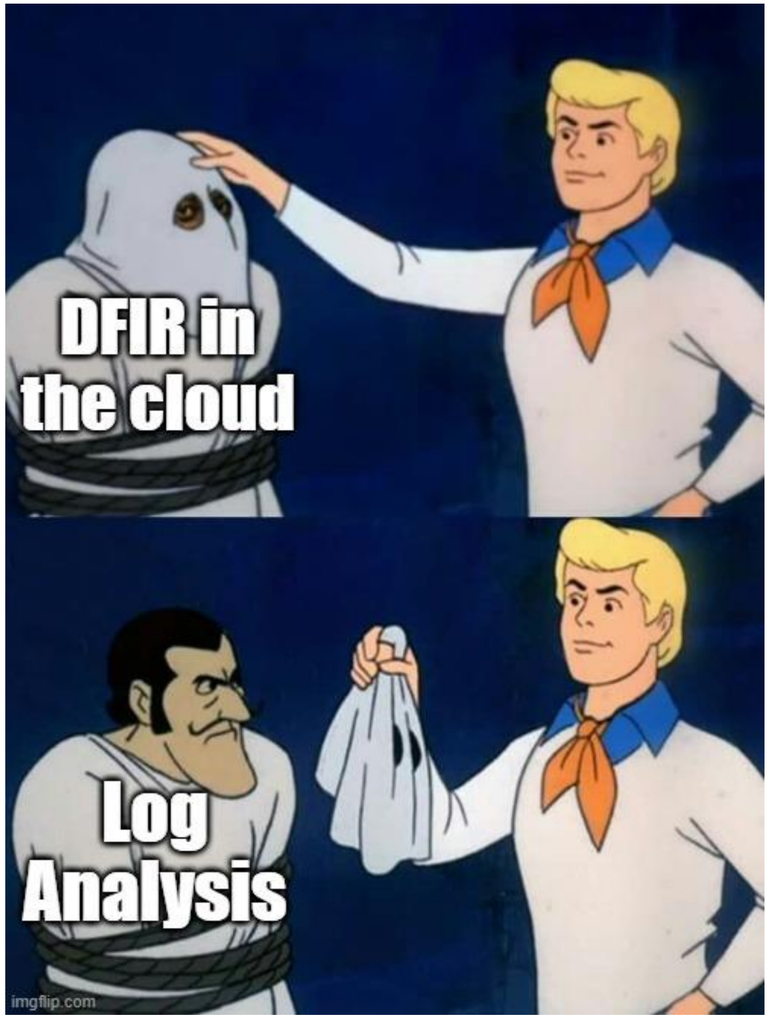

Log Collection Comparison

DFIR Approach

1. **Capturing Volatile Evidence (Memory)**
   - Memory Acquisition (FTK, RAM Capture)
2. **Network / Logs**
   - Export logs from physical systems, proxies
3. **Disk Imaging**
   - Logical (local or remote)
   - Physical (datacenter)
4. **Analysis**
   - Transfer evidence to secure storage and/or forensics lab

Cloudy Approach

1. **Capturing Volatile Evidence (Memory)**
   - AWS Systems Manager
   - Azure Monitor
2. **Network / Logs**
   - AWS CloudWatch, AWS VPC Traffic Mirroring, AWS Costs
   - Azure Monitor, Azure Network Watchers, Azure Billing
3. **Disk Imaging**
   - Snapshot instances / Services
4. **Analysis**
   - Transfer evidence using cloud provider and/or third-party tools

AON

Sample Log Sources Across CSPs

Leading Cloud Service Providers (CSPs) such as AWS, Azure, GCP

Cloud Technology[3]	Amazon Web Services (AWS)	Microsoft's Azure	Google Cloud Platform
Management Console	Console	Portal	Console
Authentication Services	Directory Service	Active Directory	Cloud Identity
Virtual Machine	Elastic Compute Cloud (EC2)	Virtual Machine	Compute Engine Virtual Machine
File Storage	Simple Storage Solution (S3)	Blob Storage	Cloud Storage
Networking	Virtual Private Cloud	Virtual Network	Virtual Private Cloud
Logging Platform	Athena	Monitor, Microsoft Sentinel, Log Analytics	Log Explorer
Log Analysis Format	CloudTrail, VPC Flow Logs	AD Audit, Sign-In, Resource, Activity, and NSG Flow Logs	Audit Logs, VPC Flow Logs
Database Services	DynamoDB, Aurora, Relational Database Service	Database, SQL Database	Datastore, Cloud Bigtable, Cloud SQL
Email	Simple Email Service	Microsoft 365	Google Workspace
Code Repositories	CodeCommit	Repos	Cloud Source
Containers	Elastic Kubernetes Service	Kubernetes Service	Kubernetes Engine
Serverless Functions	Lambda	Functions	Cloud Functions

Impact of Cloud Compromise

A single compromise in cloud environment can cascade across interconnected systems, magnifying impact.

Identity and Access Management (IAM) is the new perimeter in a cloud environment.

- Centralized Data – Unlike segmented on-premises systems, cloud services **centralize** vast amounts of data, potentially exposing extensive sensitive information.
- Rapid Propagation – Automated and interconnected nature of cloud services facilitate **rapid attack propagation**, affecting wide areas of the environment quickly.
- Access and Privilege Escalation – Compromised cloud credentials grants **wide access** across integrated clouds services, compared to isolated on-premises systems.
- Dynamic and Scalable Resources – Cloud environments **dynamically scale resources**, allowing for attackers to leverage this to amplify operations.
- Third-Party Integrations – Extensive usage of third-party services in cloud setup can not just affect internal systems, but also **external partners and services**, as opposed to on isolated on-prem systems.
- Reduced Visibility – Cloud providers control many aspects of infrastructure, and **limited visibility** can delay detection and response.

AON

Case Study: API Key Compromise

Social Engineering: The attacker performs a targeted phishing attack against Bob, a HelpDesk employee with administrator privileges. The attacker attempts to login, triggering an MFA push to Bob's phone. Bob mistakenly accepts the request, allowing the threat actor into the company's Microsoft 365 platform.

Reconnaissance: The attacker identifies that Bob's account has access to the company's GitHub/BitBucket platform. The attacker accesses a repository and identifies hard-coded API keys for a third-party payment processing service (ex. PayPal, Stripe, Square).

Actions Taken Using API Key:

- Initiate **fraudulent transactions**, charging customers' credit cards without their knowledge.
- Process **fraudulent refunds** to accounts, siphoning money from the company.
- Charge customers excessively, causing disputes and **damaging customer trust**.
- **Exfiltrate** historical transaction data and potentially, customer information.
- Manipulate API to create high volume of transactions, causing **denial of service**.

Impact: Financial losses, customer trust, operational disruption, reputational damage

AON

Case Study: Virtual Machine Compromise

VPN Vulnerability: The attacker identifies and exploits a vulnerability in a company's unpatched VPN appliance, gaining access to the internal network. The attacker exploits additional security configurations to escalate privileges to Bob's admin account.

Reconnaissance: The attacker identifies that Bob's account has elevated privileges in the company's Azure tenant. The attacker creates thousands of new VMs for malicious purposes and reconfigures its settings to allow for external RDP connections, accessing them from a command and control (C2) server.

Actions Taken Using VMs:

- Perform network reconnaissance, identifying the company's primary file server and **exfiltrating mass amounts of sensitive data** to their C2 system using a bulk network file transfer utility.
- Perform network reconnaissance, identifying the company's SQL database server. The attacker runs "SELECT *" queries to identify and **exfiltrate from tables containing sensitive data**.
- Deploy cryptomining software across VMs to mine Bitcoin, leading to a **high consumption of compute resources** and **significant financial losses** for the company.
- Deploy hundreds of high-performance VMs, **exhausting available cloud resources** and **disrupting services** for legitimate applications.
- Leverage the VMs to serve as a distribution point for malware, phishing sites, etc. leading to f**urther compromises**.

AON

Are We Prepared For What's Next?

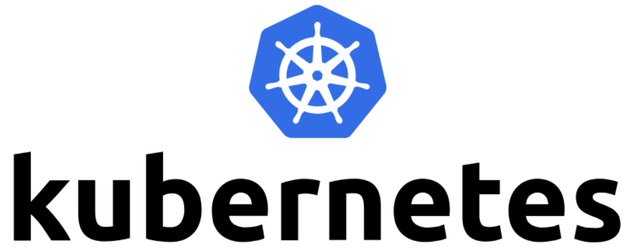

AON

Forensic Readiness in Cloud Incident Response

- Standardize forensic procedures (or workflows) for each leading cloud service provider (CSP), specifically tailored for environments such as AWS, Azure, and GCP.

- Dedicated Cloud IR specialists with deep understanding of the most pressing challenges associated with cloud forensics:

  - ✓ Restricted access to logs (including volatile data)
  - ✓ Locating forensic artifacts in large, distributed, and dynamic environments
  - ✓ Division of security responsibilities between CSP and customer (Shared Responsibility Model)
  - ✓ Complexity of multi-tenancy and cloud architecture
  - ✓ Scalability of existing forensics and IR tools
  - ✓ Encryption of data across transit
  - ✓ Third-party integrations
  - ✓ Compliance and legal challenges

AON

Endnotes

[1] “Cloud Computing Market - Global Industry Analysis, Size, Share, Growth, Trends, Regional Outlook, and Forecast 2023 – 2032.” Precedenceresearch.com, October, 2023. https://www.precedenceresearch.com/cloud-computing-market

[2] “Enable Data Access Audit Logs.” Cloud.google.com, June 3, 2024. https://cloud.google.com/logging/docs/audit/configure-data-access

[3] “Cloud Platform Log Configurations to Consider in Investigations.” Cloud.google.com, May 3, 2023. https://cloud.google.com/blog/topics/threat-intelligence/cloud-bad-log-configurations/

[4] “Forensic Readiness In The Cloud.” Cadosecurity.com. https://www.cadosecurity.com/blog/forensic-readiness-in-the-cloud

[5] “Cloud Digital Forensics: Beyond Tools, Techniques, and Challenges.” mdpi.com, January 10, 2024. https://doi.org/10.3390/s24020433

Questions & Answers