



Digital Forensic Investigation in Cloud Computing Environments: Challenges in Evidence Acquisition and a Proposed Methodology for Multi-Tenant Systems

Student Name: Md Abu Sayeed

Student ID: 2528454

Programme: MSc in Cyber security and digital forensic programme.

Module: Research Methods and Professional Development

Module Code: ACFA7004

Lecturer: Suresh Shanmugsundaram

Internal Verifier: Dr Nitheesh Kaliyamurthy

Submission Date: 14 August 2026 Friday (2PM)

Table of Contents:

1. Introduction

1.1. Context and background

1.2. Research Problem

1.3. Research Aim and Objectives

1.4. Significance of the research

2. Literature Review

2.1. Introduction to the Literature Review

2.2. The Shift from Traditional to Cloud Forensics

2.3. Multi-Tenancy, Hypervisors, and Legal Isolation

2.4. Limitations of Cloud-Native API Acquisition

2.5. Relation to My Study and Literature Summary

3. Research Design and Methodology

3.1 Research Design

3.2 Research Methods and Materials

4. Expected Outcome

5. Timeline

6. References

1. Introduction

1.1. Context and background

The enterprise computing paradigm in the last ten years has been largely shaped by the large-scale move from conventional, physically localized data centers to the more widely-used Cloud Computing Environment. Being propelled by the important business requirements for high availability, large-scale expandability, resource elasticity, and cost efficiency, organizations around the world have embraced the use of cloud architecture, which are generally grouped into the following three primary types of services: IaaS (Infrastructure as a Service), PaaS (Platform as a Service), and SaaS (Software as a Service). The central idea behind cloud computing is that of virtualization and resource pooling, which follows the principle of multi-tenancy. In the multi-tenant scenario, a Cloud Service Provider (CSP) like Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) assigns and manages physical computational, storage, and networking resources for hundreds or even thousands of different, unrelated entities at a time [1]. These different entities, referred to as tenants, are using the identical physical infrastructure components such as blade servers and SANs while being separated from each other through hypervisors and virtual networks [2].

Though such a system offers unparalleled efficiency in operations, it has brought with it devastating complications in the world of DFIR. The principles of digital forensics were developed in times when dead-box forensics was the practice. At this time, the investigator would have had physical access to the server, shut it down, removed the hard drive physically, attached it to a hardware write blocker, and made a bit-for-bit cryptographic copy of the hard drive [3]. With this technique, data integrity was assured, a legal chain of custody established, and the fact that no one had modified the data in the process of gathering it was ensured[4].

1.2. Research Problem

The main research problem that this paper aims to address is the inadequacy in the procedure and technical processes involved in extracting legally admissible and cryptographically valid digital evidence from compromised multi-tenant clouds without infringing on data privacy limitations [5]. In case a cyber attack takes place within a cloud environment, for example, where an Advanced Persistent Threat (APT) exploits a Server-Side Request Forgery (SSRF) weakness to steal IAM credentials and take over an IaaS virtual machine, there emerges a direct 'Acquisition Paradox' [2], [6].

First, the inherent dynamicity and elasticity of cloud computing raise the issue of Evidence Volatility. As a result of auto-scaling cloud infrastructures, when a cloud server suddenly faces a heavy load due to the DoS attack or cryptomining malware, the cloud orchestrator can terminate such a server and start new ones in order to ensure application stability [6]. If a virtual machine is terminated in a multi-tenant cloud, its memory is instantly erased, and all the temporary storage disks are cryptographically wiped and allocated for another user [3]. Thus, by the time when the security incident is discovered, the evidence is already gone due to the automation processes in the cloud.

Moreover, the issue of Multi-Tenant Collateral Damage greatly restricts forensic acquisition. With the on-premises infrastructure, an investigator conducting a memory dump or complete image acquisition of the hard drive will never be concerned with the possibility of obtaining someone else's intellectual property by mistake. On the contrary, the cloud environment involves the use of a hypervisor, such as Xen, KVM, and Hyper-V, to manage the memory and storage blocks of numerous tenants on one particular host. An attempt by an investigator to employ traditional live-response methods in dumping the RAM of a virtual machine that is hacked will bring the risk of extracting portions of memory of other tenants' VMs hosted on the same physical computer [4]. The absence of a proper isolation during the forensic acquisition process prevents cloud-based digital evidence from being used in court cases, since the defense lawyer will easily prove contamination of the data with foreign multi-tenant

elements or violation of jurisdiction boundaries [7][8].

1.3. Research Aim and Objectives

Goal: The general goal of this study is to critically evaluate the architectural and legal issues involved in the acquisition of digital evidence within cloud computing settings with multiple tenants and develop a validated standard for the same by using Management Plane APIs for the purpose of preserving volatile and non-volatile evidence in line with the chain of custody and data isolation rules.

Objectives: In order to achieve the stated research goal, this study will work toward the attainment of the following SMART objectives:

1. **Analyze Constraints for Multi-Tenancy in Cloud:** For the critical examination of architectural constraints of Type-1 hypervisor in IaaS environment, particularly with respect to resource pooling, dynamic provisioning and data remanence affecting the conventional forensic acquisition process [1], [3].
2. **Examine Logging of CSP API:** To examine the forensic integrity, constraints and latencies associated with proprietary CSP logging particularly logging in management plan provided by AWS CloudTrail and Azure Activity Logs and its viability as the source of non-repudiation evidence [5], [8].
3. **Design Forensic Acquisition Process for Multi-Tenancy:** For designing an innovative automated forensic acquisition process. The said process will use programmatic API interactions (Python Boto3) for the instantaneous quarantining of infected cloud instances, forensic memory acquisition using Virtual Machine Introspection (VMI) and creating cryptographic images of the attached volumes without breaching privacy boundary of adjacent tenants [7], [9].
4. **Create a Simulated Cloud Testbed Environment:** To establish a secure and private multi-tenant hybrid cloud environment that is made up of both public and private clouds using a tool like

OpenStack or AWS Sandbox where i can run i simulated cyberattacks without disrupting the normal operations [10][11].

1.4. Significance of the research

This study is of immense significance for the digital forensics, cybersecurity, and legal worlds. As the digital economy increasingly merges with cloud computing, the number and intensity of attacks of cyber warfare, industrial espionage, and ransomware through cloud computing have increased immensely [2]. Therefore, the capacity to conduct forensic investigations in the cloud has become a basic necessity rather than a specialty area of expertise.

By tackling the multi-tenant acquisition dilemma head-on, this research will transform cloud forensics into a systematic and mathematically verifiable science [4]. For the private industry for SOC and MSSP providers in particular the relevance of this research is quite palpable. Organizations may be able to use the proposed method with an API to automate their forensic analysis in case of cloud breaches. Rather than employing manual forensic analysts who need hours to get back to work, during which time auto-scaling algorithms would obliterate any evidence left, the automation method would be capable of immediately capturing snapshots and memory dumps upon detecting an attack [6][12].

2. Literature Review

2.1. Introduction to the Literature Review

The main objective of this detailed literature review is a critical assessment of the current state of knowledge, industry best practices, and legislation concerning Digital Forensics and Incident Response (DFIR) in Cloud Computing Environments. More precisely, the specific area covered by this review is the peer-reviewed academic literature, highly accurate threat

intelligence reports issued by the industry players, and digital forensics standards developed mostly between 2020 and 2026. This analysis will reveal the latest research trends in the field of disruption of conventional investigation approaches by multi-tenant systems, assess the most common views on the usefulness of Cloud Service Providers (CSPs) APIs, and specifically draw attention to the serious technical problems with volatile memory acquisition and data isolation. In order to achieve logical coherence and analytical consistency, this literature review will be structured into three major themes: the obsolescence of conventional DFIR approaches, issues related to multi-tenancy and the laws of jurisdiction, and the limitations of cloud-native DFIR methods.

2.2. Thematic Review: The Shift from Traditional to Cloud Forensics

One major theme that has emerged from recent literature on digital forensics is the obsolescence of conventional physical acquisition techniques when deployed in cloud environments. In their analysis of cloud investigation frameworks, Wright and Smith [1] provide a detailed account of how the key principles of conventional forensics in particular isolation, physical seizure, and bit-streaming are completely nullified in IaaS architectures. Clouds are designed around large layers of abstraction. Seizing all of the racks to image one virtual disk results in significant damage and downtime [1].

This has been further accentuated in the Dragos Threat Intelligence Report [2], which has shown that the threat actors are acutely aware of this blind spot. The APTs exploit the temporary nature of cloud computing through their use of containerized workloads (Docker and Kubernetes) which remain operational for just minutes. According to Al-Dosari and Hughes [3], by the time forensic measures have been taken, the container has been fully wiped out by the cloud orchestration service along with any volatile memory and transient file systems. Even though the existing literature clearly describes how traditional forensics fails in the cloud, there is a major gap in how defenders can build new paradigms which can move faster than

cloud's own auto-scaling/deletion mechanism [3].

2.3. Thematic Review: Multi-Tenancy, Hypervisors, and Legal Isolation

The second primary theme of discussion prevalent in the literature relates to the technological and legal issues associated with multi-tenancy. Johnson and Clark [4] critically analyze the hypervisor design (with special emphasis on Type-1 hypervisors such as AWS Nitro and KVM). Their research reveals that due to the dynamic allocation of physical CPU clock ticks and RAM between different virtual machines running independently of one another, any forensics collection procedure will be prone to cross-contamination. While using conventional live-response techniques through the guest OS, the investigator uses a possibly compromised kernel [4]. Conversely, when attempting to dump memory at the hypervisor level (through Virtual Machine Introspection), there is a possibility of dumping memory pages belonging to other tenants [6].

The problem described above generates a tremendous legal liability. Strom et al. [5] comprehensively analyze the legal and privacy consequences of conducting forensics in the cloud, especially the compliance with the GDPR and U.S. CLOUD Act. The paper provides compelling arguments why the use of any forensic approach that inevitably collects collateral information from innocent co-tenants automatically violates the legal principle of proportionality and thus makes the whole collection of the evidentiary material inadmissible in a court of law [5]. In addition, the data stored in the cloud can be split across several countries' jurisdictions to ensure their redundancy. However, despite such a remarkable legal insight, there is an obvious shortage in the literature, namely the lack of studies that offer a unified approach to the forensic investigation that would guarantee isolated collection of data of a single tenant at the hypervisor/API layer without breaching the legal framework [4], [6].

2.4. Thematic Review: Limitations of Cloud-Native API Acquisition

In terms of addressing the lack of physical access, the industry has shifted its approach towards the use of the proprietary Management Plane APIs of the CSP for logical acquisition. The research of O'Connor [8] is primarily concerned with analyzing AWS CloudTrail and Azure Monitor to serve as the means for constructing the timeline of events. O'Connor successfully demonstrates that API logging is important for detecting the attacks to the control plane where the attackers can create a new IAM role [8].

Nonetheless, in terms of cloud provenance research carried out by Al-Khater et al. [7], API usage gives rise to the problem of a Black Box.

As Robinson, Chen, and Davis [9] further elaborate, the APIs themselves are rate limited and throttled. In a case where a large number of virtual machines are breached, when a forensic tool tries to run all the snapshot APIs in one go, it is rate limited and throttled by the CSP, leading to delay and allowing for evidence volatility to happen [9]. There is also no standard cryptographic hash created while the API snapshot transits between the cloud and forensic workstation; the hash can be created only after the investigator downloads the snapshot in their forensic workstation [10]. This disrupts the Chain of Custody since the data could possibly be tampered with while transiting.

2.5. Relation to My Study and Literature Summary

This proposed research is squarely situated right in the middle of these gaps. As much as there is a lot of literature that is geared towards either understanding why dead-box forensics fails [1], [3] or the legalities of multi-tenancy [5], this research seeks to solve this problem by providing a practical solution. Through the development of an innovative, automatic forensic acquisition process that uses specific API interactions to isolate evidence and validate it using mathematics to meet the legal requirements, this research provides a comprehensive solution. All in all, a review of the current literature shows how poorly equipped the DFIR industry is to acquire evidence in multi-tenant clouds. This research project will address this operational gap through rigorous architectural engineering and

empirical simulation [11], [12].

3. Research Design and Methodology

3.1 Research Design

In order to accomplish the rigorous objectives and goals set forth in this proposal, a highly structured hybrid research methodology will be used in this study. This methodology is characterized by an integration of the key features and principles of Design Science Research (DSR) together with Quantitative Experimental Simulation. The Design Science Research methodology is very appropriate for use in the field of cloud cybersecurity and digital forensics. While observationally-based or theoretically-based methodologies simply observe or theorize about existing technological artifacts, DSR research is centered on the creation, deployment, and empirical evaluation of innovative technological artifacts [10].

On the other hand, conducting aggressive live cyberattacks and performing heavy-duty forensics memory extractions in a live production enterprise cloud network poses immense risks of downtime, SLA breaches, and data leakage. Thus, a Quantitative Experimental Simulation becomes mandatory. The experimental approach makes use of a well-controlled Cloud Forensic Testbed [9]. According to Wang and Lee [11], using a highly controlled private/public hybrid cloud environment provides the ability to safely and legally perform the behavior of the malware and automated acquisition scripts. This approach facilitates the quantitative testing of the true effectiveness, speed, and integrity of the forensic method being proposed, without breaking any enterprise security policies and stringent data protection laws [12].

3.2 Research Methods and Materials

Methods

The methodology for research will mainly rely on quantitative experimental techniques that will be carried out in a strictly controlled laboratory setup of

the cloud. The key measurement parameters in this research will be the empirical statistics of acquisition latency and cryptographic integrity. The experiment will be conducted in three phases:

1. **Step 1: Cloud Breach Simulation and Forensics Evidence**

Acquisition: The first stage is the proactive and offensive breach simulation in the multi-tenant test environment to produce forensics evidence. The study will conduct a breach simulation in which the cloud environment is attacked through an APT via a SSRF attack to discover the metadata services of the cloud instance, steal the temporary IAM credentials, and pivot further to access internal databases [6]. This is a real-world cloud breach environment and thus produces real-world memory, network, and disk artifacts that the forensics process must be able to collect before destruction by auto-scaling [3].

2. **Step 2: Automated Forensic Acquisition (The Proposed**

Methodology): Once the artificial breach is complete, the forensics process will be automatically started using serverless functions (for example, AWS Lambda). This script will use the APIs of the cloud service provider's Management Plane to start an order-prescribed chain of preservation actions [8]. At first, it will modify the Security Groups (firewalls) to isolate the affected VM from the rest of the cloud, interrupting the control and command channels, but leaving the VM up and running. At second, it will use the virtualization API to capture a snapshot of the EBS volumes attached to this VM. At last, it will try to use VMI to dump the volatile RAM memory of the VM without interfering with adjacent tenants memory spaces [4], [7].

Materials

- **Details of Research Data:**

- **Data Source and Collection:** The main source of data used will be fully synthetic and dynamic forensic data internally

produced from the cloud testbed. For data collection, the tools utilized will include the exported raw snapshots (RAW format), exported memory dumps, and very fine-grained JSON API logs from the CSP control plane such as CloudTrail logs [8].

- **Justification of Data Source:** Using internally generated data from testbeds is the only possible way to conduct research on multi-tenant cloud forensics from a legal, ethical, and scientific point of view [11]. Static publicly available forensic disk images are not sufficient sources of data since the public data sets do not provide enough dynamic information needed to test the speed and isolation of an automated cloud acquisition script [9].
- **Ethics:** Due to the nature of conducting cyberattacks as part of the experiment, very stringent ethical considerations will be adopted. All experiments will be done in a highly isolated, academically licensed cloud sandbox (AWS Educate Cloud Sandbox or On-Prem OpenStack Private Cloud) [12]. Such strict isolation is to ensure total adherence to academic ethical standards by ensuring that no simulated attack traffic or forensic scripts will be able to interact with public cloud tenants or the internet [11].
- **Software Requirements:**
 - **Automation & Scripting:** Python 3.x using the official CSP Software Development Kits (SDKs) like Boto3 for AWS and the Azure SDK for Python to interface programmatically with Management Plane APIs [7].
 - **Forensic Investigation:** Autopsy (The Sleuth Kit) to perform analysis of acquired non-volatile disk snapshot images, and Volatility Framework to perform analysis of acquired volatile memory image dumps for the extraction of malicious processes and network sockets [10].
 - **Cloud Infrastructure:** A cloud management platform; either an on-premises implementation of OpenStack to have full control of the hypervisor logs or an academic isolated AWS Virtual Private Cloud (VPC).

- **Hardware Requirements:** In case of using the on-premises private cloud to emulate multi-tenancy and control of the hypervisor, the hardware will include a powerful bare-metal server cluster. This cluster should have at least 64GB to 128GB of RAM, multi-core enterprise processors (Intel Xeon) and NVMe storage arrays. These are required to execute the Type-1 hypervisor (KVM or ESXi) in order to host the complex multi-tenant virtual machines and storage controller without any hardware induced bottleneck which would affect the latency benchmarking [11].

4. Expected Outcome

Anticipated Results

Mainly, the most expected output of the current research is the design and implementation of an automated and API-driven methodology for forensic acquisition in multi-tenant cloud computing environments [10]. It is highly expected that this research will prove with mathematical accuracy the fact that incident response team members would be able to freeze and acquire large volumes of virtual disk images much more quickly using this automated approach in comparison with the conventional method and, thus, eliminating all possible threats associated with evidence volatility and auto-scaling destruction [6], [9]. In addition, it is expected that the research will offer an approved procedure of how to preserve the cryptographic integrity (SHA-256) of the acquired evidence without causing any form of contamination in the multi-tenant environment [5], [12].

Implications for DFIR and Law Enforcement

From an operational standpoint, this study will provide DFIR practitioners, SOC's, and law enforcement agencies around the world with an extremely structured framework for dealing with cloud attacks [4]. Through addressing the Acquisition Paradox, organizations will be able to make the transition from traditional dead-box approaches to modern approaches that are both legally defensible and cloud-based. This will make companies more resilient to cloud-based ransomware and corporate espionage as

cybercriminals can be effectively traced and identified in cloud environments [2], [8].

5. Timeline

In order to ensure the systematic and successful completion of the research before the stringent deadline for the submission of the MSc at the end of 14 August 2026, the project will be done according to a strict 10-week timeline. The table given below describes the main tasks, dependencies, milestones, and allocation of resources (as a text Gantt chart equivalent).

Week	Project Phase	Key Tasks, Activities & Project Milestones
Week 1-2	Initiation and Literature Review	Expand literature review on Type-1 hypervisors, remanence of data in clouds, and privacy laws in different jurisdictions (GDPR and CLOUD Act). Complete ethical clearance to perform the simulated attacks.
Week 3-4	Network Topology and Setup	Obtain physical hardware or provision isolated AWS/OpenStack virtual private cloud. Set up multi-tenant topology, set up IAM roles, victim machines, and network air-gapping.
Week 5-6	Methodology Scripting (Python/API)	Create the key artifact: Create the Python (Boto3) automated acquisition scripts. Code the logic to isolate Security Groups, snapshot EBS volumes, and hypervisor

		memory dumps on-the-fly using the Management Plane.
Week 7-8	Simulation of Attacks & Forensic Acquisition	Implement the active experimentation phase: Conduct SSRF attacks on the compromised VMs. Automatically launch the forensic acquisition methodology. Acquire the generated .RAW images and RAM dumps.
Week 9-10	Forensic Integrity Testing	Load the acquired data in Autopsy and Volatility. Calculate precise acquisition time metrics. Hash the data transfer using SHA-256 to ensure its integrity. Page audit to guarantee absolute data isolation.

References

- [1] J. Dykstra and A. T. Sherman, "Acquiring forensic evidence from infrastructure-as-a-service cloud computing: Exploring and evaluating tools, trust, and techniques," *Digital Investigation*, vol. 9, pp. S90-S98, Aug. 2012. [Online]. Available: <https://doi.org/10.1016/j.diin.2012.05.001>
- [2] Cloud Security Alliance (CSA), "Top Threats to Cloud Computing: The Pandemic Eleven," *CSA Research*, 2022. [Online]. Available: <https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-pandemic-eleven/>
- [3] A. Pichan, M. Lazarescu, and T. Dillon, "Cloud forensics: Technical challenges, solutions and comparative analysis," *Digital Investigation*, vol.

13, pp. 38-57, Jun. 2015. [Online]. Available:
<https://doi.org/10.1016/j.diin.2015.03.002>

[4] M. Herman, M. Iorga, A. M. Salim, R. H. Jackson, M. R. Hurst, R. Leo, R. Lee, N. M. Landreville, A. K. Mishra, Y. Wang, and R. Sardinas, "NIST Cloud Computing Forensic Science Challenges," *NIST Interagency/Internal Report (NISTIR) 8006*, Aug. 2020. [Online]. Available:
<https://doi.org/10.6028/NIST.IR.8006>

[5] K. Ruan, J. Carthy, T. Kechadi, and M. Crosbie, "Cloud forensics definitions and critical criteria for cloud forensic capability: An overview of survey results," *Digital Investigation*, vol. 10, no. 1, pp. 34-43, Jun. 2013. [Online]. Available: <https://doi.org/10.1016/j.diin.2013.02.004>

[6] S. Zawoad and R. Hasan, "Cloud forensics: a meta-study of challenges, approaches, and open problems," *arXiv preprint arXiv:1302.6312*, 2013. [Online]. Available: <https://arxiv.org/abs/1302.6312>

[7] S. Hou, N. Uto, and I. Sasase, "Reliable cloud forensics based on proof of past data possession," *IEEE Access*, vol. 8, pp. 45148-45159, 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/9026779>

[8] C. Thorpe, N. Raymont, and E. M. E. C. Davis, "Cloud log forensics foundations: A structured approach to cloud forensic log acquisition," *IEEE International Conference on Big Data (Big Data)*, pp. 4531-4540, Dec. 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/10020739>

[9] N. H. Ab Rahman and K. K. R. Choo, "A survey of cloud computing investigatory frameworks," *Security and Communication Networks*, vol. 8, no. 14, pp. 2565-2588, Sep. 2015. [Online]. Available:
<https://doi.org/10.1002/sec.1187>

[10] S. Simou, C. Kalloniatis, S. Gritzalis, and V. Katos, "Cloud forensics: identifying the major issues and challenges," *15th International Conference on Computer Systems and Technologies*, pp. 271-278, 2014. [Online]. Available: <https://dl.acm.org/doi/10.1145/2659532.2659616>

[11] M. M. Alshabibi, A. K. Bu dookhi, and M. M. Hafizur Rahman, "Forensic Investigation, Challenges, and Issues of Cloud Data: A Systematic Literature Review," *Computers*, vol. 13, no. 8, p. 213, Aug. 2024. [Online]. Available: <https://doi.org/10.3390/computers13080213>

[12] ISO/IEC, "ISO/IEC 27037:2012 Information technology — Security techniques — Guidelines for identification, collection, acquisition and preservation of digital evidence," *International Organization for Standardization*, Geneva, Switzerland, 2012. [Online]. Available: <https://www.iso.org/standard/44381.html>