AN EXPERIMENTAL SURVEY TOWARDS ENGAGING TRUSTABLE HYPERVISOR LOG EVIDENCE WITHIN A CLOUD FORENSIC ENVIRONMENT

Sean Thorpe

Faculty of Engineering and Computing, University of Technology, Kingston, Jamaica
thorpe.sean@gmail.com

ABSTRACT

In this survey paper the author explores the technical as well as high level conceptual trust issues that arise in acquiring log forensic evidence from the virtual machine (VM) hosted operating systems within the data clouds. This specific survey work is done at the University of Technology [UTECH], Jamaica, which currently functions as its own independent private data cloud provider. The data acquisition is particular to the hypervisor system logs that can be used to track VM incidences which are later used to compile potential evidence for a cloud investigation. This work also presents a model to show the layers of virtualization trust that can arguably be used to support the collection of such log evidence. The paper provides the context for the support of such cloud digital investigations and analyzes the choices available to a forensic investigator using proof of concept experiments. The experimental work is achieved by making a comparative evaluation of popular forensic acquisition tools including Guidance EnCase and AccessData Forensic Toolkit, as to how volatile and non-volatile hypervisor log data can be collected. Finally the paper explores three solutions for the managed log evidence data acquisition phase within a cloud investigation.

KEYWORDS

Forensic, Log, Cloud, Trust, Hypervisor

1. INTRODUCTION

Discovery and acquisition of log evidence in remote, elastic, provider-controlled cloud computing platforms differ from that in traditional digital forensics, and examiners lack appropriate tools for these tasks. This work is motivated by prior work [29,30,32,34,38,39,41] and presents new foundations for supporting the compilation phase within a remote log based cloud investigation which one could argue is trustworthy and forensically sound. The view of using the logs as a part of a post incident analysis framework has been evaluated within cyber cloud investigations[37]. However supporting forensic log analysis with independent tools where the data is collected remotely is not prevalent. Hence this paper explores the issues by examining the existing forensic tool alternatives. This work also presents a platform for which law enforcement in general can evaluate confidence in compiling log evidence from data clouds**.**

While there are many important issues in this emergent field, one needs to begin with the data compilation /acquisition phase. Crimes that target or use cloud computing will present themselves in this phase, and examiners will rely on their existing expertise in commercial tools like Safeback, Guidance EnCase or Access Data Forensic Toolkit (FTK), unless alternative tools and techniques are provided. Although the author has designed a new tool called virtual machine

DOI : 10.5121/ijcsit.2012.4610

log auditor[33] used in experiments for compiling hypervisor log evidence and making analysis with this evidence [28,29,42], its verifiability as a trustable tool outside of the University environment is still in progress, and hence the existing expert forensic tools represent the better approach for comparative analysis.

Digital forensics for the data clouds presents new technical and legal challenges for the law enforcement community. Cloud forensics is different from traditional forensics, particularly given the distributed jurisdictional nature of the evidence-: in terms of the lack of physical access, and the demand for trust with respect to the integrity and authenticity of the data collected within these ubiquitous logical domains. While the goals of the potential cloud digital investigator are formulated on the skills acquired from traditional digital investigations, the unconventional difficulty of acquiring public cloud data is riddled with forensic integrity at hypervisor, detected at VM host operating source, including the need to handle the disparate nature of the system log data types and volumes that may be available from the multiple vendor platforms. This in essence makes the chain of custody and data ownership requirement a perplexing problem for users, in general and specifically for investigators and system administrators.

Seizure and acquisition of digital artifacts are the initial steps in the forensic process [2]. Two possible scenarios exist: remote investigators could collect forensic log evidence themselves from the source, or providers could deliver it. Each scenario requires a different degree of trust in the log data returned. Notwithstanding it can be maintained that both approaches should only be supported when the forensic agencies and agents are all a part of a known and trusted collaboration [43], where verification history of the trusted parties can be maintained. Further, each scenario uses different technical implementations to recover the log data. Given years of development, acceptance by the judicial system, and expertise in the field, market leaders in the commercial forensic tool space including EnCase and FTK are ideally prepositioned for the cloud forensic challenge [24]. One question that remained until now, however, was an evaluation of the ability of such tools to acquire and analyze cloud-based evidence in particular log data.

Cloud computing is based on the well defined service oriented architectural model of sharing IT resources in an on demand and elastic fashion through virtualization. Compared to its predecessors of grid and main frame computing the outline of the dynamic service provisions make the virtualized data clouds distinctively appealing for economies of scale as opposed to its predecessors [49].
Though some definitions of cloud computing include popular web-based services such as email and social networking, one limits the scope of this paper to computing resources that are treated as billable utilities. More specifically, one has to use the Platform as a Service (PaaS) model [17]. In this model, the consumer has complete control over a guest operating system running in a virtual machine (VM). The provider retains control and responsibility for the hypervisor (HV) down to the physical hardware layer in the data centre. Since IAAS is built on the robust and well established TCP/IP, it is more important to focus on the PAAS as the basis on which to build future work for the forensic investigator.

In this paper, one assumes that the target system of the forensic investigation only exists in a privately managed University cloud. The elastic nature of cloud computing makes it possible for the criminal to commit a crime and then immediately destroy the evidence, but that situation is not considered here. While some cases will involve the cloud as the instrument of the crime, others will involve the cloud-hosted service as the target of the crime, as in the case of the UTECH cloud data centre.

[2] states that “the issues of incident response and computer forensics for the cloud require fundamentally different tools, techniques, and training.” In this paper, one evaluates the validity

of this statement within the context of data acquisition from the hypervisor system logs within the UTECH environment. The contributions of this work include:

- Results from three UTECH experiments that exercise existing tools for persistent and non-persistent log data collection in a private cloud.
- Analysis of alternatives for log forensic acquisition at the PAAS layer of the virtualized stack, for cases when there is insufficient trust in data acquisition of the log data.
- A demonstration of how virtual machine introspection can be used to inject a remote log forensic agent for remote acquisition within the UTECH hosted cloud data centre.
- Exploration of three strategies for log forensic data acquisition with non trusted hypervisors

The rest of the paper is organized as follows. Section 2 reviews previous and related work. Section 3 presents the cloud forensic examination, layers of trust and choices in cloud forensics. Section 4 presents the UTECH data centre experiments in using the native capabilities of EnCase, FTK, Fastdump and Memoryze for data acquisition in EC2. Section 5 suggests alternative approaches. Section 6 discusses considerations and Section 7 concludes the work.

2. PREVIOUS AND RELATED WORK

The US federal government evaluates some of the most widely used forensic tools to ensure reliability based on the well established Daubert standards. The National Institute of Standards and Technology's (NIST) Computer Forensic Tool Testing (CFTT) project is charged with testing digital forensic tools, measuring their effectiveness, and certifying them[17,18]. They evaluated EnCase 6.5 in September 2009 and FTK Imager 2.5.3.14 in June 2008 [19, 20]. More recently several researchers have been testing and certifying the enterprise versions of these products that include remote forensic capabilities. The author has not seen this visibly in the University environments, especially where those environments are functioning as cloud providers.

NIST currently publishes a Digital Data Acquisition Tool Specification, which "defines requirements for digital media acquisition tools in computer forensic investigations" [19]. The most recent version of the specification was written in 2004, before cloud computing(as currently defined) was known to exist then.

Several researchers have pointed out that evidence acquisition is a forefront issue with cloud forensics [5, 6, 22, 26], especially in the context of log data compilation and evaluation [30, 34, 35, 41]. In Dykstra and Sherman's analysis [5, 6] they conducted two hypothetical case studies that illustrated the non-trivial issues with collecting evidence from a cloud crime.

Ruan et al [22] suggested that evidence collection should obey "clearly-defined segregation of duties between client and provider," though it was unclear who should collect volatile and non-volatile cloud data and how. Taylor et al.[26] also lamented about the lack of appropriate tools for data from the cloud, noting that "Many of these tools are standardized for today's computing environment, such as EnCase or the Forensics Tool Kit [sic]." Thorpe et.al [29,30,34,35,39,41] made conscious efforts to design a new tool for log audit analysis that addressed timestamp and event data collection from the hypervisor system log environments. This proved useful in corroborating case evidence for the UTECH cloud environments. Although the design approach for building one's own forensic tools can prove in expensive cost wise, it still lacks the trust and standard as compared to known expert tools.

Lacking from the current cloud log forensics community is the ability to support virtual machine introspection (VMI). VMI is a technique whereby an observer can interact with a virtual machine client from the outside through the hypervisor. Garfinkel et.al [9] demonstrated a technique for intrusion detection inside a virtual guest using VMI. Garfinkel et.al [9] in 2009 used Symantec and demonstrated anti-virus code injection into a virtual machine from the hypervisor [3]. From that year, researchers have proposed various applications of VMI to forensic memory analysis [4,16]. Santana [45] reports that Terremark uses introspection for monitoring, management and security for their vSphere cloud computing offering. So far no attempt has been made to inject a forensic tool, such as an EnCase servlet, into a virtual machine from the hypervisor.

In 2009, Gartner [12] published an overview of remote forensic tools and guidance for their use, particularly targeted at enterprise environments. They cited EnCase and FTK as the most widely used products, with the greatest international support. These tools, however, have their faults: in 2007, vulnerability was found in the authentication between the remote EnCase agent and the server [10]. From a legal perspective of trust, Guidance Software's own "EnCase Legal Journal" for 2011, a comprehensive examination of legal issues and decisions about electronic discovery, has no mention of judicial decisions or statutory law related to the complex legal questions surrounding remote data acquisition [11].

EnCase Enterprise and FTK include a client-server feature for remote forensics. In each case, a small executable is installed on the client machine (EnCase calls the executable a "servlet;" FTK calls it an "agent"). Fig. 1 illustrates how the server, built into or on top of the vendors' forensic analysis software, communicates with the client over a secure connection, and can command the client to return forensic data including a hard drive image , ideally as a snapshot hypervisor system log capture.

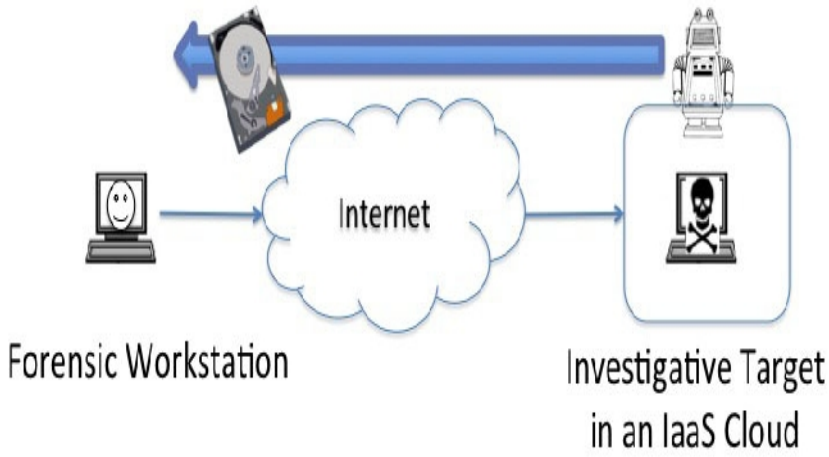

Figure 1. General technique for remotely acquiring log forensic evidence over the Internet using a remote trusted agent as a snapshot.

The examiner may conduct some log data acquisition procedure remotely on the client, or return to the target evidence server for local log analysis. Remote log forensics is employed in large enterprises where both virtual and physical machines may be geographically dispersed, but the incident response team remains centralized to the log supported investigation.

3. THE CLOUD FORENSIC EXAMINATION

In this section one explores the log forensic examination of a cloud-based crime. As a foundation, there is the presentation of a model to reason about the trustworthiness of log evidence from the cloud, since the level of trust influences the choices for how an exam should be conducted. Second, one poses choices that determine how to approach a cloud log forensic investigation.

3.1. Layers of Trust

Before evaluating tools for the log acquisition, it is important to understand trust in the log cloud environment. When brought to court, the judge or jury must ultimately decide if they believe and trust the log evidence presented to them. This choice embodies a specific confidence about whether the potential log result is accurate and reliable. In traditional forensics, where the target machine is physically present, some of the same trust issues exist, with log cloud forensics. Consider an example where a single desktop computer has been used to plan the theft of corporate information within the UTECH environment. If law enforcement removes the hard drive for imaging, they must trust their hard drive hardware to read the disk correctly. If they run forensic tools on the live computer, they may have to trust the integrity of the host operating system in addition to the hardware. If the suspect computer was hosted in the cloud, new layers of trust are inherently required. One cannot consider direct trust in the automated or human agents with responsibility for these forensic acquisition tools, since these components, while important, are outside the perimeters of cloud environment. Table 1 adopts a model of trust in PaaS for the log based cloud computing into six layers.

In PaaS, the consumer retains administrative control over Layers 5 (Guest OS) and 6 (Guest application), despite no physical access. Furthermore, the forensic acquisition activity would be different at each layer. Each layer requires a different amount of confidence that the layer is secure and trustworthy; and hence the farther down the stack, the less cumulative trust is required. In public clouds, all layers require some trust in the provider, especially trust against malicious insiders. Ultimately, it is the judge or jury that must have confidence in the data to render a legal decision.

Layer	Cloud Layer	Acquisition Method	Trust Required
6	Guest App	Depends on data	Guest OS Hypervisor Hardware, Network Hypervisor
5	Guest OS	Remote Forensic Software	Guest OS Hypervisor Hardware Network Hypervisor
4	Virtualization	Introspection	Host OS Hypervisor Hardware Network
3	Host OS	Access Virtual disk	Host OS Hardware Network
2	Physical Network	Packet capture	Network
1	Hardware	Access Physical disk	Hardware Network

Table 1. Six layers of the Paas log cloud environment and potential log forensic acquisition techniques for each, including the cumulative trust required by each layer.

Imagine a situation where the UTECH VM forensic investigator has remote access to the guest virtual machine operating system. The investigator could collect hypervisor log evidence contained inside the VM via the Storage Area Network (SAN) disk, install a forensic tool and obtain log evidence remotely, then terminate the VM session and analyze the logs offline.

Unfortunately, log acquisition at this layer requires trust of the guest operating system, hypervisor, host operating system, underlying hardware, and the network which is not often provided. The general assumptions of trust would mean that the logs produced would be complete and accurate evidence sources of data, in that they are free from intentional and accidental tampering, which compromise an investigation as suggested in prior works[29,30]. Note that with Type 1 hypervisors, such as Xen, HyperV, HyperSafe, the hypervisor is the lowest software layer, thus arguably eliminating one layer of trust, as one values the experience and track record of the VM vendors.

As a risk mitigation strategy, the VM investigator should examine log evidence at multiple layers. This technique allows an investigator to check for inconsistency and to correlate evidence. Arranging individual log contexts together into groups is similar to the archaeological approach adopted and proposed by Garfinkel[8] in examining multiple drives to correlate evidence across seemingly unrelated evidence categories, such as for identifying members of social networks etc. By extending this idea to detect VM suspicious activity corroborated by the hypervisor logs, it can be useful towards providing digital investigative log hypotheses as supported by Thorpe and Ray in their prior works[29, 30].

Investigators may be tempted to conduct their VM log investigation remotely on a running machine particularly given the size of the log data, the time and cost to retrieve a full drive image, and the propensity to conduct live log forensics. This matter is revisited in Section 4.

Currently, law enforcement asks the cloud provider for the log data. A search warrant or subpoena is issued to the provider, and the provider executes the search, collects the log data, and returns it to law enforcement. Though this process frees law enforcement from needing remote log acquisition technology and from the burden of understanding details of the cloud environment, it presents significant trust issues particularly from the perspective of malicious insiders within the cloud data centre; as well it raises questions about the data acquisition techniques used in processing the compiled log data. The implications of this concern is significant and should not be limited to the fact that only log data is considered in this study, but alternative data sets like the existing data volumes on the VM SAN disk are also a concern. Purportedly, the examiner must now trust the integrity of the cloud provider's technical staff to execute the search in a trustworthy manner, including the hardware and software used to collect the hypervisor log data, and the cloud infrastructure (at least network and hardware) to retrieve, reassemble, and report the data. Separately the issues of data collection in general are also important, as some providers employ sandboxing , manual instance relocation , failovers, server farming, address relocation, man in the middle , and let's hope for the best techniques[46] as suitable data retrieval strategies, however due to space limitations a critique of these approaches will be provided in a follow up paper.

3.2. Choices in Cloud Forensics

One now considers how to conduct a log forensic exam of PaaS cloud computing by considering the following issues.

The layers explored in Section III are also choices of where to conduct a log forensic investigation. In particular, the investigator can choose at what layer of the cloud the log forensic process will be executed. Considerations for this decision revolve first around the technical capability to conduct forensics at that level, and second the trust in the log data returned. The layer also influences what type of log forensic data are available for collection, such as packet captures at Layer 1 (Network), physical files at Layer 2 (Physical Hardware), or virtual files at Layer 3 (Host OS). The good thing is that all layers can be captured within the hypervisor system logs as argued by Ryan ko et.al [47] who propose a file-centric logger [flogger] that can compile PAAS system data as an abstraction of the cloud's system and workflow assurance layers. The work defined in[33] could however adopt the workflow assurance layers in [47] as a synergy for creating a trustable cloud log forensic framework. Ideally for each data type the data must adhere to a strict chain of custody and must include a mechanism for integrity checking. Thorpe et.al [44] propose a compressed log hash provenance checksum mechanism that uses sequential identifier labels to the log tables compiled on the target evidence servers as one such integrity-based approach.

One must choose who will conduct the log forensic exam and where will it be conducted. Possible choices for who will execute the investigation includes law enforcement, ideally an independent examiner (e.g. the local Cybercrime Units controlled by the Police) assigned to the cloud provider. Choices for where the investigation will take place include the provider's corporate headquarters, at one of the provider's remote data centres, or at a remote law enforcement facility, or even an independent third party facility local to a jurisdiction for which such an examiner has the necessary access control rights to that data centre. These choices are as much about practicality and logistics as the law and the examiner's qualifications. Requiring a non-employee of the provider to conduct an investigation on provider premises would impose an unacceptable logistical burden to some providers, unless there are government regulations that guide the search and seizure.

As discussed in Section VI, verifying the integrity and completeness of the log data is still a challenge. Cost is another choice affecting how a forensic log investigation is conducted. When log forensic data are requested, the cost in dollars and labour to preserve and produce records might be passed on to the requestor, or sold as a service by the cloud vendor.

Technical choices of how to conduct a log forensic exam for cloud computing environments are numerous but closely resemble the options available within a traditional digital investigation. The nature of the crime dictates whether the specific log forensic process will take place as well as how (i.e. either performing live or offline machine log analysis). Second, regardless of whether the log forensic data come from a physical workstation or the cloud, the forensic goal of determining what happened is the same, except that the volume and format of the log data may differ. The examiner's choice of analysis tools may be influenced by the format of the log data collected (e.g., traditional files vs. cloud "blobs"), volume of data, and data type (e.g., netflow logs, billing records, drive images).

Cloud computing provides virtual machine snapshots as the acceptable option for the data capture. With many cloud implementations that utilize virtualization it is possible to take a snapshot of a running machine and later restore and run the snapshot offline as if it were live. This offers the ability to create a historical record, as well as do "live" forensics after the fact.

4. CLOUD FORENSICS USING TODAY'S TOOLS

In this section, one measures and evaluate the ability of EnCase Enterprise and AccessData FTK to remotely acquire forensic evidence from cloud computing and measure their effectiveness. Both products are widely deployed today, benefit from tool expertise in the field, are trusted by the courts, and have a remote acquisition feature that has been targeted at geographically dispersed corporate LANs.

The goal is to evaluate the ability and scientific accuracy of these features to acquire log forensic data from cloud computing environments over the Internet. One also test live forensic acquisition tools using Fastdump from HBGary, Memoryze from Mandiant, and FTK Imager from Access-Data within the University environment. These UTECH experiments evaluate the success at gathering log evidence, the time to do so, and the trust required.

4.1. Motivation

Experimentation and testing of today's most popular forensic tools have previously been applied to cloud computing, however not specifically to support log based cloud forensics. One proposes three experiments using the PaaS cloud model, since that gives the examiner the most access and control of all cloud models. In particular, one can use a public cloud, EC2 from Amazon Web Services (AWS), as a live test bed. Experiment 1 collects log forensic data from Layer 5, the guest operating system. Experiment 2 collects log data from Layer 4, the virtualization layer. Experiment 3 collects log data from Layer 3, the host operating system. Because Experiment 2 and Experiment 3 use the Amazon cloud, one assumes that the provider is producing the correct, untampered log data. The goal of these experiments is to evaluate the ability of five tools to acquire log forensic data from cloud computing environments over the Internet, specific to a privately managed cloud like the author's UTECH environment. The investigator would likely pick the most popular volatile and non-volatile forensic software acquisition tools and seek to use them in the UTECH cloud data centre environment.

The first tools that were chosen were Guidance EnCase and AccessData FTK, for reasons previously mentioned. They have been used in thousands of trials, and withstood arguments about

their effectiveness [11]. Each product has a remote acquisition feature that has been targeted at geographically dispersed corporate LANs. Additionally a choice of three memory acquisition tools-Fastdump, Memoryze, and FTK Imager- were also selected to determine their success in the cloud.

4.2. Extracting Data from Amazon EC2

Extracting data from Amazon's EC2 implementation requires extra work, considering that the UTECH cloud servers are distributed across four (4) campuses in distinct geographic proximities. Here one explains what is learned and ultimately used to acquire log forensic data. One choice for acquiring remote, persistent storage is to download a copy of the volume, or a snapshot thereof. Amazon stores virtual hard drives, called Elastic Block Storage (EBS) volumes, in its Simple Storage Service (S3), but they are not exposed to the end user for downloading.

Two options exist to obtain the log data for an entire volume. The first is to create a snapshot from a drive being investigated, create a volume from that snapshot, attach the new volume read-only to a trusted Solaris UNIX 2.9 instance in EC2, and then create an ISO disk image of the volume that could be downloaded. The second is to detach the target volume from the host under investigation, attach it to a trusted Solaris instance in EC2, and use a low-level copying utility (e.g., the Unix data duplication tool dd) to create a block copy which can be stored in S3 and downloaded.

Amazon provides a service to export data from S3 onto a physical device and ship it to the requestor, but the customer must provide the storage device and is billed at standard costs of $80 per storage device handled plus $2.49 per data-loading hour [1].

In neither of these cases is it possible to verify the integrity of the forensic disk image. Amazon does not provide checksums of volumes in their data cloud, so one cannot positively assert that the image retrieved is identical to the original. Further, no hardware write blocker can be used to protect the integrity of the exhibit. However, it is possible to guarantee that the log data have not been modified in transit (e.g., hashing the image before export and again after it has arrived from shipping). Notwithstanding the considerable trust issues in each of the above cases as direct trust resides in the hands of the vendor to follow the compliant NIST procedures [21].

4.3. Methods

For each experiment, the author used a non-cloud based standalone control machine to evaluate the success of the test. The control was a HP workstation with 32-bit Windows 2008 R2, a single 18GB disk drive and 2GB RAM. One connected the machine to the Internet and installed the TomCat web server. This was supported by the creation of a simple Java web enabled set of pages with identifying names and content.

Some files were deleted. This artificially compromised the machine using a web-based vulnerability, and assumed that a criminal and forensic investigation had commenced. An image of the drive with EnCase and FTK was done.

Experiment 1 tested the advertised ability of popular tools to collect log forensic data remotely in the cloud at the guest OS (Layer 5). Success or failure would be measured by (a) if the tool was able to collect log evidence remotely, and (b) how accurately the data compared to those from a standalone control machine. One prepared a single, Internet-connected (proxied), forensic examiner workstation with 64-bit Windows 7 Enterprise. EnCase Enterprise 6.11, including the SAFE (Secure Authentication for EnCase), was installed according to the manufacturer's

instructions. FTK 3.2 was also installed. In Amazon EC2, one provisioned a new virtual machine to simulate the target of an investigation. This machine was an Amazon-provided Windows 2008 R2 32-bit image with a single 18GB disk drive and 1.7GB RAM. This was configured against an Amazon firewall to allow only a Remote Desktop Protocol (RDP) (tcp/3389).

Once connected to the UTECH test target machine using RDP and proceeded to exercise normal behaviour of a user configuring a web server. Then a copy of TomCat was downloaded and installed and created several web pages with identifying names and content. Some files were deleted. Then again with this artificial compromise, the machine using a web-based vulnerability assumed that a criminal and forensic investigation had commenced. EnCase Servlets and FTK Agents are the remote client programs that communicate with their host server controllers.

Each deployment can be done differently. In a corporate environment like the UTECH, agents are typically deployed to Windows machines over the network using Windows file shares. The products also allow manual file delivery (e.g., USB). In the experiment, one transferred the agent to the UTECH target virtual machine over RDP and executed it. The firewall was modified to allow communication with the agent: the EnCase servlet used tcp/4445 and the FTK agent used a user-defined port of tcp/3399. A test of FTK Imager Lite version 2.9.0. was also done. The product was copied over the Remote Desktop connection from the UTECH examiner's workstation and run interactively. FTK Imager Lite does not require installation, and runs self sufficiently once uncompressed. For this experiment one attached a second storage volume onto which one saved an image of the primary log volumes captured by FTK Imager. Finally, one ran Fastdump, Memoryze and FTK Imager to acquire images of system memory, resulting in three 1.7GB images.

Experiment 2 tested popular forensic tools at the virtualization layer by injecting an agent into the virtual machine (Layer 4) within the UTECH testbed. Success or failure was again measured by (a) the ability of the tool to collect log evidence, and (b) how accurate the log data were compared to those from a standalone control machine. Then one prepared an installation of the Eucalyptus cloud platform [1] from the Solaris 2.9 distribution on a Dell workstation. Eucalyptus supports the Xen hypervisor for managing virtual machines, and LibVMI[1] which is a library for monitoring guest operating systems in Xen. Then the LibVMI library writes to memory of the guest virtual machine. With this capability, one demonstrated injecting an EnCase Servlet and FTK Agent directly into a running guest. As with Experiment 1, the agent communicates over the closed UTECH local area network, used to facilitate this test.

Experiment 3 tested forensic acquisition at the host operating system level by exercising Amazon's Export feature (Layer 3). This experiment most closely resembles the process probably used to satisfy subpoenas and search warrants in a traditional crime scene forensics investigation or digital data centre investigation, since the data are exported from Amazon's internal network at a data centre. Additionally, AWS maintains a chain of custody for the VM storage device while it is in its custody. Then a measure of the success or failure by (a) the ability of the technique to collect hypervisor log evidence, and (b) the accuracy of the log data as compared to those from the standalone control machine. AWS Export involves a service request to Amazon and shipping them a storage device. Unfortunately, it is currently possible only to export data from an S3 bucket and not from an EBS volume. To meet that requirement, one attached the EBS volume from the compromised machine to a Solaris Unix 2.9 VM station, and used dd to store an image of the volume in an S3 bucket. Then a request from AWS for an export of this buckets, followed by the shipping of Seagate FreeAgent eSATA external hard drive. The storage device was returned with a copy of the log data.

4.4. Results

The manual installation of the EnCase Servlet and FTK Agent in Experiment 1 was successful and hence able to acquire a hard drive and memory image remotely on the closed UTECH network. Analyzing these images in the EnCase Forensic and FTK Investigator respectively correctly revealed a timeline of activity, including the installation of TomCat and the web pages that were created and deleted. The analysis revealed no unusual artefacts of the virtual environment, nor any apparent anomalies to raise doubt about the integrity of the hypervisor system log data. The speed of the acquisition process was limited by the need for learning how to use the remote agents and the UTECH network bandwidth to transfer the log data. The latter took approximately 8 hours each for EnCase and FTK to transfer the 18GB disk image and 1.7GB memory image using the university's OC-12 connection.

Experiment 2 successfully resulted in a complete image of the drive and a correct timeline. VM introspection is a powerful tool for forensics and allows live investigation of a host without revealing the presence of the investigator.

However, introspection is a special feature which must be implemented by the cloud service provider. This was the only experiment that could be verified cryptographically with the integrity of the image, since one had access to the physical disk and could compare hash values of the EnCase image and the original disk.

Table 2. Results of three experiments acquiring cloud-based log forensic evidence using popular tools, including the time to retrieve the data and trust required in the guest operating system (OS), hypervisor (HV), host operating system, host hardware, network, and Amazon Web Services (AWS) components.

Experiment	Tool	Evidence Collected	Time(in hrs)
1	Encase	Success	8
1	FTK	Success	8
1	FTK Imager (Disk)	Success	8
1	Fast Dump	Success	2
1	Memoryze	Success	2
1	FTK Imager(Memory)	Success	2
1	Volume Block Copy	Success	10
2	Agent Injection	Success	1
3	AWS Export	Success	60

The AWS Export process in Experiment 3 also successfully returned a complete image of the drive. Then a loading of this drive into EnCase and FTK with no difficulties was achieved, and hence verified the contents of the drive. An added benefit of this method is that AWS generates a log report with metadata for each file exported [42]. This report contained the following for each file: date and time of the transfer, location on the storage device, MD5 checksum, and number of bytes. These data are saved in an S3 bucket that is specified in the export request. Using expedited shipping, it took four days to receive the log data, at a cost of $125. One imagines that this process would closely mimic the steps taken by AWS when complying with a search warrant or subpoena. EnCase and FTK were easiest to use. Despite setup and learning time required to use the remote capabilities, the features of the tools were familiar and easy to execute.

The eight (8) hour time required to retrieve a disk image was significantly shorter than the 60 hours required for the AWS Export process for this data volume. AWS Export spent 4 hours loading log data, while the remaining 56 hours were spent in transit. At these rates, the most time effective choice is the export process when more than 120 GB of data will be retrieved. Table 2 summarizes the results of data acquisition in EC2. Each tool and technique successfully resulted in log evidence production, but each requires substantial trust in the cloud infrastructure at all levels.

5. ALTERNATIVES FOR FORENSICS ACQUISITION

In this section the author adopts four well known alternate solutions to acquiring cloud-based log data: Trusted Platform Modules (TPMs), the cloud management plane (management console), forensics as- a-service, and contractual support. The adoption of one or more of these alternatives would make remote hypervisor log data acquisition more trustworthy than acquisition using EnCase or FTK since trust is rooted at lower cloud layers, although a combination of all approaches may equally become warranted especially in academic environments like that of the author.

5.1. Rooting Trust with TPMs

The deployment of TPMs would root trust in cloud computing hardware. Several researchers have previously suggested this [14]. Ideally TPMs provide multiple capabilities, and no less ideal for a hypervisor based log cloud where: machine authentication, hardware encryption, key signing, secure key storage, and attestation are significant as evidence. Similar prior work on cloud attestation, and key signing was done by Popa et.al [48] using the CloudProof tool.

Previous solutions for TPMs in cloud computing focus on provisioning trusted guest VMs rather than on attestation of the host platform. If TPMs were installed in each cloud server, the hardware and associated software could validate what software is installed on each machine and verify the health and status of each machine. Despite this benefit and low cost, TPMs have limitations of their own and are not perfectly secure. It is still possible, for example, to modify a running process without detection by the TPM.

While appropriate for future consideration, one believes that the primary hindrance to this approach today is that cloud vendors have large amounts of heterogeneous, commercial hardware which is replaced as needed rather than all at once, much of which does not have a TPM. While future hardware may include a TPM, the provider cannot guarantee that each server in its cloud has one today. Nevertheless, customer demand today or in the future may drive providers to introduce trusted hardware for some or all customers including law enforcement client agencies compiling trustable log forensic data evidence.

5.2. Collection from Management Plane

Cloud computing has a powerful attribute that could be used to support trustworthy hypervisor based log forensics: consumers manage and control virtual assets via a management plane, an out-of-band channel that interfaces with the cloud infrastructure. In Amazon Web Services, this system is called the AWS Management Console. But in the case of the UTECH environment where AWS is not currently supported, ongoing work to simulate a similar system forensic auditor interface may work perfectly well as a third party plug n play tool to support the same. This might be sufficient and trustworthy for managing the UTECH private cloud, however not otherwise. The idea is that a common web-facing system interface that allows access to the

provider's underlying file systems and hypervisors, can be used to provision, start and stop virtual machines, and manipulate the firewall.

The management plane is particularly attractive because it is user driven. The provider, end users, and law enforcement could download log files, disk images, and packet captures from the management plane on demand. Further, with log forensic acquisition occurring under the hypervisor, retrieving VM images and other hypervisor log data would require trust only in Layers 3 and below. While attractive, this solution does require trust in the management plane, a potential vulnerability which does not exist with non-virtualized, physical computers. As a web interface, the management console opens a new attack surface which must be protected by the provider. Access to the management plane should be logged and strictly enforced with identity and access management. Communication between the user and the management plane endpoint should be done securely (e.g., using SSH).

5.3. Forensic Support as a Service

Provider support for log forensic acquisition is a naturally acceptable choice by law enforcement. The provider is already pre-positioned to preserve and collect the hypervisor log data since they control both the platform and infrastructure. This includes accountability for all logging mechanisms, packet captures and billable records. Technology for remote log acquisition would be moot if the provider and its infrastructure were trusted and the provider was willing and able to provide log evidence to the investigator directly. At their choosing, providers could offer these services to their clients with little effort and cost. Voluntarily doing so would demonstrate their care for security, and put reluctant security-minded clients at ease knowing that the log investigation was indeed possible. At least one provider, Terremark, offers forensic-as-a-service [27]. Potential drawbacks to a forensic support service include response time (potentially mitigated by the Service Level Agreement) and the provider's lack of knowledge about how customers are using the cloud to meet their goals.

Consider the following protocol for trust-preserving, provider-assisted log evidence production. Law enforcement serves in the author's case, the University acting as a private data cloud service provider with a search warrant for log data related to a particular IP address, including the client records for the user of that IP and the virtual machine serving content. A UTECH technician, certified as a forensic examiner by an independent third party, sits down at an offline forensic workstation connected to the backend cloud data centre, which in this context should also be treated as the cloud provider. The provider executes the warrant and gathers the log data requested, validating the data with cryptographic checksums [44]. Among the system log data requested are virtual machine CPU-ID addresses (i.e. Virtual IP addresses), access logs from the VMWare Management Console, data provenance logs, netflow records for the requested IP, and firewall logs. The data are copied to media for law enforcement. This protocol works at Layer 3, which requires trust in the VM host operating system, hardware, and network. The assumption for this protocol still requires trust in the hardware integrated with the TPM. Hence the basic assurances that the operating system, network, and technician are trustworthy are observed.

6. DISCUSSION

The nature of online remote forensics introduces security considerations. For example, a UTECH forensic examiner's workstation must have access to the Internet to acquire the evidence. While precautions such as firewalls and proxies may help shield the UTECH workstation from attack and compromise, the possibility of infection becomes more likely than if the UTECH workstation were standalone or on an isolated component of the UTECH LAN. This risk must be accepted, or remediated with appropriate monitoring and patching technologies.

One attractive feature of using existing tools which are executed by the cloud forensic examiner, as in Experiment 1, is that no changes to the UTECH cloud infrastructure are necessary, and no assistance from the provider is required. VM log introspection, as in Experiment 2, requires considerable change to the environment made by a provider, even though that feature could be exercised without the provider's intervention. Hypervisor log data export, as in Experiment 3, requires no change to the infrastructure, but must be executed by the provider. The experiments discussed by the author assume that the cloud consumer is the victim of the crime and the plaintiff in the investigation.

However, an equally likely scenario is one in which a criminal creates a system in the cloud, uses it to commit a crime, and removes the cloud system entirely. This situation demands proactive logging of data by the provider which may be of forensic relevance in the future. Shields, et al.[25] created a proof-of-concept continuous forensic evidence collection system that could be used to record the creation and deletion of cloud provisions. Finally, if the cloud provider is the criminal, the forensics service is also suspect and another alternative must be considered to investigate the crime. This specific issue presents an even larger overall behavioural trust concern with the participating cloud forensic data centres[43] and warrants through research, the need to establish suitable verification trust scoring systems to appraise these participating cloud provider agencies within forensic investigations. Hence a reasonable solution is to ensure that investigations be supported through a centrally trusted forensic collaboration agency with oversight responsibilities for the cloud data centres participatory to the investigation.
At the local level of the data centre there is the forensic shortcoming, and potential legal problem, where there is the lack of validation for the disk images or the relevant hypervisor logs that contain the metadata snapshots [41, 42]. Forensic examiners are accustomed to using cryptographic hashes to validate that the copy of a hard drive that they have taken is identical to the original. With no hash available for the original hypervisor log data source, examiners are unlikely to accept the log evidence. In the UTECH experiments presented in this survey, one is unable to verify cryptographically that the cloud images were identical to the standalone control because of differences such as different hardware and network configurations. The hypervisor log file timestamps is the only approach used in prior work to support this claim[29,30,44]. These differences did not affect the ability to reconstruct the crime.

The EnCase Servlet and FTK Agent used for these UTECH experiments had some limitations. These programs typically have System privileges, giving them undesired access privileges to both memory and disks. However, as with all software, they are vulnerable to malicious code that may have already compromised the target machine. The agent could be installed at any time in the lifecycle of the virtual machine; installing at the time the VM is provisioned, prevents the disruptive installation after an incident has taken place. Cloud providers such as Amazon employ user configurable firewalls that must also be opened to allow the agents to communicate with the command and control node. Though not inherently a vulnerability, open ports do increase the attack surface. Fortunately, EnCase and FTK employ network encryption between the client and server to provide confidentiality and authentication. Consumers must consider the cost associated with a remote log forensic investigation. Imaging and retrieving a virtual hard drive and its associated memory will incur potentially significant bandwidth costs. The experiment used an instance with a 18 GB virtual disk and 1.7 GB memory. It is well understood that Amazon currently bills outbound data transfer at $0.150 per GB, for the first 10 TB / month. Therefore, the retrieval of the disk and memory images totaled only $3.60. One TB of UTECH log data, which would cost $150.

6. CONCLUSION

In this survey paper the author presented the technical and high level conceptual trust issues that arise in acquiring forensic evidence from the virtual machine (VM) host operating systems within the data clouds. This assessment was done at the University of Technology [UTECH], Jamaica, which currently functions as its own independent cloud provider. This data acquisition is particular to the hypervisor system logs that can be used to track VM incidences which are then later used to compile snapshots of potential evidence for an investigation. This work also presented a model to show the layers of virtualization trust that can arguably be used to support the collection of log forensic data within these logical clouds. The author also provided the context for the support of such cloud digital investigations and analyzes the choices available to a forensic investigator. This was achieved by making a comparative experimental evaluation of popular forensic acquisition tools including Guidance EnCase and AccessData Forensic Toolkit, as to how volatile and non-volatile related hypervisor log data can be collected. Finally the paper explored three possible solutions for the managed hypervisor log data acquisition process within the data clouds. Ongoing work explores the proof of concept cloud simulator similar to Gridsim [43] that engages trustable remote cloud forensics within participating forensic cloud data centre agencies. Future work will be required to allow cloud clients to retrieve forensic logs and metadata (e.g., cryptographic checksums of disk volumes) directly from the online management log auditing facilities and assessing the legal merit of this. Equally there is the need for solutions to preserve log privacy as inspired by prior work[50] and prevent the loss of log forensic evidence when cloud resources are released.

REFERENCES

[1] Amazon Web Services. AWS Import/Export. Available at http://aws.amazon.com/importexport/; 2011. Last accessed December28, 2011

[2] Casey E. Digital Evidence and Computer Crime: Forensic Science,Computers, and the Internet. 2nd ed. Amsterdam: Elsevier AcademicPress, 2004.

[3] Conover M, Chiueh T. Code Injection From the Hypervisor: Removing the need for in-guest agents. In: Proceedings of Blackhat USA. 2008. Last accessed November 1, 2011.

[4] Dolan-Gabitt B, Payne B, Lee W. Leveraging Forensic Tools for Virtual Machine Introspection. Technical Report, Georgia Institute of Technology, GT-CS-11-05; 2011.

[5] Dykstra J, Sherman AT. Understanding Issues in Cloud Forensics: Two Hypothetical Case Studies. In: Proceedings of the 2011 ADFSL Conference on Digital Forensics Security and Law. ASDFL; 2011a. p. 191–206.

[6] Dykstra J, Sherman AT. Understanding Issues in Cloud Forensics: Two Hypothetical Case Studies. Journal of Network Forensics 2011

[7] Federal CIO Council. Guidelines for the Secure Use of Cloud Computing by Federal Departments and Agencies (Draft Version 0.41).2011.

[8] Garfinkel S. Forensic Feature Extraction and Cross-Drive Analysis. Digital Investigation 2006;3:71–81.

[9] Garfinkel T, Rosenblum M. A virtual machine introspection based architecture for intrusion detection. In: Proceedings of the 10th Annual Symposium on Network and Distributed System Security (NDSS 2003). 2003. p. 191–206.

[10] Giobbi R, McCormick J. Vulnerability Note VU912593: Guidance EnCase Enterprise uses weak authentication to identify target machines. Available at http://www.kb.cert.org/vuls/id/912593; 2007. Last accessed September 21, 2011.

[11] Guidance Software . EnCase Legal Journal. Available at http://www.guidancesoftware.com/DocumentRegistration. aspx?did=1000017380; 2011. Last accessed September 21, 2011.

[12] Heiser J. Remote forensics software. Gartner RAS Core Research Note G00171898; 2009.

[13] Krautheim FJ. Building Trust into Utility Cloud Computing. Ph.D. thesis; Department of Electrical Engineering and Computer Science, University of Maryland, Baltimore County; Baltimore, Maryland; 2010.

[14] Krautheim FJ, Phatak DS, Sherman AT. Trusted Virtual Environment Module: Managing Trust in Cloud Computing. In: 3rd International Conference on Trust and Trustworthy Computing. 2010. p. 211–27.

[15] LibVMI . Virtual Machine Introspection (VMI) Tools. Available at http://vmitools.sandia.gov/; 2011. Last accessed November 1, 2011.

[16] Nance K, Hay B, Bishop M. Investigating the Implications of Virtual Machine Introspection for Digital Forensics. In: Proceedings of the International Conference on Availability, Reliability and Security (ARES '09). 2009. p. 1024–9.
[17] National Institute of Standards and Technology. Computer forensic tool testing (CFTT) project overview. Available at
[18] National Institute of Standards and Technology. Digital Data Acquisition Tool Specification. Available at http://www.cftt.nist. gov/Pub-Draft-1-DDA-Require.pdf; 2004. Last accessed September 21, 2011.
[19] National Institute of Standards and Technology. Test Results for Digital Data Acquisition Tool: FTK Imager 2.5.3.14. Available at http://www.ncjrs.gov/pdffiles1/nij/222982.pdf; 2008. Last accessed September 21, 2011.
[20] National Institute of Standards and Technology. Test Results for digital Data Acquisition Tool: EnCase 6.5. Available at http://www.ncjrs.gov/pdffiles1/nij/228226.pdf; 2009. Last accessed September 21, 2011.
[21] National Institute of Standards and Technology. The NIST Definition of Cloud Computing. Available at :http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf; 2011. Last accessed May 4, 2012.
[22] Ruan K, Carthy J, Kechadi T, Crosbie M. Cloud forensics: An overview. In: Advances in Digital Forensics VII. 2011.
[23] J. P. Wilkinson, "Nonlinear resonant circuit devices (Patent style)," U.S. Patent 3 624 12, July 16, 1990.
[24] SCMagazine . Best computer forensic tool. SCMagazine 2011; Last accessed May 15, 2012.
[25] Shields C, Frieder O, Maloo M. A system for the proactive, continuous, and efficient collection of digital forensic evidence. In: The Proceedings of the Eleventh Annual DFRWS Conference. Volume 8; 2011. p. S3–13.
[26] Taylor M, Haggerty J, Gresty D, Lamb D. Forensic investigation of cloud computing systems. Network Security 2011; 2011(3):4–10.
[27] Terremark . Secure Information Services. Available at http: //www.terremark.com/uploadedFiles/Services/Security_Services/TMRKSIS_Gatefold2_4pagelayout_Screen.pdf; 2009. Last accessed November 1, 2011.
[28] United States Code. Communications Assistance for Law Enforcement Act (CALEA). 47 USC 1001-1010; 1994.
[29] Sean Thorpe , Indrajit Ray. FileTimestamps in a Cloud Digital Investigation. Proceedings of the Journal of Information Assurance and Security, pg. 495-502 ,Vol 6, Issue 6 , December 2011.
[30] Sean Thorpe, Indrajit Ray. Detecting Temporal Inconsistencyin Virtual Machine Activity Timelines. Proceedings of the Journal of Information Assurance and Security pg. 24-31,Vol. 7, Issue 1, May 2012 .
[31] Sean Thorpe, Indrajit Ray, Indrakshi Ray, Tyrone Grandison. A Formal Temporal Log Data Model For The Global Synchronized Virtual Machine Environment. Proceedings of the Journal of Information Assurance and Security . pg. 398 -406, Vol 6, Issue 5, March 2011.
[32] Sean Thorpe, Indrajit Ray , Indrakshi Ray , Tyrone Grandison, Abbie Barbir. A Global Virtual Machine Attribute Access Control Policy for Auditing Federated Digital Identities within a Compute Cloud Environment. Proceedings of the Journal of Information Assurance and Security. pg. 322 -330, Vol, 6, Issue 4, December 2010
[33] Sean Thorpe, Indrajit Ray, Tyrone Grandison, Abbie Barbir. The Virtual Machine Log Auditor ToolKit. Proceedings of the Journal of Information Assurance and Security Letters(IASL), pg.37-44,Volume 2, December 2011.
[34] Sean Thorpe, Indrajit Ray, Tyrone Grandison. "Enforcing Data Quality Rules for a Synchronized VM Log Audit Environment using Transformation Mapping Techniques". The Proceedings of the 4th Intl Conference on Computational Intelligence in Security for Information Systems, Torremolinos, Malaga, Spain. June 8-10, 2011.
[35] Sean Thorpe, Indrajit Ray, Tyrone Grandison. "Use of Schema Associative Mapping for synchronization of the Virtual Machine Audit Logs". The Proceedings of the 4th International Conference on Computational Intelligence in Security for Information Systems, Torremolinos, Malaga, Spain. June 8-10, 2011.
[36] Sean Thorpe, Indrajit Ray, Tyrone Grandison. "Enabling Security Uniformly Across Cloud Systems". ACM ASPLOS (Architectural Support for Programming Languages and Operating Systems) RESOLVE (Runtime Environments/Systems, Layering, and Virtualized Environments) Workshop. Newport Beach, California. March 5, 2011.
[37] Sean Thorpe, Indrajit Ray, Tyrone Grandison. "A Synchronized Log Cloud Forensic Framework". Proceedings of the International Conference on Cybercrime, Security & Digital Forensics. Glasgow, UK. June, 2011.
[38] Sean Thorpe, Indrajit Ray, Robert France, Tyrone Grandison, Indrakshi Ray, Abbie Barbir. "Towards Formal Semantic Annotations for Compute Cloud Forensic Investigations using Synchronized Log Audit Mechanisms". Proceedings of the International Conference on Cybercrime, Security & Digital Forensics. Glasgow, UK. June, 2011.
[39] Sean Thorpe, Abbie Barbir, Indrakshi Ray, Tyrone W A Grandison, Indrajit Ray. "Formal Parameterization of Log Synchronization Events within a Distributed Forensic Compute Cloud Database Environment". The 3rd International ICST Conference on Digital Forensics & Cyber Crime (ICDF2C). Dublin, Ireland. Oct 26-28, 2011.
[40] Sean Thorpe, Indrajit Ray, Tyrone Grandison, Abbie Barbir. "Cloud Digital Investigations based on a Virtual Machine Computer History Model". The 6th International Symposium on Digital Forensics and Information Security Conference (DFIS-12). Vancouver, Canada, 26-28 June 2012.
[41] Sean Thorpe, Indrajit Ray, Tyrone Grandison , Abbie Barbir. Cloud Computing Log Evidence Forensic Examination Analysis. Proceedings of the 2nd International Conference on CyberCrime, Security and Digital Forensics, May 14-15, 2012 , London , UK.

[42] Sean Thorpe , Indrajit Ray, Tyrone Grandison, Abbie Barbir. Cloud Log Forensics Metadata Analysis. Proceedings of IEEE Computer Forensics and Software Engineering Workshop(Co-located with COMPSAC 12) , Izmir , Turkey , July 16 -20, 2012.

[43] Sean Thorpe, Tyrone Grandison, Indrajit Ray, Abbie Barbir. Trust Among Participating Cloud Forensic Data Center Agencies. Proceedings of the Security Data Management (SDM)workshop ,VLDB12, Istanbul ,Turkey, August 27-30 2012.

[44] Sean Thorpe, Indrajit Ray, Tyrone Grandison, Abbie Barbir. Formal Hash Compression Provenance Techniques For The Preservation of the Virtual Machine Log Auditor Environment. Proceedings of the International Journal of Information Science and Computer Applications (IJISCA), Vol. 1 , 2012.

[45] Santana M. Cloud Security: Beyond the Buzz. Available at: http://www.linuxworldexpo.com/storage/10/documents/CI7\20Mario\20 Santana.pdf; 2009. Last accessed September 21, 2011.

[46] Waldo Deport , Martin Oliver , Michael Kohn. Isolating a Cloud Instance for Digital Forensic Investigation. Eighth Annual IFIP WG 11.9 International Conference on Digital Forensics, Pretoria, South Africa, January 2012

[47] Ryan K.L. Ko , Peter Jagadpramana,Bu Sun Lee. " Flogger – A File Centric Logger For Monitoring File Access and File Transfers within Cloud Computing Environments. 2011 International Joint Conference of IEEE Trust-Com

[48] Raluca A Popa , ,Jacob A Lorch , David Molnar, Helen J Wang, Li Zhuang. Enabling Security in Cloud Storage SLAs. Proceedings of Usenix 2011.

[49] Peter Mell , Timothy Grance. NIST Definition of Cloud Computing. Retrieved from: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.PDF, September , 2009.

[50] Rajeev Bedhi, Mohit Marwaha , Tajinder Singh, Harwinder Singh, Amritpal Singh. Analysis of Different Privacy Preserving Cloud Storage Frameworks. Proceedings of the International Journal of Computer Science and Information Technology(IJCSIT). Vol 3, No.6 , December 2011.

Author

Sean Thorpe holds an M.S. and B.S. degrees in Information Security and Computer Science respectively from the University of Westminster, London, UK in November 2002 and from the University of the West Indies, Mona Campus Jamaica in November 2000. He joined the University of Technology (UTECH) as a Lecturer in January 2003 with responsibility for teaching System Security at the undergraduate level. Mr. Thorpe has worked extensively in t he IT industry since 1995 as a System Programmer Analyst and Oracle DBA before joining academia. He is a 2009 recipient of the Fulbright Visiting faculty Scholarship award to Harvard University, where he explored collaborative research work in the area of Security Metrics. He is also the 2009 winner of the OOPSLA Educational Symposium Award for his innovative computer science teaching methods, and the recent 2011 and 2012 American National Science Foundation (NSF) awardee for Caribbean based research in the area of Cloud Computing and data mining respectively. His specific research interest includes cloud forensics, and security policies.