Isolating instances in the cloud

W. Delport and M. S. Olivier

2012

Note that the attached preprint of this paper was titled *Isolating instances in cloud forensics*. The title of the published paper was changed to *Isolating instances in the cloud*. The correct citation of the published version of the attached paper is therefore:
W. Delport and M. S. Olivier. "Isolating instances in the cloud". In: Advances in Digital Forensics VIII. Ed. by G. Peterson and S. Shenoi. Springer, 2012, pp. 187–200

Isolating instances in cloud forensics.

Waldo Delport and MS Olivier
Information and Computer Security
Architectures Research Group
Department of Computer Science
University of Pretoria
South Africa
wdelport@cs.up.ac.za
molivier@cs.up.ac.za

January 2, 2014

Abstract

Cloud Computing is gaining acceptance and increasing in popularity. Organizations often rely on Cloud resources to effectively replace their 'in-house' computer systems. In a Cloud environment, an instance is typically accepted to be a virtual system resource, established within that Cloud. Multiple instances can also form one logical instance. Multiple instances can be contained within a single node. The Cloud itself consists of multiple nodes.

In order to facilitate a Digital Forensic Investigation (DFI) a Digital Forensic Process (DFP) is followed. The DFP defines a finite number of steps to aid in the successful completion of the DFI. Isolating the environment is generally accepted within the Forensic Community to be an integral part of a Forensic process. We contend that this isolation is also needed in a DFI. The isolation prevents any further contamination of or tampering with possible evidence.

This paper serves as a starting point for a possible Digital Forensic Process (DFP) to facilitate an investigation in a Distributed Instance System (DiS) environment. It focuses on the process to isolate an instance. Conditions are introduced to act as a guide for the successful isolation of an instance. It also considers the complications that can occur when doing a digital DFI in a DiS.

Keywords: Cloud Computing, Digital Forensic, Digital Forensics Process, Isolation.

Published as: W. Delport and M. S. Olivier, "Isolating instances in cloud forensics", in *Advances in Digital Forensics VIII*, ser. IFIP Advances in Information and Communication Technology, G. Peterson and S. Shenoi, Eds. Springer Berlin Heidelberg, 2012, vol. 383, pp. 187 - 200.

1 Introduction

Distributed computing enables workloads to be distributed [1]. As distribution technologies grow, the network infrastructure in the world is also changed [2]. The network infrastructure is getting more reliable and faster. Virtualization inside a distributed computing environment enables virtual resources to be provided. When the network infrastructure was able to handle large volumes of data, fast and reliably, it became apparent that interaction with virtual resources over the network would be possible. Cloud Computing was thus born, enabling a service provider to provide virtual resources over the network [1].

In Cloud Computing, all interactions and workings of the Cloud are digital by nature [3]. When something erroneous happens and an investigation is required, a Digital Forensic Investigation (DFI) needs to be done [4]. When doing a DFI, a Digital Forensic Process (DFP) is followed [5]. The DFP enables admissible evidence to be gathered from the investigation. If the evidence is not admissible, there is little justification for doing a DFI.

In a cloud an instance must be isolated when it becomes apparent that an incident happened on that particular instance. This isolation helps preserve the integrity of the evidence collected from the instance. One of the problems associated with preserving the integrity of an instance is an attribute of Clouds [6]. In a Cloud, the data from one instance may share the storage of multiple instances, and may not be in a constant place in the Cloud. To preserve the integrity of the evidence, the location on the Cloud must be known and must be protected from tampering and contamination. Another complexity is that other instances on the same node may belong to other users. Users should, at least, be able to expect availability and privacy of their instance provided from the service provider [7]. The Digital Forensic Process must be done in a manner that will not result in the loss of privacy of other instances, and the availability of the instance must be affected in the smallest possible manner.

This paper introduces conditions that we argue must be met to confirm the successful isolation of an instance. The conditions are then then tested in an experiment. The necessary methods to meet each condition are explained and experimental data is given. The paper then looks at the collection and examination phase of a DFP, with the focus on the isolation of an instance.

In section 2, a brief description of Cloud Computing is given. Section3 provides an overview of two DFI procedures. In section4, the term isolation is explained. We also identify conditions that we consider necessary in order for an instance to be seen as isolated. In section5, the different service models and deployment models are discussed in terms of isolation. A description of a DiS cloud is given in section6. Section7 introduces a method to locate an instance in an Infrastructure as a Service (IaaS) Cloud. Section8 focuses on isolating the instance by blocking communication. The problems regarding

contamination and availability of evidence will be discussed in section9. In section10 the process of removing unrelated information is discussed. The second last section, section11, focuses on the problems that can arise in a DiS cloud, where multiple suspect instances exist.

2 Cloud Computing

Cloud Computing is a relatively old term but has been adopted quickly over the last couple of years [8]. Cloud Computing builds on different forms of distributed computing. It ties distributed computing together with virtualization. Cloud Computing enables a service provider to provide a flexible, cost effective and on-demand infrastructure to its clients, instead of the clients running their own infrastructure. For the purpose of this paper Cloud Computing will be defined as a distributed computing architecture providing flexible, cost effective and on-demand resources to users over some form of network by using virtualization to create virtual resources. There are three types of Cloud Computing service models namely Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS) [9].

The first service models is *Infrastructure as a Service*. The users of a Cloud infrastructure are provided a virtual computer which can be interacted with, usually through the Internet [6]. This virtual computer needs to be setup and maintained by the user. This virtual computer can also be referred to as an instance. Normally, an instance can be accessed from anywhere in the world, depending on the security setup. The instance can be a small instance, used by a single user to store backups of files, or it can be a server running the website and database of a company. A user only pays the service provider for services rendered. If the requirements of the user change in terms of computational power or storage space, it is an easy process to change the scope of the instance to accommodate the new requirements of the user. If a new instance is required the task of starting up an instance is trivial. The service provider is responsible for maintaining the Confidentiality, Integrity and Availability (CIA) of the instances on a hardware level. The user is responsible for protecting the CIA on a higher level, e.g. the content of files and the operating system [10].

The second service model is *Platform as a Service*, where the user is provided a platform that is maintained by the Cloud service provider [9]. The platform is an instance that was created with a specific focus by the service provider. The service provider may create a default platform for a web server. The user must then configure the application on the platform. The service provider may also provide the necessary tools to successfully build upon the platform.

The last service models is *Software as a Service*, where software is made

available through the use of Clouds. The application and the data of the application are seen as the resources on the Cloud [11]. The user pays to get access to an application that can be customised according to the requirements of the user. The user has no concerns related to the underlying hardware and software below the application of interest.

There are four deployment models for Clouds. They are Public, Private, Hybrid and Community [3].

In a *Public* Cloud, the infrastructure is owned by a Cloud service provider. The service provider will sell resources of the Cloud to other companies and the public. The service provider is responsible for managing the Cloud.

In a *Private* Cloud, the Cloud infrastructure is for the use of one company only. The company owns the Cloud and uses the resources. The Cloud infrastructure can be on company property or may be located elsewhere. The company, or a contracted company, is responsible for maintaining the Cloud.

If the Cloud infrastructure is for the use of several companies, it can be seen as a *Community* Cloud. The companies own the Cloud and use the resources collectively. The companies form a community with shared interest. The Cloud infrastructure can be on one of the companies' properties or may be located elsewhere. The companies, or a contracted company, would be responsible for maintaining the Cloud.

Hybrid is a combination of at least two of the above models. Each of the models used is still a separate entity in the Hybrid Cloud. Forms of technology bond the entities together. This is normally used for load balancing.

The value that can be added from Cloud Computing is significant, primarily to Small, Medium and Micro Enterprises (SMME) [12]. These enterprises have a survivalist methodology [13]. These enterprises are started without large start-up investments and with a specific knowledge base. Cloud Computing enables these enterprises to have access to servers without the initial start-up cost and they have no maintenance costs on a hardware level. As the businesses grow, their infrastructure can easily be changed to adapt to the growth. Another added benefit is that most of the knowledge required to build the infrastructure is not needed at the enterprise, as the Cloud service provider is responsible for this.

Cloud Computing is growing and is estimated to become a billion dollar industry this year [14]. The reason for this is that some of the largest IT related companies have implemented or are implementing Cloud Computing. Some of the large companies are Google, Microsoft, IBM and Amazon [10, 6]. These companies state that they will provide CIA to their customers by using various techniques.

3 Digital Forensics Process

In order to obtain admissible evidence a well-defined forensic process needs to be followed. Cohen [5] proposes a model for the digital forensic examination that consists of seven phases. The phases are Identification, Collection, Transportation, Storage, Examination and Traces, Presentation and Destruction.

The *Examination and Traces* phase consists of four subcategories: Analysis, Interpretation, Attribution and Reconstruction [5].

Documentation is a continuous process and needs to happen in all phases of the digital examination. One of the main aids to help preserve the integrity of the evidence is documentation. The documentation should at least include the name of the evidence and the place the evidence is gathered. The documentation should also include the processes followed in identifying, retrieving, storing and transporting the evidence. The documentation should also mention the chain of custody when the examination was in progress. There have been several cases where the outcome of the case was influenced by the documentation.

There are other DFPs. They include most of these phases or a combination thereof. One such prominent DFP was defined by the National Institute of Justice (NIJ) [4]. The phases defined are Collection, Examination, Analysis and Reporting.

From the two processes it can be seen that they include the same set of underlying steps. The process by Cohen is subdivided into more steps. This enables a more systematic flow of events.

The normal computer forensic process uses static or "dead" analysis [15]. In a "dead" analysis, the system is turned off as soon as the examination team acquires it and images are made of the storage mediums; the analysis is then conducted on the images. The other approach is a "live", analysis where the computer is kept on and evidence is gathered from the computer in the environment that is on the system. There are advantages and disadvantages to both approaches . The main disadvantage of a "dead" analysis is the fact that some information may be lost because it is in a buffer or the ram. A problem with a "live" analysis is that the evidence can be unintentionally destroyed or modified without the intent to do so [15].

The Cloud computing added complexity to the DFI. As stated in section 1 one problem is that in a Cloud multiple users data may be contained on the storage of alongside the data of the suspected instance [6].

4 Isolation

The term "isolate" is defined as the state of being identified and then separated from others [16]. Isolation is the process of isolating. In a "real world"

forensic process the crime scene is isolated [17]. The isolation helps protect the possible evidence from contamination and loss of continuity. To aid the admissibility of the evidence a crime scene is divided into separate parts to aid in the isolation. These parts can only be entered by authorised personnel using an authorised manner. A path is laid out where the personnel can walk in and around the crime scene. A log is kept of where personnel members are and what they are doing.

In the Digital Forensics realm isolation is used although the terminology used differs. As an example, seized cell phones are placed inside a Faraday bag [18]. This stops the cell phone from communicating with the outside word. This is a form of isolation. The cell phone is separated from the world. When doing hard drive forensics on a hard drive, a write blocker is used to enable a write free read [19]. The write free read protects possible evidence from contamination. The process of isolating is being used in the field of Digital Forensics, although the term "isolation" is not often used in this context.

Gathering evidence is one of the aims of a DFI. If there is suspicion that the evidence is invalid by any means it will not be able to serve as admissible evidence. In order to add to the evidence's admissibility, the evidence needs to be protected from contamination and tampering. The need for isolation in the Cloud environment becomes apparent when taking the evidence's admissibility into account.

After careful consideration and experimentation a list of conditions for isolation was composed. We contend that these conditions needs to be met in order to state that an instance has been successfully isolated. The list is as follows:

- The instance‘s physical location is known (Location).
- The instance is protected from outside interference (Incoming Blocking).
- The instance is blocked from communicating with the outside word (Outgoing Blocking).
- Possible evidence from the instance can be gathered (Collection).
- The possible evidence is not contaminated by the isolation process (Non-Contamination).
- Information unrelated to the incident is not part of isolation (Separation).

The name for each conditions is given in brackets. The rest of the section will describe the conditions in more detail.

In order to know the physical location of an instance, it must be known were in the Cloud the instance is. This means finding the node on which the instance resides.

The instance must be protected from other instances and other external sources. In the Cloud, all interaction with an instance is through the network. This means the network connection must be blocked. This will aid in the protection against contamination of evidence and the tampering with evidence.

The instance might suspect irregular activity and will try to send messages or try to tamper with evidence. The instance must thus be blocked from sending messages over the network. The instance must also be blocked from tampering with files and other information on it.

It must be possible to collect all possible evidence from the instance. The possible evidence may be running programs, information in the swap space and information on the hard drive. It must still be possible to collect all possible evidence after the instance has been isolated on the node.

If the evidence is contaminated on the instance, the isolation process has failed. The isolation process is conducted to protect the evidence. If it does not protect the evidence there is no reason to do the isolation.

On a node, data from multiple instances can be contained. The isolation must be done in a manner such that unrelated information is excluded from the isolation but related information is isolated. This protects the confidentiality of information.

5 Isolation on different types of Clouds

In the previous section isolation was introduced and explained. Clouds differ in deployment and service models. This difference has an impact on isolation. This section focuses on explaining isolation within each deployment and service model. Some problems regarding each model are also introduced. This paper will focus on Private and Public deployment methods. We contend that Hybrid and Community Clouds can fit into either Public or Private Clouds for the discussions to follow in this paper.

In the *deployment method* of Clouds, the difference is where the data is located, and who owns the data. The location and owner of the data are an important considerations when isolating.

When performing a DFI on a *Private* Cloud all the information on the Cloud belongs to a company. If that company is doing the DFI, the concern for confidentiality is lower. If an external company is doing the DFI, or is responsible for the DFI the concerns change. The company whose Cloud is being investigated is concerned about keeping the confidentiality of the Cloud. Only information related to the incident needs to be found. The other information should stay confidential. In the case where the company

is doing the DFI, the first five conditions of isolation needs to be met. The last condition, separation, is not of high importance. When an external company is doing the DFI all six conditions are important.

On a *Public* Cloud, the data on the Cloud belongs to different users and the users can even be in different countries. This brings a jurisdiction problem into the Cloud, as not all the information may be under the same legal system. On a Public Cloud the users pay the service provider to keep their information safe and available. The service provider is thus responsible for protecting the confidentiality and availability of their clients' resources. The party responsible for the onset of the DFI would normally be an external source. The external source is thus interested in finding evidence, and this may not align with the concerns of the Service provider. In a public Cloud, isolation is required to protect other users of the Cloud. This means all six isolation conditions needs to be met.

The `service models` of Clouds separate the data into different layers inside the Cloud. Each layer holds separate conditions for isolation.

In a *IaaS* Cloud, each instance on the Cloud may belong to a different user. Each instance can be seen as a virtual machine. The isolation of an instance on a IaaS service model Cloud involves isolating the entire instance. This is done because each instance can be unique. Another reason is that each instance is an unknown computer. The instance might destroy evidence on itself or contaminate other instances on the Cloud. The isolation is thus done to protect evidence on the instance and to protect other instances from contamination.

In a *PaaS*, Cloud the underlying platform of each instance is known. The differences between instances are the software installed on the instance and the data stored on the instance. In order to isolate the instance, the applications running on the instance and the data of the instance should be isolated. Because the underlying platform is known, most of the capabilities of the instance are known or can be controlled.

The only difference between *SaaS* instances are the setup of the application and the data stored in that instance. The possible evidence can lie in the settings and data. This means that this is the only part of the instance that needs to be isolated. All of the capabilities of the instance are known.

The difference in Cloud deployment and service models in terms of isolation has been explained. Clouds can also differ in terms of implementation on these models.

6 Distributed Instance System (DiS)

The previous section focused on isolating a single instance. This section focuses on introducing an environment where there are multiple suspect instances. A unique implementation is using a IaaS Cloud to build a type of

server farm. In Cloud Computing availability, is regarded as a necessity [3]. One model of Cloud Computing providing availability is a DiS. In a DiS, multiple instances on the Cloud are combined to form one logical resource. This combination aids in protecting the availability of the resource. When one of the instances malfunctions, it is discarded. A new instance may be started to take its place. This means that the instances are dispensable.

A server farm is a multi-node system [20]. In web server farms, the web-site is split over two or more nodes. The user interacting with the website only sees the functionality of a single server. In the server, multiple nodes are used to deal with the website. The server farm uses some form of routing to route requests from users between nodes from users. The server farm uses distribution technologies to enable this service. This distribution aids in the quality of service of the web-site. There is no single point of failure. When a node fails, the router will stop sending requests to that node.

In the Cloud, the DiS can be located over several nodes. This also enables the service provider and the client to provide continuous availability. If a node fails, it will have little to no effect on the overall DiS. We will now look at the isolation process of a single instance. In section 11, we will look at isolation in a DiS environment.

7 Locating an Instance

The first condition proposed is to locate the instance. In order to locate an instance, the node on which the instance is running needs to be found. One method involves using the Cloud management software. The software may provide the functionality to locate an instance. Another method is subnetting where a sub network is created for each node. The subnetwork is then used to trace the instance to a node.

We will use subnetting to do our experiment. For the purpose of this paper Nimbula Director will be used to run experimental tests [21]. Nimbula Director is a Cloud operating system providing IaaS in a private network. In the experiment an instance was launched on the Cloud. We then set out to find the node on which the instance resides.

We know the IP of the instance, how the sub-networking is done and what a wire address is. The IP was provided by Nimbula when the instance was started. In our experiment the IP is 10.128.0.10. A sub-network (Sub-net) is created by Nimbula for each instance group on the Cloud. The wire address is the network connections address and has the lowest IP in the network [2].

To find the instances using the information we have, we need to find a link between the networks of the instance and nodes. We will look at all the network information of all the nodes and the network information of the instance. To start, `ifconfig` command is run on each node.

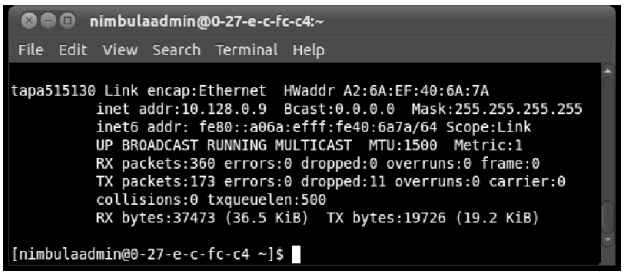

Figure 1: The ifconfig command on a node

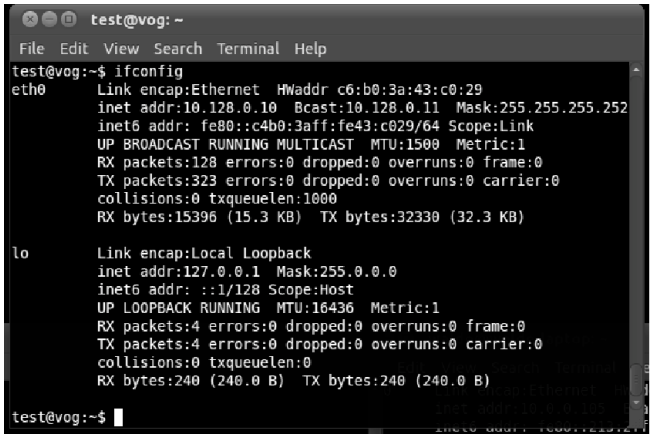

Figure 2: The ifconfig command on the instance

Figure1 contains a part of the `ifconfig` output of a node in the Cloud. The figure shows the IP address of one of the network interfaces on the node.

In Figure2, the `ifconfig` output is given of an instance running on the Cloud. The figure shows the network interface IP address and other network attributes.

As illustrated in Figure2, the Subnet mask for the instance is 255.255.255.252, thus `11111111.11111111.11111111.11111100` [2]. The netmask uses the first 30 bits as the network ID. The remaining two bits can be used as the address. This causes the availability of three IP addresses on the instance's private network.

In Figure2, the instance‘s IP address is 10.128.0.10, and the broadcast address of the instances network is 10.128.0.11. It can thus be inferred that because there are three addresses available, and two of the addresses have been used, that the wire address is 10.128.0.9. This can be inferred when it is taken into account that the wire address can be calculated when all the address bits are set to zero and then adding one to the address [2]. This means that the wire address is the first address in the Subnet.

In Figure1, the IP address of one of the devices of the node is 10.128.0.9. This corresponds to the instance's wire address. It can then be logically inferred that the instance is located on the node with a device which has the same address as an instance's wire address. Thus, to find an instance, one must find a node with a network device that has an IP one lower than the IP of the instance.

In the experiment, the location of the instance was discovered by accessing nodes and an instance. The experiment outcome concludes that an instance can be found without accessing the instance. This enables the search to avoid detection by the instance. This will aid in protecting the possible evidence from possible contamination.

8 Blocking Communication

The first condition was to locate the instance. The location has now been discovered. Conditions two and three are to block ingoing and outgoing communication. In the Cloud most communication is through the network [3]. In order to protect the instances from outside interference and block communication with the outside world one can stop the network. Multiple methods exist to block network communication. One option is to bring down the the network. Another option is to use a piece of software to block the network traffic, e.g. a firewall.

In-order to bring down a network on a Linux computer the network the command `ifconfig eth0 down` can be used. The command `ifconfig tap8d6cb50`

`down` was run on a node with a network interface to an instance. This stops the device from sending or receiving data. If one attempts to establish an ssh session into the instance, it fails because there is no network traffic allowed to the device. This method also stops the instance from communicating with the outside world. This may not be ideal as all communication is stopped. The Cloud operating system and the investigators will not have access to the instance.

One of the other options is to use software. A firewall can be used. The firewall can be implemented in one of two manners. The one is on the node itself. The other is running on the Cloud itself. If communication is blocked using the firewall on the node the communication on a specific network device must be stopped. The network device is the virtual connection to the instance as described in Section7. The other option is a single large firewall to control the network on the Cloud. The instance IP is given to the firewall to block all communication to and from the device. Using a central firewall may be easier to maintain and to setup to block communication. An added advantage of using a firewall is that it may be implemented in a manner that will allow the DFI to gather evidence from the traffic that was blocked, as it can be logged.

Other software that can be used to block communication: the Cloud operating system is an example. This is not implemented in the Cloud operating system used for the current experiments. The administrator can state that a specific instance is blocked and the system is then responsible for the traffic blocking. The system might also provide some information. The information might include: to what other instances the suspect instance is communicating and the kind of communication used.

The methods used to block communication generally block all network traffic. The list of possible methods can be expanded to include using a piece of software to block specific communication. This will aid a "live" DFI investigation. All web traffic may be blocked as it is not required to do a DFI. Communication left open may include ssh connections. The problem is that there are unmonitored connections left open. These connections need to be monitored to avoid unauthorised use. To do this a specialised firewall might be used or some specific piece of software to allow certain connections from the DFI team but block all other communication.

9 Gathering Possible Evidence

Conditions four and five are that it must still be possible to gather evidence and that the evidence is not contaminated. When doing an isolation, the instances must be protect against contamination. When finding the instances, there is no direct communication or interaction with the instances. The communication blocking is done in a way that will minimise interference

with the instance. The precautions taken in the previous steps will aid in protecting the instances from contamination.

One of the outcomes of a DFI investigation is evidence. If the instances were protected from contamination the possible evidence on the instances will also be protected from contamination. The process of gathering possible evidence from the isolated instances is outside the scope of this paper. There are other studies focusing on these [22].

10 Separation

To create a clean crime scene the node must contain only the suspect instance. This can be done by moving the instance to a clean node or by moving the other instances from the node. Moving the other instances has the advantage that the suspect instance is unaware of the moving process. The process to move an instance can be done using a variety of methods [23]. In some cases the Cloud operating system can be used to move the instances.

11 Isolation of Cooperating Suspect Instances

In a DiS infrastructure it can also be the case that there are multiple suspect instances. The possible evidence can then be distributed over multiple instances. The problem is explained by a basic DFI problem [5]. The problem is when to stop gathering evidence. The instances used in the DFI will form a set of instances. In a DiS with multiple suspect instances there are three options for the size of the set used. One option is to gather evidence from all instances. The second option is to gather evidence from one instance. The third case is some middle ground between the two other options.

When all the instances are used in the set all the instances need to be isolated. In order to isolate all the instances all the instances need to be found. If one is missed the data set is incomplete. A possible problem with finding all the instances is the time needed to find them. The time that is used to find all the instances can be used by the suspect instances to contaminate evidence. When finding the instances a subset of the instances might find proof that a DFI is being carried out. This might cause the instance to respond harmfully towards the possible evidence. Another problem is the order in which the instances are isolated. The instances need to be isolated in manner that will not raise suspicion amongst the other suspect instances.

The second option is to only find one instance on the Cloud. One does not investigate or even search for other suspect instances. The other suspect instances are left untouched by the investigation. There are a few problems with this method. One of the problems is the amount of evidence. Will sufficient evidence be found on a single instance? If the DiS is setup in

such away that instances mirror each other, it may be assumed that enough evidence will be gathered. On the other hand, if the DiS is setup in a distributed manner and each instance only contains a small subset of the evidence, the total amount of evidence that may be gathered can be less than the small subset of evidence available. This causes evidence loss and an ineffective amount of evidence may be gathered. In a DiS, the loss of an instance is expected. When we isolate an instance, the other suspect instances may ignore the loss and continue normally.

The last option is to find some middle ground between the two previous options, were we only find a subset of the instances. The problems here are how many do we find and the speed at which we do it. If we find too few of the instances, we may not gather enough evidence. If we try and find too many we may be wasting resources. When isolating the instances one-by-one, it might be detected by other suspect instances and they may react by contaminating the evidence. There is no golden answer for how many instances need to be isolated or how many resources can be dedicated. This varies per investigation and also according to the investigation time frame itself.

When doing a DFI, it must be decided whether a "*live*" or "*dead*" forensic investigation will be followed [15]. On a Cloud four options are available: the "*live*" and "*dead*" as mentioned. The third is a new option, let us call it "*halfdead*". This is where the instance is dead but the node is alive. The node is used to gather evidence and is trusted. The fourth option where the node is dead but the instance is alive, we will call it "*resurrected*". This is where the node has been killed but the instances were restated in a new controlled environment.

When doing a DFI on a DiS Cloud a combination of "*live*" and "*dead*" forensics can be used. We called it *Community Live*. To do, this we isolate some instances and monitor others. The instances we isolate are subsets of the suspect instances. The isolated instances are used in a "*dead*" forensic manner. To gather "*live*" forensic evidence we monitor the network traffic of some of the other instances. The content of the ram and running processes on the instances can also be used. This means that there are five types of forensic investigation that can be followed on a DiS Cloud.

There are problems with this approach. Some of the problems have been mentioned previously but are applicable here. If too few instances are isolated there may not be enough "*dead*" forensic evidence. If too many instances are isolated, it might be noticed by other suspect instances and they may contaminate evidence. This means the "*live*" forensics may lead to inadmissible evidence.

The combination of both "*live*" and "*dead*" forensic evidence can prove a vital resource to a DFI in a DiS Cloud.

In Table1 an analysis of the network traffic is given that was captured in another experiment. In this experiment five instances were setup to work

Table 1: Analysed network traffic.

Address A	**Address B**	**Packets**	**Bytes**	**Bytes A->B**	**Bytes A<-B**	**Duration**
10.128.0.22	10.128.0.26	1046	81006	44051	36955	303.5644
10.128.0.26	10.128.0.34	1100	85200	38900	46300	294.145
10.128.0.26	10.128.0.30	1102	85332	38900	46432	292.0971
10.128.0.18	10.128.0.26	1122	86904	47226	39678	298.7616

together. They send traffic to a controlling instance and it responds with an answer. The captured data was retrieved from the node on which the first suspect instance is located. The suspect instance has the IP 10.128.0.26. The analysis shows that the network traffic can be used to detect the instances that are working together. In Table1, the second row shows that 10.128.0.22 sent 44051 to 10.128.0.26 and
10.128.0.26 sent 36955 bytes back to 10.128.0.22. The other rows can be read in the same manner. After the analysis we concluded that the instances with IP's 10.128.0.{18,22,30 and 34} are working together.

12 Conclusion

Cloud Computing will continue to expand and usage of Clouds will increase. On the other hand we know that crime will happen. This means that we need to be prepared to do a successful DFI on a Cloud. One of the sub problems for doing a DFI on a Cloud is isolation. This protects the evidence and other users.

In this paper we introduced six conditions which we argue need to be met in order to confirm to the successful isolation of an instance. We also provided some focus on isolation and explained its use.

The conditions were then explained. The methods to do each condition were introduced. Some experimental data was also given about the condition outcomes.

We considered the DiS environment. This is where multiple instances are used to form one logical resource to aid in protecting availability.

We also explained what problems can exist in a DiS environment when isolating. Solutions were given for the number of instances that need to be isolated in a DiS Cloud with multiple suspect instances. The "*live*" and "*dead*" forensics model was expanded to include three other methods. They were "*halfdead*", "*resurrected*" and "*Community Live*".

Future work includes providing a complete DFP that includes the isolation explained in this paper. This process can then be used to aid in facilitating a DFI in a Cloud.

References

[1] I.Foster, Y.Zhao, I.Raicu, and S.Lu, "Cloud computing and grid computing 360-degree compared," in *Grid Computing Environments Workshop, 2008. GCE '08*, November 2008, pp. 1 –10.

[2] D.Barrett and T.King, *Computer networking illuminated*, ser. Jones and Bartlett illuminated series. Jones and Bartlett, 2005.

[3] P.Mell and T.Grance, "The NIST Definition of Cloud Computing, Recommendations of the National Institute of Standards and Technolog," National Institute of Standards and Technology, Tech. Rep., 2011.

[4] J.Ashcroft, *Electronic Crime Scene Investigation: A Guide for First Responders*, Technical Working Group for Electronic Crime Scene Investigation, July 2001.

[5] F.Cohen, *Digital Forensic Evidence Examination*, 2nded. Livermore, CA: Fed Cohen & Associates, February 2010.

[6] S.Biggs and S.Vidalis, "Cloud computing: The impact on digital forensic investigations," in *Internet Technology and Secured Transactions, 2009. ICITST 2009. International Conference for*, November 2009, pp. 1 – 6.

[7] R.Chow, P.Golle, M.Jakobsson, E.Shi, J.Staddon, R.Masuoka, and J.Molina, "Controlling data in the cloud: outsourcing computation without outsourcing control," in *Proceedings of the 2009 ACM workshop on Cloud computing security*, ser. CCSW '09. New York, NY, USA: ACM, 2009, pp. 85 – 90. [Online]. Available: http://doi.acm.org/10.1145/1655008.1655020

[8] M.Vouk, "Cloud computing - issues, research and implementations," in *Information Technology Interfaces, 2008. ITI 2008. 30th International Conference on*, June 2008, pp. 31 – 40.

[9] C.Binnig, D.Kossmann, T.Kraska, and S.Loesing, "How is the weather tomorrow?: towards a benchmark for the cloud," in *Proceedings of the Second International Workshop on Testing Database Systems*, ser. DBTest '09. New York, NY, USA: ACM, 2009, pp. 1 – 9. [Online]. Available: http://doi.acm.org/10.1145/1594156.1594168

[10] R.Lu, X.Lin, X.Liang, and X.S. Shen, "Secure provenance: the essential of bread and butter of data forensics in cloud computing," in *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security*, ser. ASIACCS '10. New York, NY, USA: ACM, 2010, pp. 282–292. [Online]. Available: http://doi.acm.org/10.1145/1755688.1755723

[11] Nitu, "Configurability in SaaS (software as a service) applications," in *Proceedings of the 2nd India software engineering conference*, ser. ISEC '09. New York, NY, USA: ACM, 2009, pp. 19 – 26. [Online]. Available: http://0-doi.acm.org/10.1145/1506216.1506221

[12] G.Reese, *Cloud Application Architectures : Building Applications and Infrastrucure in the Cloud*, 1sted., A.Oram, Ed. O'Reilly Media, 2009.

[13] M.M. Khosa, *Infrastructure Mandate for Change 1994-1999*, M.M. Khosa, Ed. Human Sciences Research Council, 2001.

[14] K.Ruan, J.Carthy, T.Kechadi, and M.Crosbie, "Cloud forensics: An overview," *IFIP International Conference on Digital Forensics*, vol.7, 2011.

[15] M.A. Caloyannides, N.Memon, and W.Venema, "Digital forensics," *Security Privacy, IEEE*, vol.7, no.2, pp. 16 – 17, March 2009.

[16] "Definition of Isolate," August 2011, oxford Dictionary. [Online]. Available:http://oxforddictionaries.com/definition/isolate

[17] P.White, *Crime Scene to Court: The Essentials of Forensic Science*, 3rded., P.White, Ed. Royal Society of Chemistry, 2010.

[18] N.Lim and A.Khoo, "Forensics of computers and handheld devices: identical or fraternal twins?" *Commun. ACM*, vol.52, pp. 132 – 135, June 2009. [Online]. Available:http://0-doi.acm.org/10.1145/1516046.1516080

[19] J.R. Lyle, "A strategy for testing hardware write block devices," *Digital Investigation*, vol. 3, Supplement, no.0, pp. 3 – 9, 2006, the Proceedings of the 6th Annual Digital Forensic Research Workshop (DFRWS '06). [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1742287606000624

[20] E.Casalicchio and S.Tucci, "Static and dynamic scheduling algorithms for scalable web server farm," in *Parallel and Distributed Processing, 2001. Proceedings.* Ninth Euromicro Workshop on, 2001, pp. 369 –376.

[21] Nimbula, "Nimbula Director," Computer Program, version 1.0.3. [Online]. Available:http://nimbula.com/products/overview

[22] D.J. Ras and M.S. Olivier, "Finding droplet in the cloud: Finding file fragments in the cloud." Accepted, Eighth Annual IFIP WG 11.9 International Conference on Digital Forensics.

[23] W.Delport, M.S. Olivier, and M.Köhn, "Isolating a cloud instance for a digital forensic investigation." Research in Progress, 10th Annual ISSA Conference.