Cloud Separation: Stuck inside the cloud.

Waldo Delport and Martin S. 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

Abstract. When something erroneous happens happens in digital environment, a Digital Forensic Investigations (DFIs) can be used to gather information about the event. When conducting a DFI, Digital Forensic Procedures (DFPs) are followed. DFPs provide steps to follow to ensure the successful completion of the DFI. One of the steps in a DFP is to isolate possible evidence in order to protect the evidence from contamination and tampering. The introduction of Cloud computing complicated the isolation process because there is a shared layer between users. This means that the methods used to isolate evidence must be adapted and reworked to work in the Cloud environment. In some cases new procedures need to be introduced to address the isolation problem.

In this article we introduce the idea of Cloud separation to isolate a part of the Cloud. We argue that the separation process consists of methods to move instances, as well as methods to divide the Cloud. The paper also introduces methods to accomplish the movement of instances and the division of the Cloud. The paper reports on the finding of testing the dividing methods on different Cloud operating systems in experimental conditions. The experimental outcome was that some of the methods are not applicable to Cloud separation and the methods to be used will depend on the circumstances of the DFI. Out of the experiment some lessons were learnt which should be considered when doing Cloud separation.

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

Published as: W. Delport and M. S. Olivier, "Cloud separation: Stuck inside the cloud" in *Trust, Privacy and Security in Digital Business*, ser. Lecture Notes in Computer Science, S. Fischer-Hbner, S. Katsikas, and G. Quirchmayr, Eds. Springer Berlin Heidelberg, 2012, vol. 7449, pp. 36 - 49.

1 Introduction

Cloud Computing is a fast growing industry and is becoming part of most enterprises [1]. Cloud computing builds on advances in both the network industry and in virtualization [2]. As network infrastructure becomes faster and more reliable, it is also becoming better able to handle large volumes of data, fast and reliably. Virtualization also enables virtual resources to be provided. The process of creating and maintaining virtual resources is being simplified and optimized. Cloud Computing enables a provider to provide virtual resources over the network [4].

When something erroneous happens an investigation may be required. In the Cloud Computing environment the resources are virtual and most interactions with the Cloud are digital in nature [5]. When conducting an investigation on digital artifacts, a Digital Forensic Investigation (DFI) may need to be performed [6]. When doing a DFI, a Digital Forensic Procedure (DFP) is followed [7], which enables admissible evidence to be gathered from the investigation. In the virtual Cloud environment a DFP is followed to conduct an investigation.

In previous work we introduced a Distributed Instance System (DiS) environment, in which multiple instances form a single resource [8]. This is accomplished when multiple instances work together to achieve a common goal. The previous work introduced conditions for isolating single instances to protect the evidence. When working within a DiS environment it is preferable to isolate all the instances at once in order to protect the evidence from tampering and contamination.

In this paper we propose methods to isolate a set of DiS instances. This set of isolated instances can then be used in a Digital Forensic Investigation.

We look into a subset of the proposed methods and provide feedback on them. The results were gathered from empirical experimentation using different Cloud operating systems.

The remainder of the paper is structured as follows, in section 2 cloud computing is explained. Section 3 explains the Digital Forensic Procedure (DFP) followed when doing a DFI. The reasons for Cloud isolation are given in section 4. The methods that can be used for Cloud separation are introduced in section 5. In section 6 considerations are introduced when doing Cloud separation on different Cloud models. Experimental results are reported in section 7.

2 Cloud Computing

Cloud Computing builds on different forms of distributed computing tying together distributed computing and virtualization [1]. Cloud Computing enables a service provider to provide a flexible, cost effective and on-demand infrastructure to its clients, freeing clients from running their own infrastructure. 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 and can be contained within a single node. The Cloud itself consists of multiple nodes. The Cloud can be described by service and deployment models, where the service models describe what service the Cloud offers and the deployment models specify the physical deployment of the Cloud. There are three types of Cloud Computing service models, namely the Infrastructure as a

Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS) models [**?**]. Each of the service models will be explained below.

The first service model 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 [3]. This virtual computer needs to be set up and maintained by the user and can also be referred to as an instance. If the requirements of the user changes 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 with 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 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 model 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.

As mentioned, the Cloud has different deployment models. There are four deployment models for Clouds. They are Public, Private, Hybrid and Community models [5]. In a *Public* Cloud, the infrastructure is owned by a Cloud service provider and 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 exclusive use of one company, therefore 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, forming a community with shared interests. 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.

The *Hybrid* model 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. This is normally used for load balancing.

Cloud Computing is growing and is estimated to become a billion dollar industry this year 2012 [12]. The reason for this is that some of the largest IT related companies have implemented or are implementing Cloud Computing. Some of these large companies are Google, Microsoft, IBM and Amazon [10] [3]. 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 [7] proposes a model for the digital forensic investigation that consists of seven phases, namely the Identification, Collection, Transportation, Storage, Examination and Traces, Presentation and Destruction phases. The *Examination and Traces* phase consists of four subcategories: Analysis, Interpretation, Attribution and Reconstruction [7].

Although not previously mentioned, documentation is a continuous process that needs to take place in all phases of the digital examination [6] [7]. 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 where the evidence is gathered. It 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 alternative DFPs to Cohen's proposed model for digital forensic investigation. The other models include most of these phases or a combination thereof. One such prominent DFP was defined by the National Institute of Justice (NIJ) [6]. The phases defined are Collection, Examination, Analysis and Reporting. The two processes include the same set of underlying steps. Cohen's process is subdivided into more steps. This enables a more systematic flow of events.

The process of isolation forms part of a DFP [8], and is especially important in the collection phase. Examples of isolation methods used in DFPs are when seized cell phones are placed inside a Faraday bag [13] and when doing hard drive forensics on a hard drive, a write blocker is used to enable a write-free read [14]. The isolation helps protect the possible evidence from contamination and loss of continuity.

4 Cloud Isolation

Previous work has been done on isolating single instances [15] [8]. We proposed conditions that we argue need to be met in order to identify instances as successfully isolated. The conditions are, the instance‘s physical location is known, 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 conditions can be expanded to the isolation of a sub-part of the Cloud.

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.

In order to isolate a Cloud we isolate a sub-part of the Cloud. This is done to keep the isolated part of the Cloud in a Cloud environment [15] [8]. In this paper the focus is not on isolating a single instance or a small sub-set of instances but rather a part of the Cloud. This sub-Cloud will have the normal functionality of a Cloud. The instances running on the Cloud will not be aware of the change of Cloud to sub-Cloud. This separation is done to tie together cooperating instances and to exclude unrelated instances. The separation also aids in the admissibility of the evidence. Once the Cloud is separated the DFI is done on the isolated part of the Cloud without any disruption of service to the other clients of the Cloud provider.

5 Cloud Separation

Cloud separation can be argued as a vital part of a DFI on Clouds since, as stated above, the isolation process can aid the admissibility of the evidence. The Cloud separation forms part of the Collection phase of a DFI, the separation is done to prepare the Cloud for an the investigation. We argued that the conditions for isolation as stated in section 4 need to be met in order to state that the separation was done successfully. After careful consideration while creating the condition we discovered the notion of Cloud separation can be separated into moving the instances and dividing the Cloud.

Moving the instance involves relocating the instances from one node to another. This moving of instances should move all the involved instances to a certain part of the Cloud and all non-related instances to another part of the Cloud. The movement is done as a starting point to do the isolation explained in section 4. The movement can be done using one of several methods, the fist option is that the instances can be moved from one Cloud to another directly. The second option is to move the instances to an external Cloud and then from there to the other Cloud. The third option is to move the instance to an external Cloud, then move it two one or more other external Clouds and finally move it to the other Cloud. The fourth option is to use the Cloud operating system to move the instances. The last option is to just identify the nodes which contains suspect instances but we do not move them.

The division of the Cloud is done to complete the isolation. This division can be done in several ways: the first option is to separate the nodes by creating two separate networks from one network, the second option is to create two virtual networks on one logical network and the third option is to create sub-clouds inside the actual Cloud. The last option for Cloud division is using the movement methods to move the instances to a Cloud dedicated for the DFI. The movement and division methods together form the Cloud separation. This means different Cloud movement and division methods can be used together in different combinations depending on the specific requirements. The remainder of this section will expand on the movement and dividing methods.

When using the first option to move instances, from one cloud to the another, it can be done using two different methods. The first method is to mirror an external Cloud then send instances from the main Cloud to the external Cloud. The external Cloud is thus not external but a controlled Cloud that was setup to accept instances. An overview is given in Figure 1. Some Cloud operating systems have the functionality to send instances from their Cloud to another Cloud. This functionality will be used to transfer the instances. It must be known how the instances are sent and what is required while sending it. This will make it possible to mirror one of these external Clouds and receive instances from the main Cloud. The advantage of this method is that instances can see this movement as a normal Cloud operation activity. A second method is when the Cloud operating system allows the sending and receiving of instances. The Cloud operating system is used as an aid in the movement of the instances. An example of a Cloud operating system that can send and receive instances is VMWare [16].

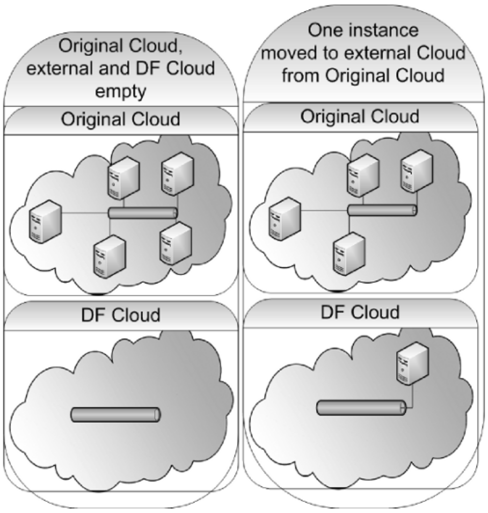
Original Cloud, external and DF Cloud empty
Original Cloud
DF Cloud
One instance moved to external Cloud from Original Cloud
Original Cloud
DF Cloud

Fig. 1. Moving an instance from one Cloud to another

When using an external Cloud in the process of moving instances, the same methods as suggested above, can be used. The instances are sent to an external Cloud. This external Cloud can accept instances from the main cloud and can be assessed by the DFI team. Once the instance is on the external Cloud it is sent to the controlled Cloud. Cloud operating systems like Nimbula and VMWare can send instances to external Clouds. This external Cloud can be hosted by other companies or be another Cloud owned by a company. Figure 2 explains the steps. Automated methods can be used to move these instances. In the case where there are no automated methods, one of the methods proposed to move an instance can be used [15].

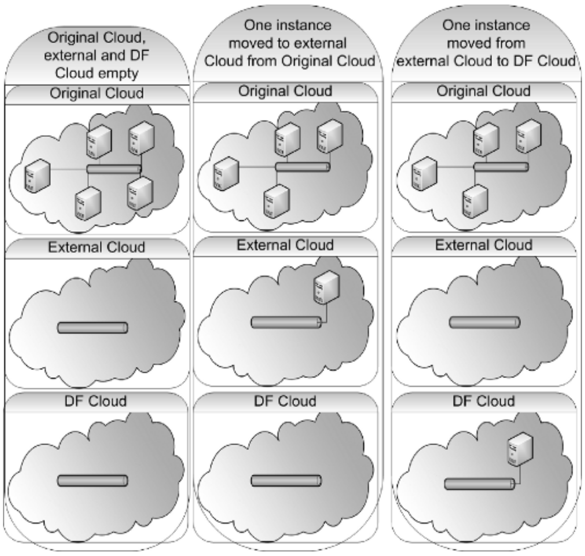
Original Cloud, external and DF Cloud empty
One instance moved to external Cloud from Original Cloud
One instance moved from external Cloud to DF Cloud
Original Cloud
External Cloud
DF Cloud

Fig. 2. Moving an instance from one Cloud to another using an external Cloud

The option where multiple external Clouds are used is the same as the above but there exist multiple external Clouds between the two Clouds. This method can be employed when no middle ground exists between the main Cloud and the other Cloud. The external Clouds are used to link the two Clouds.

The Cloud operating system can also be used to move instances. Some Cloud operating systems provide the functionality to migrate instances while they are running between nodes. The last option where the nodes are only identified will be used if there are no methods available to move the instances, or if there are only suspect instances on the node.

The first option when dividing the Cloud is using the self-healing characteristic of Clouds that will be used to create two Clouds. If a node or nodes malfunction in Nimbula the Cloud itself will continue to operate. In this option the first step is to identify the nodes that need to form part of the new Cloud. The second step is to move all non-related nodes from these Clouds. The next step will be explained by means of an example: if the Cloud has six nodes and three of them need to move to the new Cloud, the process is as follows: Connect two switches to each other, the one has all the nodes connected to it. Systematically move the suspect nodes' network wire/VLAN one by one to the other switch. Once all the suspect nodes are connected to the other switch the connection between the two switches is broken. Then the Cloud operating system will create a new Cloud using the self-healing ability. The process is illustrated in figure 3.

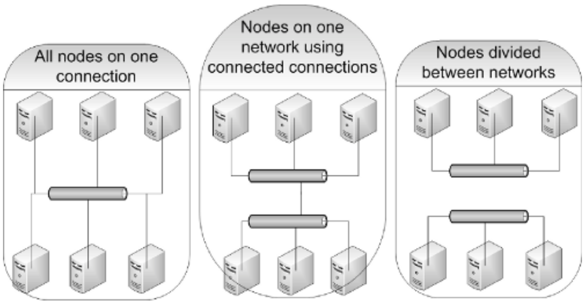
All nodes on one connection
Nodes on one network using connected connections
Nodes divided between networks

Fig. 3. Creating two Clouds from a single Cloud

The second option for Cloud dividing is to create two Clouds on one network. A high level overview is given in figure 4. This category can be separated on the notion of knowing which node belongs to which Cloud or not knowing which node belongs to which Cloud. Each Cloud runs its own Cloud operating system that will control it. There are different methods to create two Clouds on the same network. One option is to create separate Subnet masks for each instance of the Cloud [2]. This will enable the installation and operation of each Cloud on a separate Subnet mask. A possible alternative option is to use a Cloud operating system that enables the selection of the controlling node, and by using this strategy, a new master is set up on the network and some nodes are allocated to it.

The third option to divide the Cloud, is to create sub-Clouds. The Cloud is logically broken up into separate parts. The same Cloud operating system controls them. The sub-Cloud is a fully functional Cloud and it just runs on the main Cloud but is interacted with as if it is a normal Cloud. Some Cloud operating systems allow for the creation of sub-clouds inside the Cloud itself. It is used to sell a Cloud to the service provider's customers. To do this a sub-cloud is created on the main Cloud and then instances are moved to the sub-cloud. This sub-Clouds is implemented on the same hardware as the base Cloud. The moving functionality is provided by the Cloud operating system. Figure 5 shows the main Cloud's hardware and the virtual Clouds created on that hardware.

The last option to divide the Cloud is to use any of the instance movement methods to move instances to an already divided Cloud. This Cloud can be a

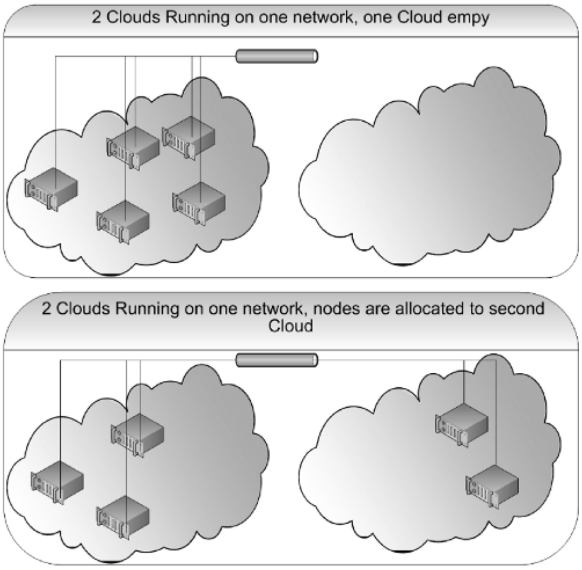
2 Clouds Running on one network, one Cloud empy
2 Clouds Running on one network, nodes are allocated to second Cloud

Fig. 4. Creating two Clouds on one network

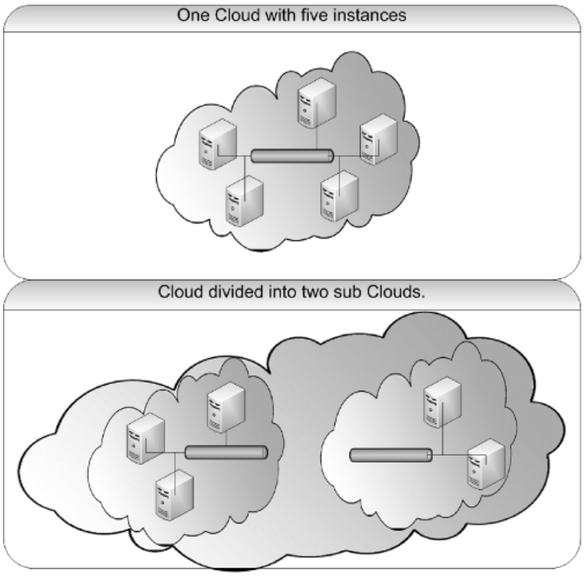
One Cloud with five instances
Cloud divided into two sub Clouds.

Fig. 5. Creating two sub-Clouds

Cloud prep to do Cloud forensics. The Cloud can also be located on the Cloud providers premises or an external Cloud on the DFI teams premises.

While doing the separation all steps must be documented as part of the DFI, this created a audit trail which can be used to prove the viability of the methods followed.

6 Cloud Separation on different types of Clouds

In the previous section we introduced methods for Cloud separation. As stated in section 2 Clouds can be divided into different service and deployment models, which have different impacts on isolation. Some considerations need to be taken into account when doing Cloud separation for the different models. The important consideration is the Confidentiality, Integrity and Availability (CIA) of instances.

The difference between service models is in who owns what part of the instance. The instance can usually be divided into the hardware, the hypervisor, the operating system, applications and data. In an IaaS model, the service provider is responsible for the hardware and hypervisor whereas the client is responsible for the rest. If the service provider is requesting the DFI it must get the cooperation of the client to gather evidence from the operating system, applications and other data residing on the system. If the client is requesting a DFI they must get the cooperation from the service provider in gathering evidence form the hardware. The client and service provider are both responsible for the availability of the system. It is easily possible for clients to have multiple instances working together without the knowledge of the service provider.

In a PaaS model the service provider is responsible for the hardware, hypervisor, operating system and some applications. The client is responsible for applications and data on the system. The service provider must ensure that high availability is maintained. The client can provide evidence from their own applications and stored data. It is possible, but more unlikely than in IaaS, to have cooperating instances.

When doing a DFI on a SaaS the service provider is responsible for the hardware, hypervisor, operating system and applications. The client is responsible for configurations of applications and data on the system. The service provider must ensure high availability of the systems. Clients are only responsible for their data. It is very unlikely for a client to have cooperating instances.

For the purpose of this paper we will only look at public and private development models, and we argue that hybrid and community development models have the same considerations as public and private development models. When doing Cloud separation on a public development model the Cloud service provider is responsible for protecting the CIA of their clients. When separating a part of the Cloud it must be confirmed that only data related to suspect instances are separated. The separation must also protect the admissibility of the evidence. All unrelated instances should not be affected by the separation and must thus stay available. If the service provider is doing the DFI the provider must protect the privacy of its clients and inform its clients of the investigation. If an external company is doing the DFI the company must protect the privacy of the service provider and its clients.

When doing Cloud separation on a private development model all data should belong to one company. The separation is done to protect the admissibility of the evidence. If the owners of the Cloud are responsible for conducting the investigation, the main focus is not on protecting the privacy of the information. If an external company is responsible for doing the investigation the separation should also protect the privacy of the owner's data. The owner is responsible for deciding the importance of the availability of the Cloud.

It can be argued that Cloud separation is valid for IaaS and PaaS models. Cloud separation can be an integral part of a DFI on a public Cloud but can also be important in a DFI on a private Cloud.

7 Experimentation

In this section our experimentation results are given. The experiment was limited to the dividing methods, moving an instance can be done by the Cloud operating system making it part of normal Cloud operation or if there is no functionality by using one of the methods proposed in the paper by Delport et al [15]. In the experiment we tested two dividing methods, the methods were: creating two Clouds using the network hardware and creating sub-Clouds. The experiment used VMware and Nimbula Director. This was done to get some comparison between the methods' feasibility. The reasons why VMware and Nimbule were chosen is that VMware is a widely used platform to provide Cloud resources and Nimbula focuses on providing private Cloud infrastructure. This gives us better coverage for both public and private Cloud computing.

In order to create sub-Clouds one needs more than one layer of abstraction. In the experiment VMware was used to create the sub-Clouds. There were two base nodes running VMware, which are known as ESXi hosts. These nodes have Intel i5 processors and 2GB DDR3 memory. On each of the hosts two other ESXi hosts were created. The virtualized ESXi hosts were used to form a Cloud on each main host. A vCenter management instance was created on each virtual Cloud. VCenter is used to control the Cloud [16]. Two instances running CentOS 6.0 minimal were also created on the virtual Cloud. The layout then is vCenter running on Windows 2008 server and two CentOS minimal instances running on two ESXi hosts. The ESXi hosts are running as virtual machines on a base ESXi host.

The setup and maintenance of this experiment was relatively easy. The Clouds where stable and there were no apparent problems with the recursive virtualization. In the testing environment there was some loss of performance: this occurred because some resources are used to run the other virtual hosts and another reason is that there are two controlling layers.

To test the performance loss a single sub-Cloud was created on a Dell PowerEdge R710 with two i7 processors and 97GB of memory. On this node the performance decrease was not noticeable. The performance drawback should not be noticeable on most of the powerful infrastructure used by most Cloud Providers. The instances might notice the loss in performance on the node and might start self defense mechanisms, while this can be done the performance on the Cloud environment is inherently unstable because of resource over committing that is part of most Cloud environment [16].

VMware also aids in creating sub-Clouds. VMWare allows the movement of instances from the main Cloud to the virtual Cloud and from the virtual Cloud to the main Cloud. The user must link the virtual Cloud's vCenter to the underlying infrastructure. This allows the user to move instances between the layers of virtualization. The drawback is that there is a connection created from the virtual Cloud to the underlying Cloud. This can be used to tamper with evidence. The advantage is that a virtual Cloud can be created at a later stage and instances moved to it from the main Cloud. Once a digital investigation is required the instances can be moved to sub-Clouds, one for uninvolved instances and one for suspected instances. If instances are no longer suspected in the suspect Cloud it can be moved to the other Cloud.

An experiment using Nimbula director was conducted to create sub-Clouds. The Cloud consisted of three nodes, on each two virtual nodes were created. It was possible to access each of the sub-Clouds separately. The problem was that the virtual node should be on a virtual network. If they are not on a virtual network they cannot communicate with the other virtual nodes. It was possible to create instances on the virtual nodes. Because of a limitation in Nimbula instances can not be moved from the main Cloud to the sub-Cloud directly. The movement methods proposed in previous work needs to be used to move the instances to the sub-Cloud.

The next experiment created two Clouds using network hardware. This experiment was done using Nimbula director and using six nodes with 2GB of RAM and i5 processors. The experiment was done as described in section 4. Access was lost to the Control centre of Nimbula on the one part but the instances were still running. A possible problem is that the control centre holds information about all running instances. If the Cloud is broken up the control centre loses communication with the other instances that are running on the other part of the Cloud. They will show as being in an error state. The instances can then be "deleted" from the control centre as they are not applicable to it. The problem continues because the Clouds cannot be joined later. There are two control centres running each with its one instance. In the experiments' experimental conditions it seemed impossible to join the Cloud back together. Although connection was lost with the control centre the Cloud still functioned proving that the self-healing characteristics of Nimbula are intact.

The last experiment was done using VMware to create two Clouds using the network hardware. The same procedures were followed as for Nimbula. The experiment was successful although a few problems occurred and configuration changes were needed. The problems were in vCentre assuming that host failure occurred and it tried to relaunch the lost instances. This happened because high availability was enabled on the cluster, the job of HA is to recover lost instances.

It failed because the instance storage was on the direct attached storage. On the other part of the Cloud a new vCentre needed to be created because there was no management over the new cluster.

	vmWare	Nimbula	SAN	DAS
Cloud seperation using sub-Clouds	✓	X	✓	✓
Cloud separation using Network Hardware	✓	X	X	✓

Table 1. Experiment Summary.

From the experimentation it can be seen that the method where a sub-Cloud is created using network hardware is then not advisable as it would require a lot of re-setup to put the Cloud together again, it is advised against the use of this method. The other experiment shows it is more reliable to have sub-clouds for cloud separation. Table 1 contains a summary of the experimentation.

From the experiment the following lessons were learnt: Performance is affected on less powerful Clouds, HA needs to be turned off before starting with Cloud separation, recombining the Cloud after the DFI can be hard to impossible.

An overall possible problem that must be considered with all methods for Cloud separation is where the instances storage is located. As a basic example the storage can either be on a SAN or the DAS. Creating sub-Clouds when using a SAN is not possible as connection to the SAN may also be lost. Creating a sub-Cloud can still be done when using a SAN because the nodes can still communicate with the SAN. Both methods are applicable when using DAS. Another problem with SAN's is that multiple instances share the resource, this can be avoided by using a SAN dedicated for the storage of suspect instances. Another problem is the IP address of the instance.

When moving an Instance the IP of that instance should be constant to correlate the IP with gathered network evidence. In the experiments the instances had static IP's which did not change if the instances moved. If a dedicated firewall is used to assign the IP the IP should stay the same if he instances moves. When the IP of the instance is manage by the node it's residing on the IP might change if the instances is moved to aid in correlation of evidence the IP before and after the move must be noted.

8 Conclusion

As Cloud computing grows it will become easier for individuals to create DiS resources. If the DiS resource is used in a form of a crime, methods must exist to start a DFI on the DiS without disruption the other users of the Cloud.

In this paper we introduced the notion of Cloud separation, which consists of moving instances and dividing the Cloud. We explained methods to move instances around in the Cloud as-well as moving instances out of the Cloud. We also explained the methods that can be used to divide the Cloud.

We conducted experimentation on the division methods and discovered that the methods used will depend on the circumstances of the DFI. We saw that the method that uses the network hardware to create two Clouds might not be desirable to use and the method to create sub-Clouds might be a valid choice.

Future work includes testing the methods on more Cloud operating systems to better test all the methods and discover some pitfalls. If we discover that the methods do not work on all platforms we plan to find other methods that will work on specific platform. There is also a a need to investigate the performance loss when conducting a DFI.

References

1. Vouk, M.A.: Cloud computing - issues, research and implementations. In *Information Technology Interfaces, 2008. ITI 2008. 30th International Conference on*, Pages 31 – 40, June 2008.
2. Barrett, D., King, T.: *Computer networking illuminated.* Jones and Bartlett illuminated series. Jones and Bartlett, 2005.
3. Biggs, S., Vidalis, S., Cloud computing: The impact on digital forensic investigations. In *Internet Technology and Secured Transactions, 2009. ICITST 2009. International Conference for*, Pages 1 – 6, November 2009.
4. Foster, I., Zhao, Y., Raicu, I., Lu, S.: Cloud computing and grid computing 360-degree compared. In *Grid Computing Environments Workshop, 2008. GCE '08*, Pages 1 –10, November 2008.
5. Mell, P.,Grance, T.: The NIST Definition of Cloud Computing, Recommendations of the National Institute of Standards and Technolog. Technical report, National Institute of Standards and Technology, 2011.
6. Ashcroft, J.: *Electronic Crime Scene Investigation: A Guide for First Responders.* Technical Working Group for Electronic Crime Scene Investigation, July 2001.
7. Cohen, F.: *Digital Forensic Evidence Examination.* Fed Cohen & Associates, Livermore, CA, 2 edition, February 2010.
8. Delport, W, Olivier, M.S.: Isolation, stuck inside the cloud. (In Press) *Eighth Annual IFIP WG 11.9 International Conference on Digital Forensics*, 2012.
9. Binnig, C., Kossmann, D., Kraska, T., Loesing, S.: How is the weather tomorrow?: towards a benchmark for the cloud. In *Proceedings of the Second International Workshop on Testing Database Systems*, DBTest '09, Pages 1 – 9, New York, NY, USA, 2009. ACM.
10. Lu, R., Lin, X., Liangand, X., Shen X.: 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*, ASIACCS '10, Pages 282–292, New York, NY, USA, 2010. ACM.
11. Nitu., I: Configurability in SaaS (software as a service) applications. In *Proceedings of the 2nd India software engineering conference*, ISEC '09, Pages 19 – 26, New York, NY, USA, 2009. ACM.
12. Ruan, K., Carthy, J., Kechadi, T., Crosbie, M.: Cloud forensics: An overview. *IFIP International Conference on Digital Forensics*, 7, 2011.
13. Lim, N., Khoo, A.: Forensics of computers and handheld devices: identical or fraternal twins? *Commun. ACM 52*, Pages 132 – 135, June 2009.

14. Lyle, J.R.: A strategy for testing hardware write block devices. *Digital Investigation*, 3, Supplement(0):3 – 9, 2006. The Proceedings of the 6th Annual Digital Forensic Research Workshop (DFRWS '06).
15. Delport, W, Olivier, M.S., Köhn, M: Isolating a cloud instance for a digital forensic investigation. In *2011 Information Security for South Africa (ISSA 2011) Conference*, 2011.
16. Vmware inc. Computer Program, 2011. vSphere 5.0. Available online: http://www.vmware.com (Accessed 26 May 2012)