UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA

COMBATING DATA LEAKAGE IN THE CLOUD

by

MOSES THANDOKUHLE DLAMINI

Submitted in fulfilment of the requirements for the degree

Doctor of Philosophy (COMPUTER SCIENCE)

in the

FACULTY OF ENGINEERING, BUILT ENVIRONMENT AND INFORMATION TECHNOLOGY

at the

UNIVERSITY OF PRETORIA

Supervisors: Prof. H.S. Venter and Prof. J.H.P. Eloff

April 2020

ABSTRACT

COMBATING DATA LEAKAGE IN THE CLOUD

by

Moses Thandokuhle Dlamini

Supervisors:	Prof. H.S. Venter and Prof. J.H.P. Eloff
Department:	Computer Science
University:	University of Pretoria
Degree:	Doctor of Philosophy (Computer Science)
Keywords:	Authentication, Cloud Computing, Conflict-based VM Placement, Data Leakage, Digital Forensics, Digital Forensic Readiness, Information Security, Multifactor Authentication, Risk-based Multifactor Authentication, Security of Cloud Computing, Virtual Machine, Virtual Machine Placement

The increasing number of reports on data leakage incidents increasingly erodes the already low consumer confidence in cloud services. Hence, some organisations are still hesitant to fully trust the cloud with their confidential data. Therefore, this study raises a critical and challenging research question*: How can we restore the damaged consumer confidence and improve the uptake and security of cloud services?* This study makes a plausible attempt at unpacking and answering the research question in order to holistically address the data leakage problem from three fronts, i.e. conflict-aware virtual machine (VM) placement, strong authentication and digital forensic readiness. Consequently, this study investigates, designs and develops an innovative conceptual architecture that integrates conflict-aware VM placement, cutting-edge authentication and digital forensic readiness to strengthen cloud security and address the data leakage problem in the hope of eventually restoring consumer confidence in cloud services.

The study proposes and presents a conflict-aware VM placement model. This model uses varying degrees of conflict tolerance levels, the construct of sphere of conflict and sphere of non-conflict. These are used to provide the physical separation of VMs belonging to conflicting tenants that share the same cloud infrastructure. The model assists the cloud service provider to make informed VM placement decisions that factor in their tenants' security profile and balance it against the relevant cost constraints and risk appetite.

The study also proposes and presents a strong risk-based multi-factor authentication mechanism that scales up and down, based on threat levels or risks posed on the system. This ensures that users are authenticated using the right combination of access credentials according to the risk they pose. This also ensures end-to-end security of authentication data, both at rest and in transit, using an innovative cryptography system and steganography.

Furthermore, the study proposes and presents a three-tier digital forensic process model that proactively collects and preserves digital evidence in anticipation of a legal lawsuit or policy breach investigation. This model aims to reduce the time it takes to conduct an investigation in the cloud. Moreover, the three-tier digital forensic readiness process model collects all user activity in a forensically sound manner and notifies investigators of potential security incidents before they occur.

The current study also evaluates the effectiveness and efficiency of the proposed solution in addressing the data leakage problem. The results of the conflict-aware VM placement model are derived from simulated and real cloud environments. In both cases, the results show that the conflict-aware VM placement model is well suited to provide the necessary physical isolation of VM instances that belong to conflicting tenants in order to prevent data leakage threats. However, this comes with a performance cost in the sense that higher conflict tolerance levels on bigger VMs take more time to be placed, compared to smaller VM instances with low conflict tolerance levels. From the risk-based multifactor authentication point of view, the results reflect that the proposed solution is effective and to a certain extent also efficient in preventing unauthorised users, armed with legitimate credentials, from gaining access to systems that they are not authorised to access. The results also demonstrate the uniqueness of the approach in that even minor deviations from the norm are correctly classified as anomalies. Lastly, the results reflect that the proposed 3-tier digital forensic readiness process model is effective in the collection and storage of potential digital evidence. This is done in a forensically sound manner and stands to significantly improve the turnaround time of a digital forensic investigation process. Although the classification of incidents may not be perfect, this can be improved with time and is considered part of the future work suggested by the researcher.

ACKNOWLEDGEMENTS

Writing this thesis has been a long and winding journey, yet fascinating and rewarding at the same time. I would not have been able to finish it without the support of my supervisors and others. Therefore, it is gives me great pleasure to pass a few words of thanks to some of the people who have helped me in so many different ways to finish the race.

First of all, I am grateful to the CREATOR of all created things for giving me good health, inspirational study leaders, a loving family and friends, and the courage to start and finish this tough journey.

My very special thanks goes to my main study leader, Professor H.S. Venter for his belief in my research capabilities. Thank you for your support, guidance, mentorship and encouragement, without which I would not have been able to deliver this thesis. Without your direction and support this would have been a very lonely journey.

A special word of thanks goes to my co-study leader, Professor J.H.P. Eloff, for all the support he has given me over all these years. I thank you for believing in me even when I doubted myself. You never gave up on me, but always challenged me to do my best in everything I do. Your thought-provoking questions, invaluable guidance and support have made it easier for me to keep going, even when the road seemed so difficult. You really brought the best out of me. I am forever grateful to you.

A special word of thanks also goes to Prof. M.M. Eloff. I thank you for your insightful reviews and comments.

Finally, I would like to thank and dedicate this thesis to my family (Nosihle, Yenzo, Fezo and Wenzo) for their support, encouragement and understanding over the entire duration of this study.

TABLE OF CONTENTS

LIST OF ABBREVIATIONS

BIaaS	Business Intelligence as a Service
CAPEX	Capital Expenditure
CBVMP	Conflict-based Virtual Machine Placement
Clu	Cluster
CTL	Conflict Tolerance Level
DC	Data Centre
DFR	Digital Forensic Readiness
DFRWS	Digital Forensic Research Workshop
EC2	Elastic Compute Cloud
GPS	Global Positioning System
IaaS	Infrastructure as a Service
ICT	Information Communication Technology
IDE	Integrated Development Environment
IDMaaS	Identity Management as a Service
Loc	Geo-location
MFA	Multifactor Authentication
NIST	National Institute of Standards and Technology
OOB	Out Of Band
OPEX	Operational Expenditure
OTP	One Time Password
PaaS	Platform as a Service
PN	Physical Node

ROI	Return On Investment
S3	Simple Storage Service
SaaS	Software as a Service
Sec-aaS	Security as a Service
VM	Virtual Machine
VMM	Virtual Machine Monitor
XaaS	Anything as a Service

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1 INTRODUCTION

1.1 INTRODUCTION

The ever-increasing popularity of cloud computing has seen many organisations widely pondering about migrating their platforms and critical business data from on-premise systems to cloud-hosted services. Some researchers (Fahideh and Beydoun, 2018; Khan and Al-Yasiri, 2018; IDG Research, 2016; Miteva, 2018) report that cloud computing is fast reaching its initially anticipated adoption rates. For example, Miteva (2018) reports that 60% of organisations are reported to have at least one cloud-hosted application. This is so that they can access their applications and/or data from anywhere, at any time and using any Internet-enabled device. While this thought is often motivated by numerous cloud benefits such as low capital expenditure, predictable operational costs, low total cost of ownership, economies of scale, enhanced scalability, flexibility, dynamic provision of IT resources and better disaster preparedness, it also poses new security risks (Khan and Al-Yasiri, 2018; Ma, 2015; IDG Research, 2016; Dlamini et al., 2011).

A key security issue that comes as a result of migration to cloud services is the serious and subtle data leakage threat that seems to be on the increase year on year. The Gemalto and Ponemon Institute reports that in 2018, out of about 63% organisations who share their sensitive corporate data in the cloud with third parties, 10% still do not implement any protection measures (Gemalto and Ponemon Institute, 2018). This means that such organisations leave the security of their cloud-hosted confidential data in the hands of the cloud service provider. Consequently, there has been an increase in the number of data leakage threats year after year. Furthermore, reports suggest that data leakage threats are three times more likely to occur in the cloud than in on-premise systems (Ponemon Institute, 2014; Ma, 2015). The likelihood of a data leakage or breach somehow increases in the cloud. Numerous corporates have suffered the brunt of high-profile data leakage incidents in the cloud, such as JP Morgan Chase, Dropbox, LinkedIn, Yahoo, Target, Ebay and River City Media (Quick et al., 2017). The latest incidents affected organisations like CloudFlare, Interpark, Dailymotion, Snapchat, Equifax and others (Quick et al., 2017; EquiFax, 2018). However, these are only the reported incidents. There are probably many more data leakage cases that never get reported and remain unknown to the public or consumers of cloud services. The indication is that trillions of data records have already been compromised through data leakages and breaches (Cheng, 2017). The unfortunate part is that this includes sensitive data like encryption keys, user credentials, and authentication tokens in plain text (CSA, 2017).

Of particular interest to this study is the incident involving Dropbox that happened back in 2011 (cited in Mariani, Pezze and Zuddas, 2018). Dropbox reported an authentication bug. The bug allowed a number of unauthorised and authorised clients to log into Dropbox accounts with only the email address and any password of their choice for about four hours (Ferdowsi, 2011). During the estimated four hours of free access, 25 million Dropbox accounts, along with clients' cloud-hosted data, were widely exposed. Dropbox claims that only about one percent of the 25 million accounts (i.e. approximately 250 000 accounts were active during the four hours of free access) could have been compromised by this vulnerability (Mariani et al., 2018; Ferdowsi, 2011). Despite the seemingly small number of allegedly affected accounts, and even if it were only for a minute; any undue exposure of clients' cloud-hosted data to unauthorised third parties is totally unacceptable. A similar case is that of LinkedIn, reported in CSA in 2017. In this particular case, attackers compromised 1,5 billion encrypted user accounts over a space of three years (between 2014 and 2016) (CSA, 2017). Though these had been encrypted by LinkedIn, the attackers were still able to decrypt and post them in a malicious website. The decryption of the user credentials was made possible by LinkedIn's failure to use a good one-way hash algorithm.

In order to truly understand the impact of such incidents, one needs to consider the amplified damage and ripple effects that they may have in a global company that segregates and shares their sensitive corporate data with thousands of partners all over the world using a Dropbox account. In this particular scenario, a disgruntled partner – equipped with just the email address of the global company – could gain unlimited access to sensitive data without proper authentication. Incidents like this seriously hamper the significant efforts already made towards increasing the adoption of cloud services. They are the main reason why there is still a low consumer confidence in cloud services. There is also the fear that cloud service providers do not have the necessary security technologies to thwart data leakage threats. Furthermore, up to about a year ago, in most countries of the world (save for the US), cloud service providers were not subjected to data leakage/breach disclosure laws (Stevens, 2012). However, since the implementation of the European Union's General Data Protection Regulation (GDPR) which came into effect in May 2018, this has been changing. Cloud service providers – at least those in the US and those that process personally identifiable information (PII) from any of the European Union countries – are now obligated in terms of the GDPR's Chapter 4 article number 33 to notify their customers if their data has been breached (GDPR, 2016; Gemalto and Ponemon, 2018; Park et al., 2018). However, some organisations are still hesitant to move to the cloud, despite the hype and numerous benefits that lie therein.

The data leakage threat seems to be intensifying with the rise and frequent use of cloud computing. The threat is particularly exacerbated by the public cloud's resource-sharing capability (Almutairi et al., 2012) and it capitalises on the very important multi-tenancy and virtualisation features of public cloud computing, which makes sharing of cloud resources possible. The multi-tenancy feature allows cloud clients to store their confidential corporate data on multiple disjoint virtual machines (VM) that may technically be placed on shared physical hardware next to that of their competitors or adversaries (Ristenpart et al., 2009; Zhang et al., 2012; Wang et al., 2018). Moreover, public cloud computing provides a single point of entry (i.e. cloud service provider) to multiple clients' data, which are only logically separated by a hypervisor. A flaw in the services provided by the cloud service provider could easily compromise every tenant's critical data and wrongly put it into the hands of adversaries or competitors as showed in the Dropbox case.

The foundational work of Ristenpart et al. (2009) (cited in Wang et al. (2018)) has empirically shown that it is possible to locate another client's VM on the underlying cloud service provider's physical infrastructure. The work by Ristenpart et al. (2009) also shows that an adversary can have as much as a 40% chance to strategically place their malicious VM on the same physical infrastructure as that of their target clients. Once an adversary or competitor does this, all it takes for them would be to penetrate the hypervisor isolation between the VMs. This allows them to gain unlimited access and to manipulate and extract confidential data belonging to other co-resident clients. Furthermore, given the sharing of resources, it also becomes hard to separate access logs per user, as any evidence source from such a setup will always contain residual data of other co-resident tenants. This thesis therefore argues that multi-tenancy and virtualisation as core features of public clouds present unique challenges with regard to security (e.g. authentication and access controls) and digital forensics. This is mainly due to the fact that these features facilitate cloud resource sharing among potentially untrusted tenants, which results in an increased risk of data leakage. The remainder of this section discusses the research problem.

Unwary organisations may significantly increase their risk of data leakage incidents if they blindly adopt and make use of cloud services to store and process critical business data (Filkins et al., 2016; Symantec, 2013; Pearson and Charlesworth, 2009). Many organisations are still hesitant to make a move to the cloud, mainly because of concerns around confidential data leakage as reflected in the Dropbox case. The message is clear from surveys conducted by IDG Research Services, IDC Research (Axway, 2013), the Ponemon Institute (2014) and Ma (2015).

The surveys (Axway, 2013; Ponemon Institute, 2014; Ma, 2015 and Panko, 2017) have shown that there is a lack of consumer confidence in cloud services. For example, Panko (2017) reports that only 55% of their respondents are confident in their knowledge of cloud services and the security issues thereof. In some of the studies (Ponemon Institute, 2014; Ma, 2015 and Axway, 2013), the lack of consumer confidence in cloud services is directly related to the numerous security concerns related to securing cloud services. These are listed below – in no particular order (Axway, 2013; Ponemon Institute, 2014; Ma, 2015 and Panko, 2017).

- Data leakage – exposure of confidential data
- Inadequate authentication and inappropriate access controls that lead to unauthorised access
- Lack of availability
- Vendor lock-in
- Loss of control over data
- Lack of monitoring activities on shared resources – e-discovery of digital evidence
- Trust and privacy issues
- Legal and regulatory compliance issues – liability and accountability issues
- Business continuity and disaster recovery issues

The above issues are not necessarily an exhaustive list of all the security concerns that might need to be addressed in the cloud, as they also cut across all the information security services (i.e. confidentiality, integrity and availability). Axway (2013) and Ma (2015) prioritised and ordered the concerns of cloud services. Therein, it is reported that the top concerns with regard to cloud services include the following: the leakage of sensitive data resources; data loss or theft; loss of control over data; and unauthorised access and usage. The complete list is as shown in Figure 1.1:

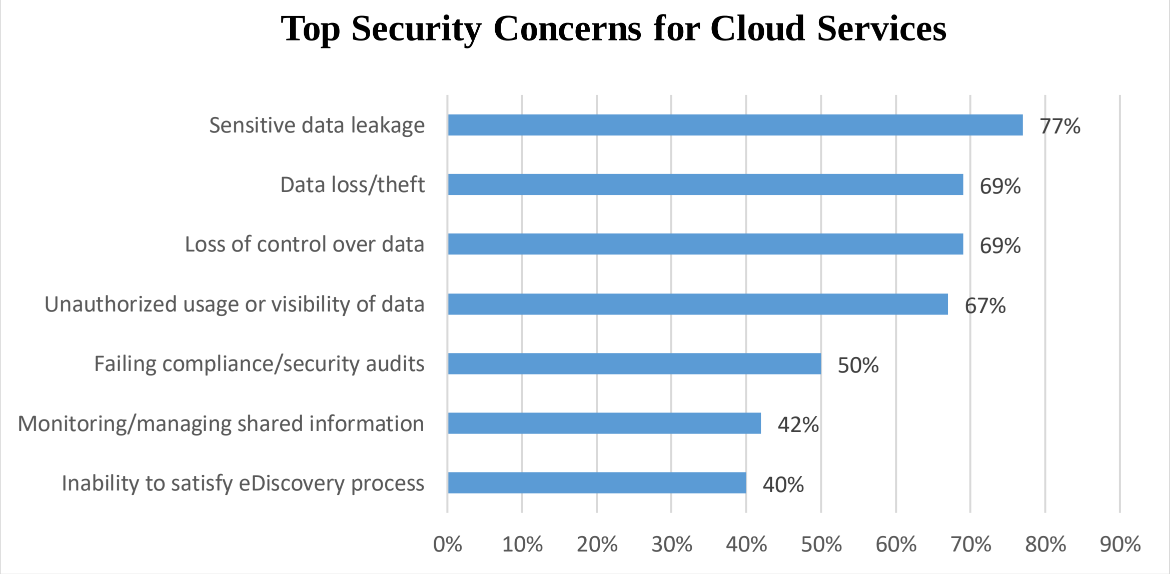

Figure 1.1:- Top Security Concerns for Cloud Services (Axway, 2013; Ma, 2015)

Moreover, the Top Threats Working Group of Cloud Security Alliance (CSA) listed the treacherous top twelve cloud computing threats as follows (CSA, February 2016; 2017):

- Data leakage threats
- Insufficient identity and access management
- Insecure APIs and shared technology issues
- System and application vulnerabilities
- Account hijack
- Malicious insiders that abuse cloud services
- Denial of service

These security concerns are somewhat similar to the list provided by Hashizume et al. (2013); Ali, Memon and Sahito (2018); Singh (2014) and Ma (2015). For example, Hashizume et al. (2013) place account hijack at the top of the list, followed by data leakage, denial of services and others. Ma (2015) places data leakage at the top of their top ten list of concerns for cloud computing. Ali et al. (2018) and Singh (2014) assume the same listing as that of CSA (2016, 2017) and places data leakage at the top of the list. This is an indication that data leakage threats are indeed a major concern for potential cloud service users looking to exploit the benefits of cloud computing. Consequently, this threat will have to be addressed before we can see a mass adoption of cloud services.

1.2 PROBLEM STATEMENT

The key problem that this research aims to address is the data leakage concerns in cloud computing infrastructures. The data leakage threat has featured in the top ten lists for the past six years now, beginning from 2012 to 2017. Hence, this thesis places strong emphasis on addressing data leakage threats and focuses specifically on cloud computing. The next section takes a look at the research question that this study aims to answer.

1.3 RESEARCH QUESTION

The challenge now is, in spite of the numerous security concerns reflected above, ***how can we restore the damaged consumer confidence and improve the uptake and security of cloud services?*** Ideally, we should address all of the challenges in the previous section. However, this would be an impractical and insurmountable task and be unable to achieve in one research effort. Based on this reason and the security concerns raised in Axway (2013), Ponemon Institute (2014), Ma (2015) and CSA (2016; 2017), this research focuses only on solving the problem of the leakage of confidential and/or sensitive data to unauthorised and/or unintended third parties in the cloud. This thesis addresses the problem from three dimensions, i.e. VM placement (Dlamini, Eloff and Eloff, 2014), authentication (Dlamini et al., 2012; Dlamini et al., 2015) and digital forensic readiness (Dlamini et al., 2014). Solving the problem on these three dimensions would directly or indirectly address some of the other concerns too. For example, real-time monitoring of the cloud-hosted resources based on e-discovery (Hook, 2018) would restore the lost control over clients' data. This would also provide auditors and compliance authorities with a clear view of where tenants' data sits at any particular moment.

In order to answer the main research question as stated above, we need to first answer some subquestions that originate from the main research question, as they could help to address the main challenge.

1.3.1 Subquestions

The subquestions are as follows:

- *How can we improve and provide VM placement that prevents the co-location of conflicting VMs on cloud computing?*
- *How can we secure cloud-based resources in a manner that prevents their leakage to unintended parties? In other words, how can we improve and provide appropriate authentication that would be suitable for cloud computing?*
- *How can we prepare the cloud to become ready for e-discovery of digital evidence?*

The study in hand answers each of these questions as attempts to prevent the leakage of cloud-hosted data to unintended and/or unauthorised third parties in the cloud. The next subsection discusses the study's research objectives.

1.4 RESEARCH OBJECTIVES

By answering each of the above research questions and subquestions, the study aims to tackle the following research objectives:

- Objective 1: Critically analyse current cloud security trends with a specific focus on VM placement, authentication and digital forensic readiness.
- Objective 2: Investigate, design and develop an innovative architecture that integrates conflict-aware VM placement, cutting-edge authentication and digital forensic readiness to address data leakage threats in the cloud.
- Objective 3: Implement and evaluate the proposed solution (an innovative model) as a proof-of-concept on a real cloud platform to demonstrate its practicability and suitability to prevent data leakage threats.

This thesis aims to demonstrate and prove that VM placement, traditional authentication and digital forensic readiness can be improved and seamlessly integrated to help prevent data leakage threats in the cloud. In the quest to tackle the above research objectives, this study is limited in scope to cover only specific aspects of cloud computing. The next section discusses the scope and context of this research.

1.5 RESEARCH SCOPE AND CONTEXT

The current research focuses on cloud security and emphasises on a public cloud platform with a specific focus on an Infrastructure as a Service (IaaS) cloud service delivery model. The proof-of-concept is limited to CloudSim – a simulated cloud, and OpenNebula and OwnCloud public cloud computing platforms.

1.6 RESEARCH METHODOLOGY

This study adopted a pragmatic research method which begins with conceptual research method, followed by an in-depth empirical method and ends with an argumentative approach. A conceptual research methodology lays out key factors, constructs, concepts or variables, as well as the relationships between them (Jabareen, 2009). The idea is to use the conceptual research method to break down cloud security into the concepts of VM placement, authentication and digital forensic readiness. The aim was to gain a deeper understanding of each of these concepts before they can be combined to solve the complex data leakage problem. The study therefore conducts a critical review and analysis of existing literature to uncover prevailing cloud security trends with a specific focus on VM placement, authentication and digital forensic readiness concepts of information security. It aimed to identify trends within these three concepts of information security in order to assess their significance and to best position this research. Furthermore, the critical analysis made of existing literature on these three concepts was meant to identify research gaps in the field. The identified research gaps are the foundation for this thesis, and our research improves VM placement, authentication and digital forensic readiness in the new cloud environment so as to counter the widespread data leakage threat. Furthermore, these research gaps show how future research could contribute to the field of information security.

The output of the conceptual research method affirms the significance and positioning of this study in the body of knowledge. Furthermore, this yields a list of system requirements that are gleaned from the research gaps detected in existing literature as well as current cloud security trends. The system requirements are used to design and create a conceptual architecture that integrates VM placement, authentication and digital forensic readiness to solve the data leakage problem in the cloud.

Conceptual research methodology was used in conjunction with an empirical research methodology. The empirical part of the study provided a proof-of-concept implementation. The proof-of-concept is a test-bed that yields experimental results to evaluate the conceptual architecture. The empirical research methodology demonstrated the conceptual architecture’s practicability and suitability to solve the research problem. This is followed by an argumentative research approach where the author uses abductive reasoning to rigorously test and evaluate the proposed solution. The next section discusses the terminology used in the thesis.

1.7 TERMINOLOGY

To avoid ambiguity or misunderstanding, this section ensures that readers have a clear understanding of what each term, as used in the thesis, means.

Risk-based authentication – a scalable method of authentication that takes into account the risk profile of a user attempting to access a system and appropriately determines the correct complexity of the security challenge required to authenticate the user (Dlamini et al., 2015; Dlamini et al., 2016; Goode, 2015; Misbahuddin, Bindmadhava and Dheeptha 2017).

Multi-factor authentication – a method of authentication that adds a second or third layer of security to user login credentials and requires the use of more than one user identity verification factor (Dlamini et al., 2012; Strom, 2015; Dostalek, 2019).

Digital forensics – defined as the use of scientifically derived and proven methods towards the identification, collection, preservation, validation, documentation, analysis, interpretation and presentation of digital evidence derived from digital sources for the purpose of facilitating or furthering the reconstruction of events suspected to be of a criminal or malicious nature or assisting to anticipate unauthorised actions (ISO/IEC 27043:2015, 2015; Dlamini et al., 2014).

Digital forensic readiness – the proactive preparation and planning involved in the identification, collection, preservation, validation, documentation, analysis, interpretation and presentation of digital evidence derived from digital sources for the purpose of facilitating or furthering the reconstruction of events suspected to be of a criminal or malicious nature or assisting to anticipate unauthorised actions before they actually happen (adopted from ISO/IEC 27043:2015, 2015; Dlamini et al., 2014).

Conflict-aware VM placement – a method that places virtual machines in the cloud infrastructure, based on a tenant's conflict of interest, conflict tolerance levels (CTLs), risk exposure and cost associated with the placement in order to reduce confidential data leakage threats (Kulkarni and Annappa, 2019; Dlamini, Eloff and Eloff, 2014; Su et al., 2015; Ratsoma et al. 2015).

1.8 THESIS LAYOUT

Figure 1.2 depicts the structure of the thesis, which is divided into five parts.

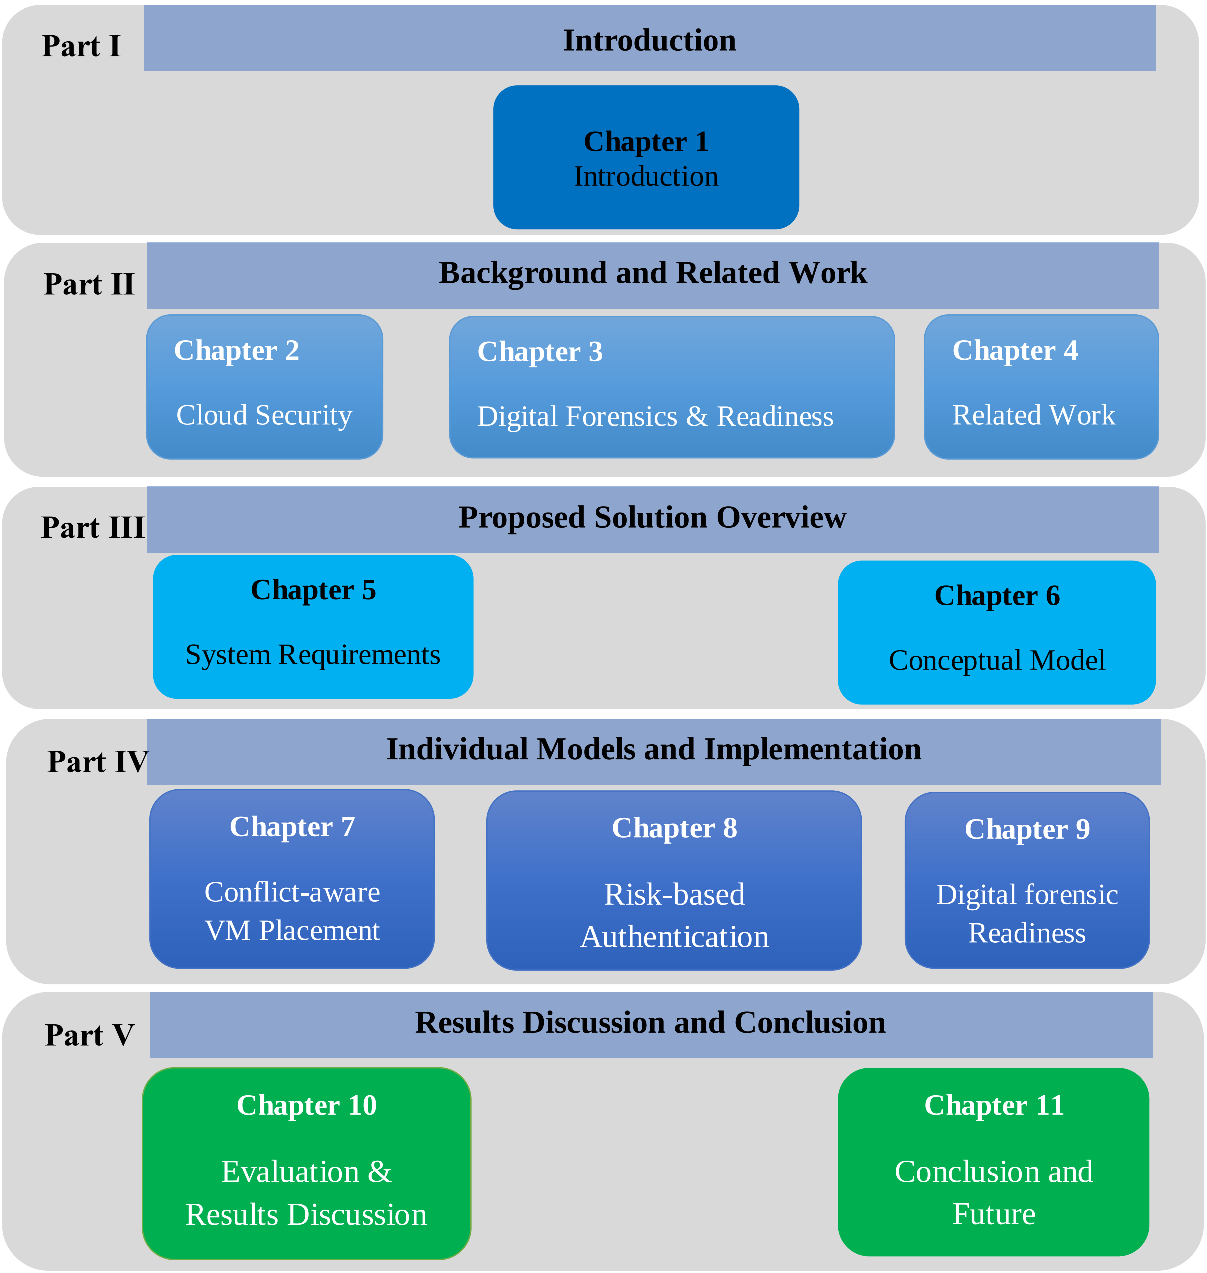
Part I
Introduction
Chapter 1
Introduction
Part II
Background and Related Work
Chapter 2
Cloud Security
Chapter 3
Digital Forensics & Readiness
Chapter 4
Related Work
Part III
Proposed Solution Overview
Chapter 5
System Requirements
Chapter 6
Conceptual Model
Part IV
Individual Models and Implementation
Chapter 7
Conflict-aware VM Placement
Chapter 8
Risk-based Authentication
Chapter 9
Digital forensic Readiness
Part V
Results Discussion and Conclusion
Chapter 10
Evaluation & Results Discussion
Chapter 11
Conclusion and Future

Figure 1.2:- Thesis Structure

Part I hosts the introduction in Chapter 1, which highlights the research problem. It also outlines the research hypothesis, main research question and sub-questions to be answered. Furthermore, this chapter discusses the research objectives, defines key terms and outlines the overall thesis structure.

Part II provides background work (in Chapters 2 and 3) with regard to cloud security, digital forensics and digital forensic readiness to set the scene. Chapter 4 in Part II provides a critical analysis of existing work to help position this study within the body of knowledge. Chapter 4 is also aimed at identifying and reflecting on currently existing research gaps. The thesis derives its system requirements from the identified research gaps within Chapter 4.

Part III comprises of two chapters, i.e. Chapter 5 and 6. Chapter 5 outlines and discusses the system requirements as gleaned from the existing research gaps in Chapter 4. Chapter 6 presents the conceptual architecture design, based on the system requirements discussed in Chapter 5.

Part IV comprises of three chapters i.e., conflict-aware VM placement in Chapter 7, risk-based MFA (Dlamini et al., 2016; Dlamini et al., 2017) in Chapter 8, and Digital Forensic Readiness (Dlamini et al., 2014) in Chapter 9. These chapters provide more details on each component of the conceptual model in Chapter 6. Chapter 7 introduces a conflict-aware VM placement model (Dlamini, Eloff and Eloff, 2014) that places VMs in the cloud, based on their conflict tolerance levels. Chapter 8 provides strong risk-based authentication. Chapter 9 presents a digital forensic readiness component that proactively captures digital evidence. It stores the evidence in a forensically sound manner or sends an alert to an investigator if there is a need to act in real time to stop an attack from happening.

Part V comprises of two chapters i.e., Chapter 10 which provides an evaluation and discussion of the results, and Chapter 11 which concludes the thesis and discusses potential future work.

1.9 CONCLUSION

Data leakage threats are a major concern for organisations that intend to move their confidential data to the cloud. The increasing number of reports on data leakage incidents is eroding the rather limited consumer confidence in cloud services, and some organisations are still hesitant to fully trust the cloud with their confidential data. Therefore, the study reported on in this thesis intended to make a plausible attempt to address data leakage threats in the cloud. This work attempted to address the data leakage problem on three fronts, namely conflict-aware VM placement, improved authentication and proactive digital forensic readiness. The idea was to strengthen cloud security in the hope of eventually restoring consumer confidence with regard to cloud service.

The next two chapters discuss background work before related work is presented in Chapter 4. Together the three chapters lay a solid foundation for the work to follow in the rest of the thesis.

CHAPTER 2 BACKGROUND: SECURITY OF CLOUD COMPUTING

2.1 INTRODUCTION

Today’s business environment is characterised by ever-increasing business competition and a rapidly changing ICT landscape. If organisations are to cope with today’s fast-paced digitally driven business environment, they must embrace the rapid technological changes and be ready to digitally transform their business models. Cloud computing is one technological change that organisations are required to embrace if they are to cope with the ever-increasing business competition and digital transformation. The true potential of cloud computing lies in its capacity to radically transform business models and help organisations respond to the fast-changing business landscape.

Cloud computing offers a new model for provisioning and obtaining computing resources as a service. Cloud computing creates a highly dynamic environment, where everything is delivered, provisioned and consumed as a service. Everything in the cloud is provided as a service (Dlamini et al., 2017; Duan et al., 2015; Gornaik et al., 2010). This model offers an on-demand and dynamic access to computing resources such as storage, applications, platforms and computing power as a service over the Internet. Cloud computing has compelling benefits that include significant cost savings, agility, high resilience and service availability (Fortinet, 2016). However, cloud computing has not yet reached the expected adoption rates. In the words of Baudin (2010), just like an antique light bulb, it still generates more heat than light.

The weak adoption and integration of cloud computing by most organisations result from the fact that some of them are still spectators who watch the early adopters from the side lines (El-Gazzar, Hustad and Olsen, 2016; Khan and Al-Yasiri, 2016; Sabi et al., 2016). This is an indication that cloud computing has not yet been fully exploited by the majority of its potential customers, probably due to numerous problems. One such problem is related to the misconceptions regarding cloud computing security issues, which must be clarified to improve its adoption rates. There is a growing need to separate real and pertinent security concerns from possible over-reaction, and from the hype and fear of the unknown that currently prevail within cloud computing (Dlamini et al., 2012). Over-reaction, hype and fear of the unknown have led to a gross generalisation that "security concerns are the biggest challenges of cloud computing" (Mell and Grance, 2009; Cloud Security Alliance, 2009).

Hence, this chapter takes the current discussions of cloud security beyond the over-reaction, hype and fear of the unknown to clear the fog hovering over such a promising computing paradigm. The goal, however, is not just to help potential cloud customers see beyond the fog that surrounds cloud computing security, but also to inform and make potential cloud computing customers aware of the overarching security issues that they should worry about.

This chapter is structured as follows: it first discusses background in terms of cloud computing and its benefits. This is followed by a discussion of the background of cloud computing security and of the taxonomies of security issues that need to be addressed in cloud computing environments. This chapter ends by highlighting the real security issues that this study is attempting to tackle.

The section below presents the fundamentals of cloud computing. It starts off with a definition of cloud computing and its basic concepts to provide a better understanding of the milieu of this research. The basic cloud computing concepts are discussed in the context of security, wherever it is possible to do so. The chapter ends by discussing the benefits of cloud computing.

2.2 DEFINING CLOUD COMPUTING

Cloud computing can be loosely defined as a new computing paradigm that makes possible the utilisation of computing infrastructure at a level of abstraction as an on-demand service, made available over the Internet (Mell and Grance, 2009). However, the most comprehensive definition of cloud computing is found in the NIST (Mell and Grance, 2011). This definition is well accepted and widely used by most researchers and experts in the cloud computing arena (Jula, Sundararajan and Othman, 2014; Hashem et al., 2015; Lewis, 2017; Subramanian and Jeyaraj, 2018; Noor et al., 2018; Ritchey, 2011). It defines cloud computing as "a model for enabling ubiquitous, convenient and on-demand access to a shared pool of configurable computing resources and services that can be rapidly provisioned and released with minimal management effort or service provider interaction" (Hashem et al., 2015; Subramanian and Jeyaraj, 2018; Noor et al., 2018). This model promotes and facilitates resource and service availability. It is based on five essential characteristics (on-demand self-service; broad network access; resource pooling; rapid elasticity; measured service). It also adopts three service models (software as a service; platform as a service; infrastructure as a service), and four deployment models (private; community; public; hybrid cloud) as depicted in Figure 2.1.

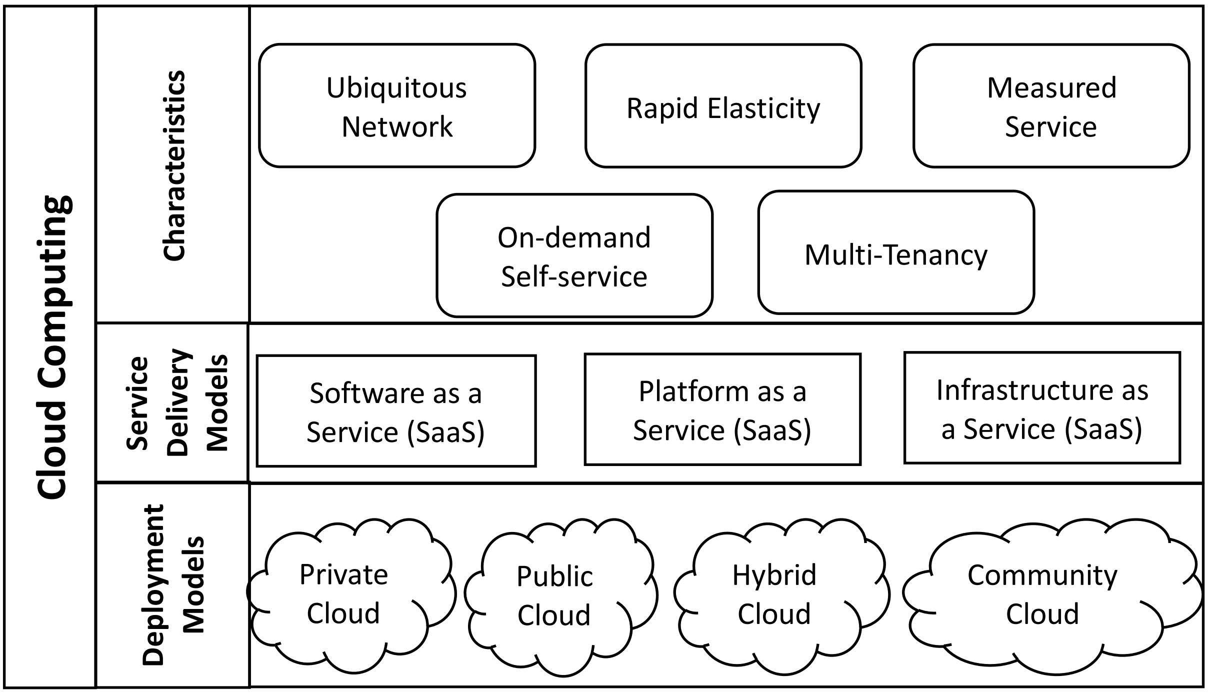

Figure 2. 1:- NIST Cloud Computing Definition Framework (Mell and Grance, 2009; CSA, 2009)

In order to grasp the importance of cloud computing, there is a need to understand its fundamental concepts and how they relate to security. Below, the author discusses each of these key concepts in relation to security.

2.3 CLOUD DEPLOYMENT MODELS

Cloud services are provisioned using four deployment models, namely public, private, hybrid and community clouds. Each of these has different advantages, disadvantages and constraints. All four deployment models also provide different levels of security. Below, the author discusses each of these deployment and service delivery models in the context of security. This discussion is necessary to point out the different security challenges that are associated with each of the key concepts of cloud computing. It is important to discuss these and show that there is no one-size-fits-all solution to addressing security concerns in the cloud. The discussions follow a bottom-up approach and start with a discussion of the deployment models, followed by discussions of service delivery models, and lastly, the characteristics of cloud computing in general.

2.3.1 Public Cloud

This deployment model offers highly scalable and shared public cloud services. Public cloud services are always located off-site and they are normally accessible through web interfaces. Furthermore, public cloud services offered by this model are available to the general public at a low cost. Public cloud deployment models abstract IT resources and make them appear limitless. Examples are Amazon Elastic Compute Cloud (EC2), Google AppEngine, Microsoft Windows Azure, Microsoft Office 365, Gmail and Dropbox. Each of these is accessible to the general public through a web interface over the Internet. However, public clouds offer cloud services that appear to be less secure and more risky than other deployment models (Fernandes et al., 2014). This is mainly because the cloud service provider has total control over the security of services deployed on a public cloud infrastructure. The customer has little or no control over the security of public cloud services. For instance, a customer may encrypt a file before hosting it on Dropbox. However, the same customer is not able to stop the same file from being leaked to a third party in its encrypted format, due to a vulnerability in the underlying infrastructure. The lack of security or customer control in public cloud deployment models has necessitated the research in hand.

Therefore, by focusing on this deployment model, this research is making a real contribution towards addressing a pertinent problem that has been noted by many other researchers.

2.3.2 Private Cloud

A private cloud deployment model delivers cloud services within the confines of a single entity, usually behind a firewall. This somehow gives users full control of what comes in and goes out of the cloud services. However, private clouds are normally associated with huge capital and operational costs, and they require highly skilled technical staff to manage them in terms of security, performance, reliability and compliance. The user base of private cloud deployments is only limited to users within an entity's walls. Access to cloud services deployed in a private cloud does not necessarily require a web interface. Such services can be accessed directly or through a dedicated virtual private network (VPN). In a private cloud deployment, users have total control of all cloud services running on their cloud. For instance, if a user encrypts a file and stores it in the private cloud, they can also control who has access to it. This somehow improves the level of security as compared to a public cloud deployment. Although this thesis does not directly address security challenges in private cloud deployments, the results reported in this research could also be usable for private cloud deployments.

Given the huge costs that are associated with private cloud deployments, consideration is given to taking on-site private clouds and hosting them with external third party providers (Fernandes et al., 2014). These are called external private clouds (Van Winkle, 2012; Kaur, 2017). An external private cloud deployment model differs from a public cloud deployment in that its cloud services are offered by a third party and to a single entity. For this deployment model, the dedicated cloud services run on the entity's dedicated hardware and software infrastructure. This means that the user base is not the general public, but a specific user group belonging to just the one entity. In this model, there is no sharing of resources with users outside the intended user base. Furthermore, the responsibility to secure external private cloud deployments is shared between the client entity and the hosting third party. Hence, this model can be argued to provide an even higher level of security compared to private clouds. However, the external private cloud deployment model is a quite novel idea and it has not yet been widely accepted. At the time of writing this thesis, only Microsoft Windows Azure was found to be experimenting with it (Van Winkle, 2012). Therefore, this thesis does not consider the external private cloud deployment models.

2.3.3 Hybrid Cloud

Hybrid clouds are formed from a combination of public and private clouds. This deployment model offers benefits of both the public and private model in one solution – i.e. hybrid clouds (Fernandes et al., 2014; Leavitt, 2013; Bittercourt and Madeira, 2011). Technologies like CloudSwitch, OpenStack and Eucalyptus are already designed with a capability to facilitate hybrid clouds. For example, OpenStack is designed for Amazon's EC2 services (Leavitt, 2013). A hybrid cloud provides low cost, elasticity and scalability of public clouds coupled with private cloud's customisation and high levels of security. The result is a combination of control over security, more like an external private cloud deployment model. An entity using a hybrid solution can host their sensitive applications on the private side of their deployment and their less sensitive applications on the public deployment side. This yields a greatly improved level of security for sensitive applications or data, whilst also making use of the low cost benefits for less sensitive applications. It provides entities the best of both worlds at a better cost, compared to fully private cloud deployment and better security than an entirely public cloud deployment. However, a hybrid cloud deployment model for some cloud services like SaaS requires careful consideration when determining applications that execute in the public and those that execute in private (Lewis, 2017; Bittercourt and Madeira, 2011). This thesis does not directly address security challenges of a hybrid cloud deployment. However, this research has an indirect connection with regards to addressing the challenges in a public cloud. It can be argued that this thesis is addressing only the security implications of a public-private hybrid cloud deployment.

2.3.4 Community Clouds

This deployment model is normally shared by multiple entities with a common interest. For example, the entities could be universities collaborating on some niche research area or companies investigating a new drug or epidemic disease. An example of a community cloud is Amazon's GovCloud and Microsoft's Azure Government (Lewis, 2017). Community clouds allow multiple entities to access the cloud, more like in a public cloud. However, in a community cloud, access is only limited to a close set of entities. This model is normally controlled by the group of entities or a third-party entity. Furthermore, it could be deployed off-site (in a third party) or on-site at one of the participant entities. A community cloud deployment model reduces the security risks associated with public clouds. Furthermore, this model ensures that the high cost of private clouds is shared among the participant entities. The security of this model is dependent on the security of the participants and the hosting third party. The community cloud deployment model also falls outside of the scope of this thesis.

2.3.5 Summary of Cloud Deployment Models

In summary, public clouds are associated with a high level of security risks because they are open to the general public. Furthermore, users lose control of their data once it is stored in a public cloud. This makes users to rely solely on the cloud service providers for security and traceability of their data. Public cloud deployments are therefore more susceptible to data leakage threats which could at times happen without any traceability. This is followed by hybrid clouds which provide an improved level of security; especially on the private side of the hybrid deployment. The security risks on the public side of a hybrid cloud will however remain. A hybrid cloud inherits the challenges of public clouds, but these are somehow neutralised by the benefits of private clouds. Users of hybrid cloud deployments could decide to store their critical data on premise within the private cloud and push their public data to the public cloud deployment. Therefore, data leakage threats are more likely to occur on the public side of the hybrid cloud. Fortunately, the impact would be minimal because of the public nature of the data that is stored on the public cloud. The community cloud deployment model closely resembles a hybrid cloud. The only advantage is that a community cloud is open to a small group of entities and no part of it is open to the general public, as the case may be in hybrid clouds. Hence, data leakage threats are not likely to happen in community clouds. Private clouds are somehow more secure than public, hybrid and community clouds. However, private clouds are argued to be less secure than external private clouds. External private clouds, though not yet extensively covered in existing work, seem to be the most secure option of all the deployment models, which makes them least likely to suffer data leakage threats.

The next subsection discusses the different cloud service models, namely Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). Each of these service delivery models is associated with different security challenges and it is important to show such differences. This is also key for the reader so that they may not assume that a solution to one might apply to the other service delivery models.

2.4 CLOUD SERVICE DELIVERY MODELS

Cloud services are mainly provisioned using three traditional service delivery models – IaaS, PaaS and SaaS. These service delivery models differ in the manner in which responsibility and accountability are shared between the cloud service provider and the consumer of the provided services (Thilakarathne and Wijayanayake, 2014). For example, in a PaaS, the service provider manages the development environment and the underlying infrastructure, and the consumer controls over-the-top cloud applications. In an IaaS, the consumer has control over the virtual infrastructure, whilst the service provider controls the hardware below the virtualised infrastructure. In a SaaS, the service provider controls both the underlying infrastructure and the applications, while the consumer takes minimal control of the application configurations. Each of these service delivery models has different advantages, disadvantages and constraints. However, the author only discusses each of these service delivery models in the context of security and give examples in each case.

2.4.1 IaaS

This service delivery model offers scalable, elastic and on-demand physical or virtual computing resources such as storage, networking or computational processing power (Thales and Ponemon Institute, 2019). For example, Amazon EC2 offers on-demand virtual machines, paying only for what is being used (Fernandes et al., 2014). A cloud service provider could offer some basic level of security using a virtual machine monitor (VMM).

The VMM ensures a logical isolation between all VM instances running on the same cloud infrastructure. Amazon's IaaS EC2 takes responsibility for physical security, environment security, and virtualisation security up to the VMM. The cloud consumer must take responsibility for the security of its VM instances. However, there are several security issues concerning VMM. For example, a compromised VMM renders all VMs under its control vulnerable. It is even worse when the host machine gets compromised, as the security of the entire virtual space becomes questionable (Thilakarathne and Wijayanayake, 2014). Thus, there is a need for mechanisms that ensure stronger physical isolation guarantees (Takabi, Joshi and Ahn, 2010). This thesis emphasises strong physical isolation of cloud resources, which is closely linked to VM placement approaches. Furthermore, some cloud providers like Amazon's EC2 do not allow its administrators to log in and access guest VM instances belonging to their clients. This is one strategy to prevent malicious insiders from tampering with clients' data and applications. However, and by default, Amazon S3 (Simple Storage Service) does not encrypt client data at rest. This gives users an opportunity to take responsibility for the security of their data hosted in their rented VM instances. Users can encrypt their data before moving it to Amazon's IaaS S3 for an extra level of security (Mosola et al., 2016; Thales and Ponemon Institute, 2019). The IaaS service model forms the basis for all other service models. Lewis (2016) argues that many SaaS cloud service providers run on IaaS, and therefore, if data leakage threats are targeted to the underlying IaaS, all over-the-top applications and platforms would also suffer. This thesis places a strong focus on the security of the IaaS layer as a foundation layer on which all other service models are built.

2.4.2 PaaS

The PaaS service delivery model is a middleware that delivers services, systems and environments like runtime program development tools, platforms, libraries and frameworks on top of which developers can build, compile, test and run their own cloud applications (Linthicum, 2017; Thales and Ponemon Institute, 2019). It provides developers with readily available tools and services that offer an end-to-end life cycle of developing, testing, deploying and hosting cloud applications as a service (Rimal, Choi and Lumb, 2009).

Developers just focus on their applications. This approach can help to reduce development times (Linthicum, 2017). The service provider manages and provides updates and patches for the underlying hardware or software infrastructure on which the applications are developed. Developers are responsible for the security of their applications and service providers are responsible for the underlying infrastructure. Google App Engine is an example of a PaaS service offering. It offers software development kits (SDKs), and integrated development environment (IDEs) for programming in Python, Java and Go (Fernandes et al., 2014). Security challenges arise from different applications sharing the same computing resources. For example, unsafe thread termination or insecure system calls can result in inconsistent object states. If the development environments are not properly isolated, applications often suffer from a resource starvation problems. This model is also vulnerable to malicious insider threats. Bad software development practices from one consumer of PaaS is likely to affect other consumers. Hence, the provider must ensure that all consumers adhere to good coding practices. Since the PaaS service delivery model is not so much related to the data leakage threat, this thesis does not focus on it.

2.4.3 SaaS

This service delivery model abstracts software and provides it over a web application without the need for installation, customisation and configuration for use as a service on demand (Thales and Ponemon Institute, 2019). For example, Microsoft Office 365 is offered in an on-demand manner and paid for on a pay-per-use basis. As SaaS and web applications are closely coupled, they share their individual security threats. The fact that web applications are a gateway to SaaS, makes it an attractive target that can compromise the entire SaaS service delivery. The consumer is required to ensure that its web interface is secure whilst the cloud service provider is responsible for security of its SaaS. The user has limited or no control over the security of the provided software. The service provider is required to prevent breaches due to the security vulnerabilities in the software. Ford (2012) raises an interesting security issue around non-transparent layering of structures where cloud services may appear independent but share deep and hidden resource dependencies that may create unexpected failure. For example, consider Microsoft Office 365 running on Amazon S3. To the consumer of MS Office 365 the underlying Amazon S3 infrastructure is hidden. However, a failure on the Amazon S3 layer has cascading effects on the MS Office 365. Ford (2012) refers to these as stability risks from interacting cloud services. They are some of the security concerns that have not yet been studied and are not well understood, yet they raise major challenges. This research nevertheless does not focus on the SaaS model, but on the underlying IaaS.

2.4.4 Summary of Service Delivery Models

The indication from the above three traditional cloud service delivery models is that cloud computing resources are delivered and consumed as a service. From these three, anything or everything is now offered in the form of a service. For example, Security-as-a-Service (Sec-aaS) (Furfaro, Garro and Tundis, 2014; Ghazi et al., 2016) and identity management-as-a-Service (IDMaaS) (Nuñez and Agudo, 2014). Some researchers have even coined the "Anything-as-a-Service" (XaaS) concept (Fernandes et al., 2014; Duan et al., 2015; Miyachi, 2018). The research in hand only considers the traditional three with a specific focus on the IaaS as the underlying model. It does not go into great detail about the new service delivery models (XaaS). These are added here as a mere confirmation that the author is well aware of them.

The next section briefly highlights the benefits of cloud computing before moving on to discuss cloud computing security. Highlighting the benefits is necessary to reflect on the value proposition that cloud computing can provide to companies that are willing to adopt it. It is important for this thesis to highlight the benefits to strengthen my argument on why this study is required and relevant. If the benefits of cloud computing are not pointed out, readers might ask, why is it important to do this research? The researcher therefore argues that the study will help companies to embrace cloud computing more easily and without worrying about any security concerns.

2.5 BENEFITS OF CLOUD COMPUTING

Cloud computing presents a number of significant benefits that can greatly improve the efficiency of ICT operations that will help organisations gain the competitive edge that is necessary to survive in today's challenging business environment. The most cited benefit of cloud computing is its significant cost saving (Hashemi and Aerdakani, 2012; Mell and Grance, 2009; Mell and Grance, 2011; Ponemon Institute, 2010). (See also Figure 2.2.) According to (Vasiljeva, Shaikhulina and Kreslins, 2017), access from anywhere and cost savings stand out as the two main benefits of the cloud. The others are faster deployment times, backup/disaster recovery, flexible pay-as-you-go model, autonomous updates, reduced on-site infrastructure and reduced workloads, to name a few.

In terms of cost savings, cloud computing lowers or removes intensive capital IT costs and transforms them into operational expenses to run core business operations. For example, Google, Amazon and Facebook data centre platforms have generated 55% savings on capital expenditure (CAPEX), up to 75% saving on operational expenditure (OPEX) and a whopping 138% return on investment (ROI) over a period of five years (Mainstay, 2016).

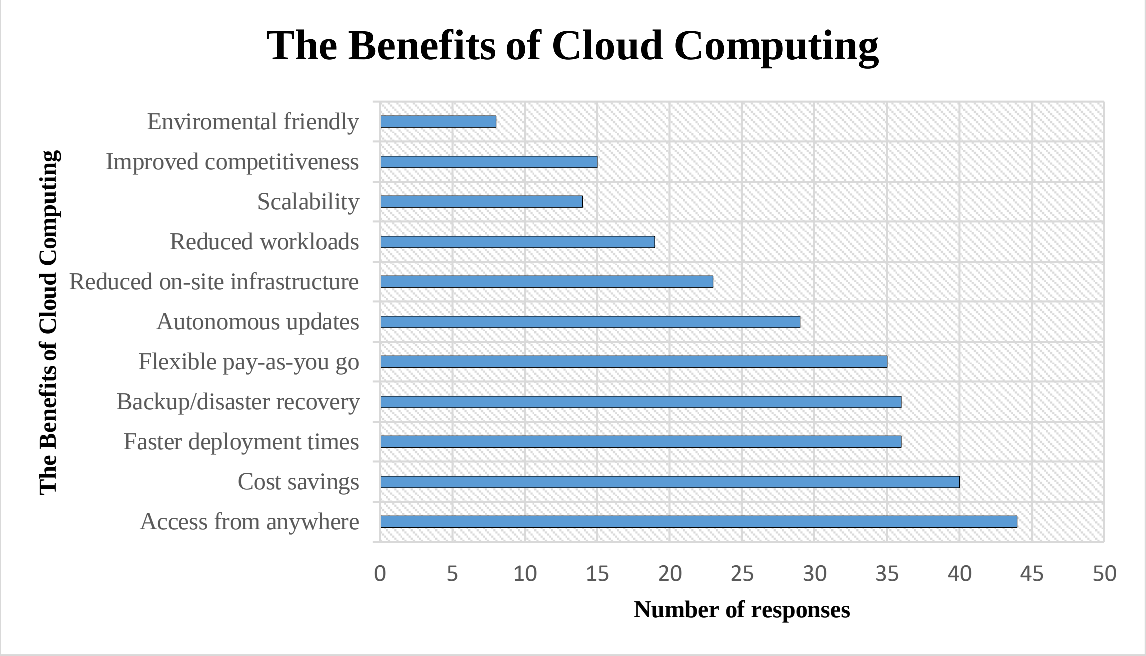

Figure 2.2:- The Primary Benefits for Cloud Computing (Vasiljeva et al., 2017)

Cloud computing greatly reduces the time to get businesses up and running. This is in relation to the benefit of faster deployment times. It also reduces or removes upfront costs of acquiring ICT infrastructure and offers instant access to flexible computing resources (Gregg, 2011; Ponemon Institute, 2010, Vasiljeva et al., 2017). The main beneficiary of cloud computing is most likely small businesses without economies of scale (Jansen and Grance, 2011). The other benefit of cloud computing is its high degree of redundancy that provides a high degree of service availability (Catteddu and Hogben, 2009; Vasiljeva et al., 2017). The level of redundancy provided in cloud computing makes the cloud infrastructure more resilient to security threats, failures and natural disasters.

Cloud computing also provides on-demand, elastic and scalable access to computing resources (Mell and Grance, 2011). On-demand means that each service in the cloud is requested as a need arises on a need-to-use (demand) basis and paid for on pay-as-you-use basis using a utility pricing model. Using the utility model, users pay for only what they have used or consumed. For instance, a customer pays for storage per gigabyte or terabyte per hour that their data is stored by the cloud storage service provider. Scalability and elasticity mean that the provisioned computing resources can either be increased or decreased, depending on a customer’s demand. To cloud computing customers, the computing resources appear as if they are unlimited. In support, Cable & Wireless Worldwide (2011) agrees that cloud computing provides infinitely (unlimited) flexible computing resources.

With public cloud computing, businesses need not worry about having the necessary expertise to run and maintain computing resources in-house. This responsibility is transferred to the cloud service providers. The aim is to reduce operating overheads for cloud customers and allow them to focus on their core competences. For example, with the IaaS and PaaS cloud model, applications and software running in the cloud infrastructure are maintained and managed by the cloud service provider (Gornaik et al., 2010). This facilitates timely updates and patches to quickly eliminate vulnerabilities before they can be exploited. Some researchers argue that cloud computing comes with better security; especially for small businesses without the technical security know-how. However, since the increasing number of security breaches that have occurred in the cloud in recent years, researchers are finding out that cloud computing is not as secure as has been anticipated. Hence, they have started to intensify their efforts towards cloud security.

The above by no means constitutes an exhaustive list of cloud computing benefits. However, the list is sufficient to illustrate the point about the glaring benefits of cloud computing. Cloud computing is believed to pose a revolutionary potential to radically change the way business is conducted and most organisations are well aware of its compelling benefits and potential. However, the recent spate of data leakages and breaches happening in the cloud is slowly eroding the rather limited consumer confidence in cloud services (Pandey, 2018; Gemalto, 2018). For example, the National Security Agency (NSA), the Pentagon and Accenture misconfigured their Amazon Web Services S3 buckets and in the process exposed hundreds of gigabytes of their data (Gemalto, 2018). The one major cause for concern about the recent data leakages and breaches is that Gemalto (2018) and Pandey (2018) report that only 1% of the 2.6 billion compromised, stolen or lost records in 2017 had been encrypted. Hence, securing data in the cloud has become a key concern.

The next section consequently discusses security concerns of cloud computing in detail. Security concerns are the major stumbling block for companies moving their data and applications to the cloud and have caused companies to be reluctant to adopt cloud computing services. There is a need to discuss the pertinent security concerns in the cloud. The point of departure is to identify these concerns and then move on to propose plausible solutions. The identified security concerns are covered in the next section.

2.6 CLOUD COMPUTING SECURITY

Numerous research efforts concur that cloud computing has considerable benefits for today’s organisations (Avram, 2014; Botta et al., 2016; Dokras et al., 2009; Oliveira, Thomas and Espadanal, 2014; Vasiljeva et al., 2017; EZComputer Solutions, 2018). Despite this, security concerns are cited as one of the biggest stumbling blocks for most organisations that consider moving their critical applications and sensitive data to the cloud (Dlamini et al., 2011; Diaz, Martin and Rubio, 2016; Ismail, Hassen and Zantout, 2016). Sceptics use security concerns as an excuse not to move to the cloud. This hesitancy has been exacerbated by the fact that cloud computing is still to prove its worth to its end users. A fear of the unknown has also contributed to the slow adoption of cloud computing. Much has been reported on security concerns about cloud computing and it has become difficult to separate real security concerns from hype, fear and confusion. To clear the air, the remainder of this section discusses some of the real security issues regarding cloud computing.

Cloud computing is a game-changing paradigm that has significantly changed the threat landscape with regard to service resilience, availability and the protection of sensitive data and applications (Alenezi, Atlam and Wills, 2019; Gornaik et al., 2010). Gornaik et al. (2010) assert that cloud computing depends on multiple independent cloud service providers (i.e. applications, data, infrastructure and platform), which create multiple points of failure. A failure in one provider could have cascading effects on other cloud services running on top of it. For example, a failure on Amazon’s IaaS (which provides storage and compute cloud services) could result in a failure of all cloud-hosted applications running on such an infrastructure. This creates an instability risk due to the unintended coupling of independent cloud services provided by different cloud service providers (Ford, 2012). Ford (2012) provides a detailed analysis of these ‘less understood’ cloud computing security risks.

Furthermore, cloud services share the same infrastructure, i.e. they operate on a public or hybrid cloud computing platform. The shared nature of cloud services (multi-tenancy) is cited as the key factor that contributes to serious security concerns in the cloud. Segregation of co-located data, computing resources and storage (among others) rely on software controls. This raises concerns about unintentional data leakages. Gornaik et al. (2010) report that access controls and identity management concerns turn out to be more complex within cloud computing environments. For example, in the cloud it is not enough to authenticate users based only on their credentials. The cloud requires that applications, end-user devices and the infrastructure that they are running on be authenticated as well.

Furthermore, cloud customers are heavily dependent on cloud service providers (CSPs) to manage most of their business services. Cloud customers surrender control of their data and applications to the cloud service providers. This is more of a problem in public clouds. Giving away control limits cloud customers’ situational awareness on looming security threats that could possibly affect their data (Gornaik et al., 2010). This has created the need for a solution that gives control back to the users.

Another interesting view raised in Gornaik et al. (2010) is that of audit logs and forensic data in case of a reported security breach. Cloud service providers have a legal obligation to protect and preserve the privacy of its customers, yet on the other hand they must provide audit logs and forensic data to law enforcers. The issue of data life cycle management is also raised in Gornaik et al. (2010). Cloud computing requires security measures that can securely create, process and destroy client data residing in the cloud. Most of the issues raised in Gornaik et al. (2010) are real and pertinent security concerns that are specific to a cloud computing environment.

A special publication (Jansen and Grance, 2011) by NIST provides an overview of the security challenges (related to system complexity, shared multi-tenancy, Internet-enabled services and loss of control) that are pertinent to cloud computing. Of note, this publication raises one of the most overlooked points in cloud computing: it has grown out of the combination of already existing technologies or paradigms such as distributed systems, service-oriented architecture and pervasive computing, with already known security issues being cast in a new environment (Gornaik et al., 2010; Jansen and Grance, 2011). In addition, Gornaik et al. (2010) and Catteddu and Hogben (2009) agree that cloud computing might be a new way of delivering computing resources, but it is definitely not a new technology. This suggests that there is nothing more to fear, as most security threats are already known, and effective security measures are in place. Cloud computing security issues may be more of a hype and over-reaction from sceptics who would wish to stifle the adoption of cloud computing.

In Jansen and Grance (2011), the complexity of the cloud computing environment is acknowledged as one of the factors that exacerbate cloud computing security issues. This complexity is due to the combination of existing technologies and many components (such as virtual machines, databases, supporting middleware, resource metering and billing, data replication, etc.) that combine to make cloud computing a reality. Most of these technologies and components already have known security issues with known mitigation strategies.

However, complexity comes when they interact, which raises new security concerns. Jansen and Grance (2011) argue that security is inversely proportional to complexity, i.e. greater complexity exponentially increases vulnerabilities. Their work ends with a number of recommendations for organisations that are planning to move their data, applications or infrastructure to the cloud. Most of the cloud computing security issues raised by Jansen and Grance (2011) are unaddressed concerns related to applicable characteristics of cloud computing such as system complexity, multi-tenancy, Internet-enabled services and loss of control.

Cable & Wireless Worldwide (2011) claims that security concerns are the final barrier for most organisations looking to adopt cloud computing. According to this company, organisations have understandable concerns about security issues. The reasons cited are that cloud computing is still considered a vague and intangible paradigm, and that clients cannot be certain about the whereabouts of their data or how it is handled, stored or transmitted. This argument raises reasonable and understandable doubts about the safety of their data. Therefore, Cable & Wireless Worldwide (2011) argues that most organisations need the following assurances before they would consider adopting cloud computing:

- Cloud computing will not compromise their security.
- Their sensitive data and intellectual property will be protected.
- They can easily retrieve their data should a need arise to change service providers.
- They can still maintain their standards and competitive performance.

The work done by Cable & Wireless Worldwide (2011) further examines and assesses the perceived security concerns about cloud computing to determine whether they are justified or not. The company recommends controls with the potential to make cloud computing security a reality.

Gregg (2011) in turn raises ten security concerns for cloud computing: geographical location of data in the cloud; regulatory requirements; access controls; classification and separation of data from different multi-tenants; handling security breaches; service level agreements; auditing; long-term viability of the cloud service provider (Bartolini et al., 2018); training and disaster recovery; business continuity plans. Most of these security concerns mentioned by Gregg (2011) are not necessarily specific to cloud computing and some have already been addressed in other environments. All that needs to be done now, is to tailor them for the cloud computing environment.

Dubey et al. (n.d.) discuss the problem of resource metering as one of the security issues in the cloud and relate it to the services that cloud service providers render to their customers. They also propose a solution that seeks to ensure that cloud customers are billed in accordance with the service rendered by the cloud service provider. Metering and billing have been successfully implemented in pay-as-you-go and pre-paid models for telecommunication and utility companies, e.g. Cisco VoIP Prepaid billing solution. With a minor adjustments, the same principles could also be used in cloud computing resource metering. For example, software as a service can be billed based on the features of an application that a customer uses and the time taken using it.

Dlamini et al. (2012) agree that security is an important issue in the cloud but argue that this issue has been blown out of proportion. They continue to show that some of the security concerns are nothing more than hype and fear of the unknown, and they demonstrate by using a few examples how existing security solutions could be tweaked and repackaged for successful application to the cloud environment (Dlamini et al., 2012; Dlamini et al., 2016).

Lastly, Cattaddu and Hogben (2009) also agree as others like Alenezi et al. (2019) that security is a top priority for cloud computing customers. They argue that cloud consumers make buying choices on the basis of confidentiality, integrity, availability and resilience of the security services offered by a cloud service provider. Catteddu and Hogben (2009) furthermore make recommendations on priority research areas that could help improve the security of cloud computing technologies.

In summary, it appears that researchers agree about security being viewed as a major concern for the wide adoption of cloud services. As a first step in the right direction, several research efforts have already been directed at identifying the actual security concerns in respect of cloud computing. As noted above, the identified security concerns vary a great deal and do not concern just the data leakage threat pointed out in Chapter 1. The current study notes the following key and pressing security issues that require attention and must be dealt with in the cloud:

- Multiple independent cloud service providers (i.e. applications, data, infrastructure and platform) working one on top of one another create multiple points of failure with cascading effects. This creates a need to secure the underlying IaaS delivery layer on which all other service delivery models hinge.
- The shared nature of cloud services (multi-tenancy) raises concerns about unintentional data leakages and causes segregation of co-located data, computing resources and storage to be key in improving the security of public clouds.
- Cloud users continue to lose control of their data and applications, which limits their situational awareness. Hence, there is a need for a solution that gives control back to the users.
- Concrete and credible audit logs and forensic data are required in case of a reported security breach.
- It is not enough to authenticate users based on their credentials in the cloud. It is essential to authenticate applications, end-user devices and other infrastructure.

Furthermore, the above and other existing studies lack a high-level and more holistic perspective of all the security concerns regarding cloud computing environments (Fernandes et al., 2014). Hence, the next section reflects on how some researchers have tried to provide a holistic picture of security concerns in the cloud. It must be pointed out, as noted by Fernandes et al. (2014), that research in this space is insufficient and scanty.

2.7 TAXONOMY OF SECURITY CONCERNS IN CLOUD COMPUTING

Despite insufficient research findings in this space, some researchers have already attempted to classify the security concerns of cloud computing in a form of a taxonomy (Fernandes et al., 2014; Hashemi and Ardakani, 2012; Iqbal et al., 2016) in order to provide a complete picture of security landscape in cloud computing.

For example, the work of Iqbal et al. (2016) provides a taxonomy of security concerns in the context of cloud service delivery models. Therein, it is argued that each cloud service delivery model is associated with specific security concerns. Figure 2.3 depicts a taxonomy of cloud security issues associated with their service delivery models.

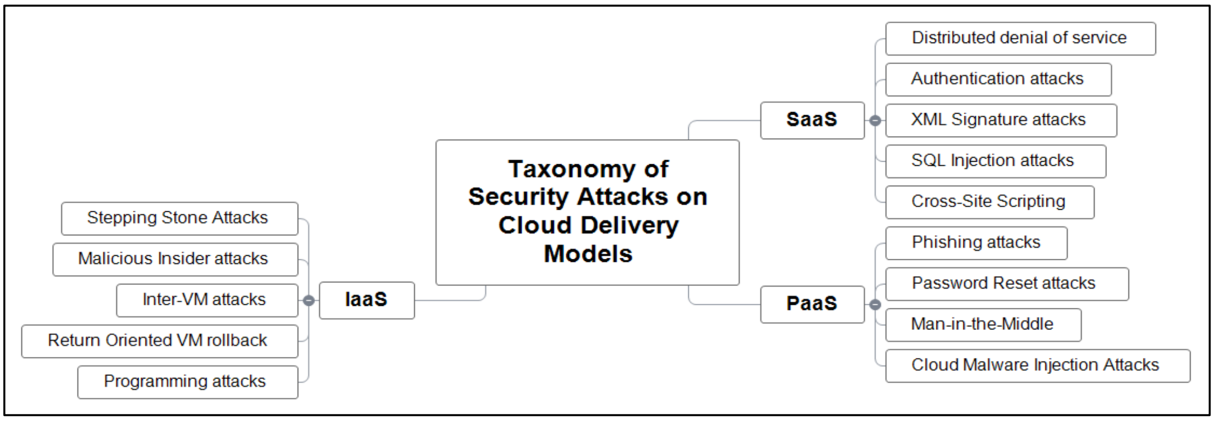

Figure 2.3:- Taxonomy of Attacks on Cloud Service Delivery Models (Iqbal et al., 2016)

Figure 2.3 shows that inter-VM attacks are mainly associated with the IaaS service models, while authentication attacks are associated with the SaaS delivery model. In the interest of brevity, we do not discuss all threats. For more details on each of these attacks, the reader is directed to the work of Iqbal et al. (2016). (The same applies to Figures 2.4 and 2.5.)

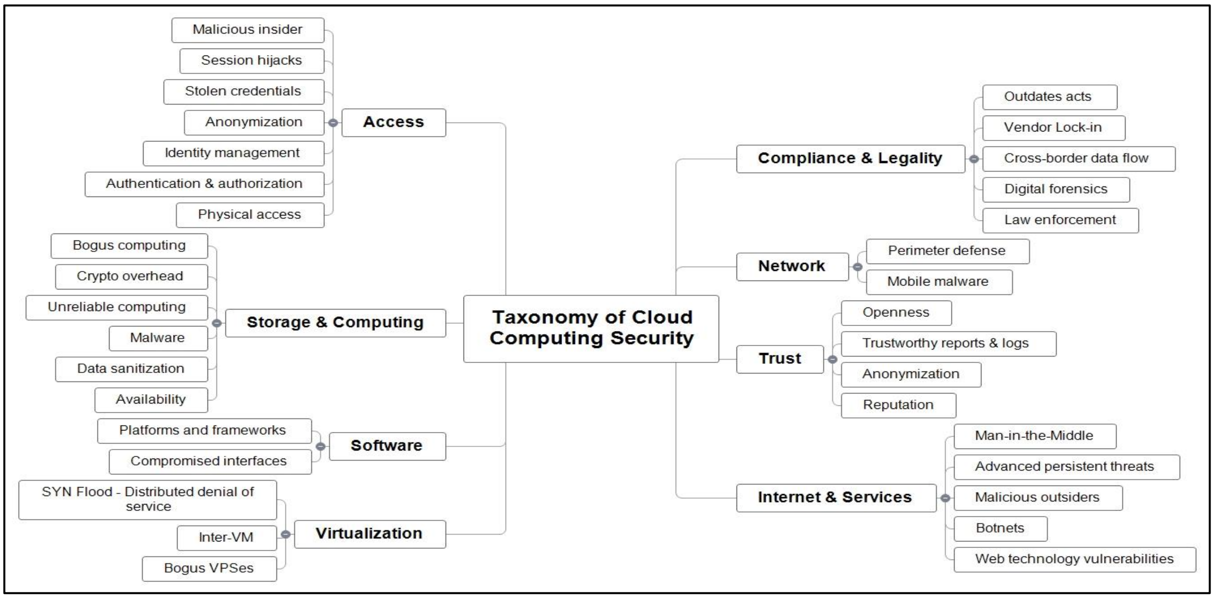

Figure 2.4:- Taxonomy of Security Concerns in Cloud Environments (Fernandes et al., 2014)

Figure 2.4 depicts the taxonomy of cloud computing security issues associated with eight main categories, i.e. software; storage and computing; virtualisation; Internet and services; network; access; trust; compliance and legality. For example, Figure 2.4 shows that inter-VMs are associated with virtualisation of the cloud. An interesting thing to note is the aspect of digital forensics, which is not included in Figure 2.3. Figure 2.5 presents a more comprehensive taxonomy that touches on the overarching research problem of this thesis, namely addressing data leakage threats. Again, inter-VM attacks are part of the taxonomy in Figure 2.5. In principle, inter-VM attacks appear in all three taxonomies. Since this is an indication that more research efforts are still required to address inter-VM attacks, this thesis focuses specifically on addressing the threat of inter-VMs.

Most of the work covered agrees that cloud computing has attractive and compelling benefits for those organisations that seek to embrace it. It seems that researchers concur that security concerns remain a challenge for the adoption of cloud computing. However, some of the studies, except the work of Cable & Wireless Worldwide (2011); Dlamini et al. (2012) and Ismail et al. (2016) fail to acknowledge that some of the perceived security concerns of cloud computing could just be an over-reaction or result from a fear of the unknown by sceptics. Failure to acknowledge and show that cloud computing security is an issue that has been blown out of proportion, as well as failure to pin-point pertinent security concerns regarding the cloud might just worsen the current situation and direct future research efforts towards invalid security concerns. It is in the wake of these and other issues that this chapter has attempted to reflect on the overall picture of security concerns in the cloud so as to bring clarity to a complicated cloud computing threat landscape. This landscape is quite often filled with hype, fear of the unknown, incomplete and oversimplified information, all of which have led to a gross generalisation that “security remains the biggest challenge for cloud computing”.

Given the plethora of security issues regarding the cloud (as depicted above), it would really be impossible to address all of them at once. Hence, the research in hand focused on addressing data leakage problems from three perspectives only, i.e. authentication flaws, inter-VM attacks arising from inadequate VM placements, and digital forensics. The idea was basically to strengthen authentication as a first point of entry to cloud services. We have noted that cloud computing requires a different approach to traditional authentication.

Furthermore, and of note is that inter-VM attacks seem to be cutting across all discussions of security issues in the cloud. Therefore, this research took the initiative to propose a solution to this problem and addressed the problem from a VM placement perspective. Moreover, and in order to strengthen the proposed solution, this study also investigated the use of digital forensic readiness to ensure that all access activity for investigation purposes is captured for future use.

2.8 CONCLUSION

Extensive industrial and academic research efforts have cited security concerns as one of the biggest challenges preventing organisations from migrating their data and applications to the cloud. Surely, security is an important issue in the cloud, but many researchers argue that this issue has been blown out of proportion, mainly because of the hype and fear of the unknown associated with cloud computing. A careful look reveals that most of the security concerns do not really pose something new that we should be worried about. Some of the concerns have already been addressed in different environments like virtualisation and would probably need to be repackaged with a few tweaks to fit into the new cloud environment. For example, authentication is an old security mechanism that requires some modifications to fit the cloud environment. In summary, Chapter 2 reflected on the security levels that vary according to the cloud deployment and service delivery models. This could allow consumers of cloud services to take calculated incremental steps towards their adoption of cloud computing. For, example, consumers could start off by using a public cloud only for non-critical data and applications.

Chapter 3 takes a look at background work in terms of digital forensic readiness.

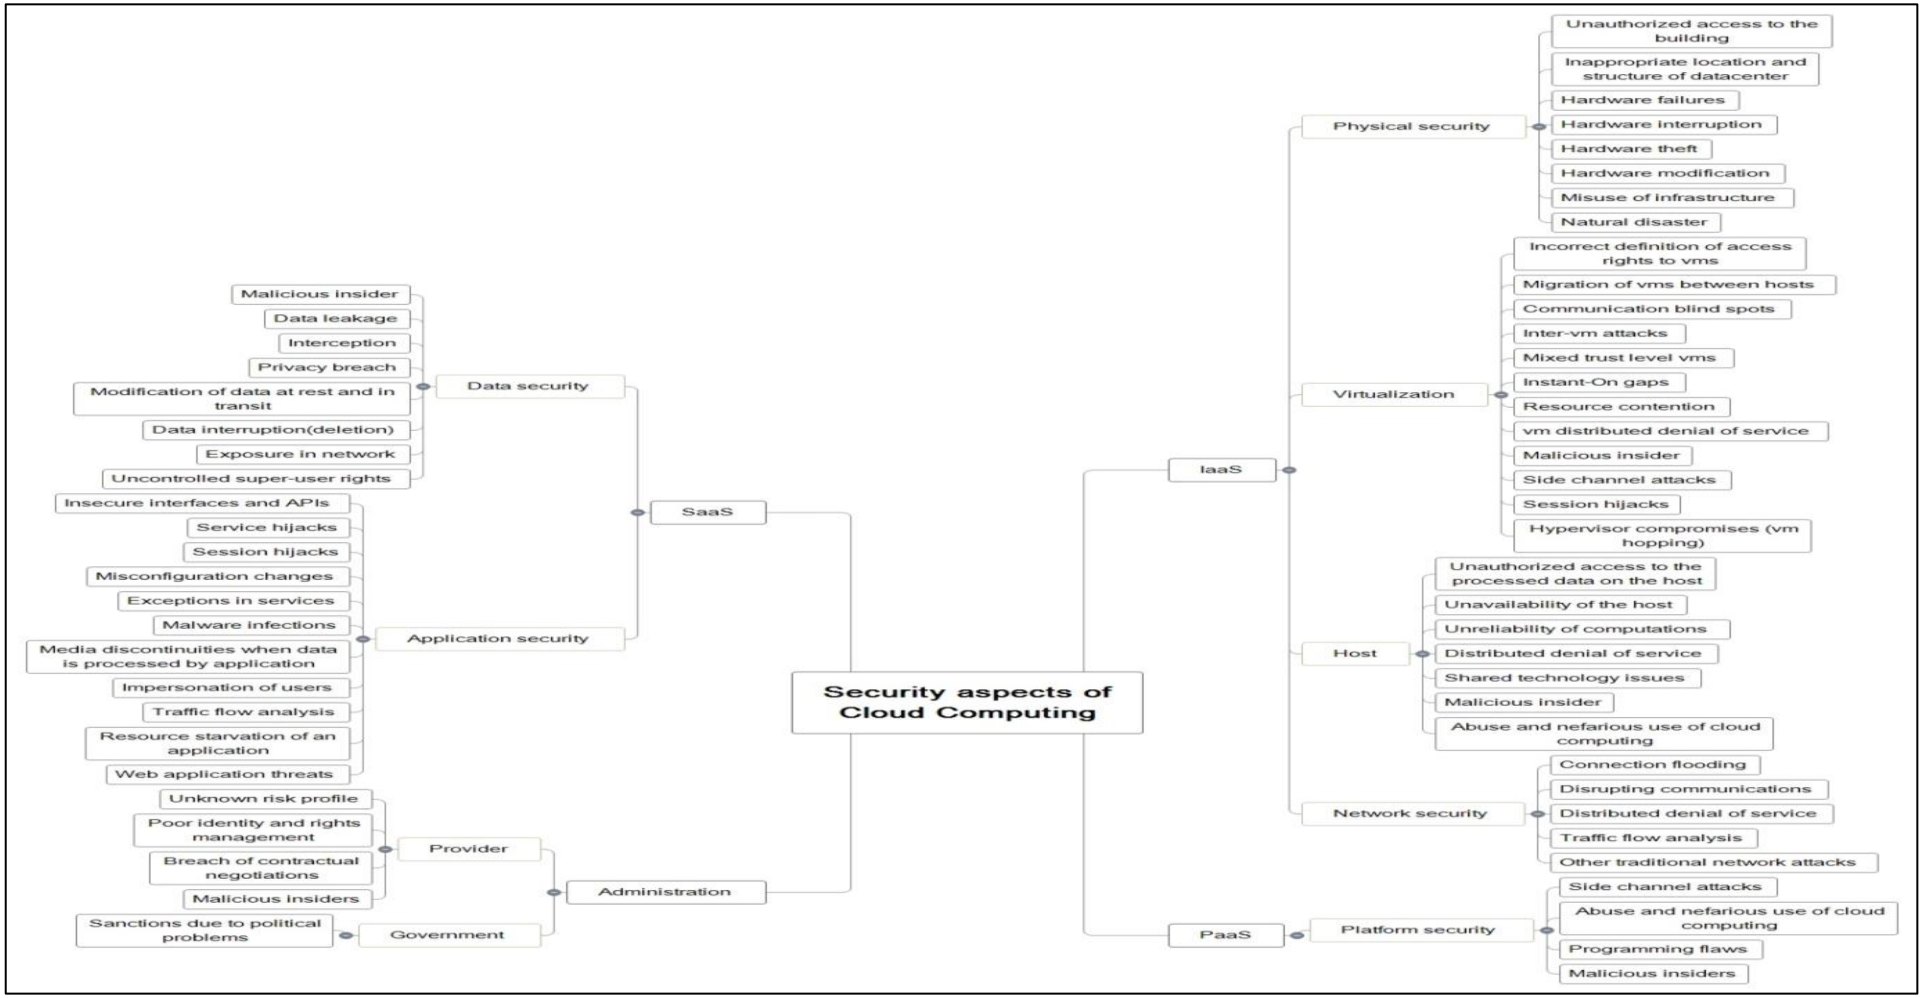
Security aspects of Cloud Computing
SaaS
Data security
Malicious insider
Data leakage
Interception
Privacy breach
Modification of data at rest and in transit
Data interruption(deletion)
Exposure in network
Uncontrolled super-user rights
Application security
Insecure interfaces and APIs
Service hijacks
Session hijacks
Misconfiguration changes
Exceptions in services
Malware infections
Media discontinuities when data is processed by application
Impersonation of users
Traffic flow analysis
Resource starvation of an application
Web application threats
Administration
Provider
Unknown risk profile
Poor identity and rights management
Breach of contractual negotiations
Malicious insiders
Government
Sanctions due to political problems
IaaS
Physical security
Unauthorized access to the building
Inappropriate location and structure of datacenter
Hardware failures
Hardware interruption
Hardware theft
Hardware modification
Misuse of infrastructure
Natural disaster
Virtualization
Incorrect definition of access rights to vms
Migration of vms between hosts
Communication blind spots
Inter-vm attacks
Mixed trust level vms
Instant-On gaps
Resource contention
vm distributed denial of service
Malicious insider
Side channel attacks
Session hijacks
Hypervisor compromises (vm hopping)
Host
Unauthorized access to the processed data on the host
Unavailability of the host
Unreliability of computations
Distributed denial of service
Shared technology issues
Malicious insider
Abuse and nefarious use of cloud computing
Network security
Connection flooding
Disrupting communications
Distributed denial of service
Traffic flow analysis
Other traditional network attacks
PaaS
Platform security
Side channel attacks
Abuse and nefarious use of cloud computing
Programming flaws
Malicious insiders

Figure 2.5:- Taxonomy of Security Aspects of Cloud Computing (Hashemi and Ardakani, 2012)

CHAPTER 3 BACKGROUND ON DIGITAL FORENSICS AND DIGITAL FORENSIC READINESS

3.1 INTRODUCTION

As reflected in Chapter 2, security concerns have been coined as a major stumbling block for most organisations intending to move their critical data, applications and systems to the cloud (Dlamini et al., 2011). Surely, in today's complex and inter-connected business ecosystem, security threats are inevitable. Moreover, security is an essential requirement for IT resources and no organisation would want its data and applications to be stored on insecure systems that would expose them to unauthorised users (Reilly et al., 2011). It is on this premise that most research efforts have been directed at security challenges of the cloud for all that the subject merits. In the meantime, the research community has left an open area of research of equal importance: that of digital forensics investigation in the cloud. Du, Le-Khac and Scanlon (2017) consequently argue that digital forensics is still in its infancy.

Surely, it is vital to protect and secure cloud systems from inevitable security threats. However, it is equally important to be able to deal with the aftermath once an incident has occurred. Organisations must be able to proactively collect and preserve digital evidence to detect malicious activities in the hope of identifying the responsible culprits. This could help to hold malicious culprits accountable for their actions. Although digital forensic investigation fills that gap. It quite often comes into play after an incident has already occurred, as a reactive approach.

Within the agile and multi-tenant cloud environment, organisations cannot afford to follow a reactive approach to digital forensic investigations. For example, the agility and transitory nature of cloud services make it easy for criminals to commit a crime, immediately destroy any traces of digital evidence, and migrate to another cloud service provider where they could do the same and leave without a trace (Dlamini et al., 2014). This makes it very difficult to identify criminals, let alone to acquire digital evidence. Furthermore, there is absolutely no need for the criminals to own any cloud infrastructure. For example, a criminal could subscribe to a cloud service provider that offers infrastructure as a service (IaaS) and easily create a virtual machine (VM) to perform their criminal activities and then destroy the VM and end their subscription (Dykstra and Sherman, 2012; Sibiya et al., 2012; Dykstra, 2015). Such a scenario makes it exceedingly hard for digital forensic investigators to gather credible digital evidence to help apprehend criminals in the cloud.

Based on the above and other difficulties, conducting an effective digital forensics investigation in the cloud environment has remained a big challenge. Even legal frameworks such as the United State of America’s (USA) Federal Rules of Discovery (Battaglia, 2016) that were designed with an inherent flexibility and applicability to technological developments, are not flexible enough to embrace the paradigm shift to the shared cloud environments (Araiza, 2011). Consequently, some researchers note that there is still insufficient research on tools, processes and methodologies required to acquire defensible digital evidence in the cloud (Chung et al., 2012; Harrington, 2012; Butterfield et al., 2018). For example, Butterfield et al. (2018) argue that existing tools for conducting digital forensics are manual. They therefore propose an automated digital forensic process.

Yet, some researchers (Martini and Choo, 2012; Daubner, 2018) argue that a lack of understanding of the complexity and challenges brought about by cloud computing greatly hinders the job of digital forensic investigators. Hence, this chapter takes the initiative to add a better understanding of the implications of cloud computing on the field of digital forensics to the body of knowledge.

The author of this thesis acknowledges that the rise of cloud computing introduces a digital forensic dilemma. In order to fully understand this dilemma, the remainder of this chapter is structured as follows: Section 3.2 aims to provide a common understanding of the definitions of digital forensics that are found in literature. It is vital to have a common understanding before we can tackle the implications of cloud computing for digital forensic investigators in Section 3.3. Section 3.4 outlines a harmonised taxonomy of the implications and challenges of conducting an effective digital forensics investigation in the cloud, so as to provide a high-level snapshot of the digital forensic challenges that have been identified by different researchers and investigators in the cloud. Section 3.5 introduces digital forensic readiness, its definition and benefits, and briefly discusses its suitability to investigations conducted in cloud computing environments. Section 3.6 concludes this chapter.

3.2 DEFINITION OF DIGITAL FORENSICS

Digital forensics is a fairly new disciple that emerged as a result of the ubiquity of digital devices and their increasing use in malicious activities (Reilly et al., 2011; Martini and Choo, 2012). The discipline of digital forensics has seen tremendous developments in the recent past in terms of investigation procedures, techniques and toolkits to support law enforcement agencies and other organisations to resolve disputes (Dykstra, 2015). Arguably, over the past couple of years, digital forensics has changed, evolved and adapted to keep up with rapidly changing technologies (Du et al., 2017; Bollo, 2017). Du et al. (2017) claim that digital forensic investigations must move towards cloud-based digital evidence processing. They argue that this will expedite the investigation process and free up investigators and law enforcement authorities to handle other tasks. Bollo (2017) agrees that digital forensics must keep up with technical changes and changing times.

Along with the technological changes came different definitions of digital forensics. The definition challenges were recognised and noted by Casey (2012), and it was only recently that a definition of digital forensics was adopted and standardised by the research community

(ISO/IEC 27043, 2015). This section discusses the different definitions of digital forensics suggested in literature and shows how the definition has changed over the years up until the point of standardisation. A number of disparate definitions of digital forensics existed in literature due to different researchers such as McKemmish (1999), DFRWS (2001), Reith, Carr and Gunsch (2002), Ruan et al. (2013), Raghavan (2013), ISO/IEC 27043 (2015) and Daubner (2018) having had differing views on what digital forensics pertains. The discussion in this section reflects the changing scope of digital forensics.

Martini and Choo (2012) cite one of the first definitions of digital forensics, i.e. the one suggested by McKemmish (1999, p.1). It refers to digital forensics as "a process of identifying, preserving, analysing and presenting of digital evidence in a manner that is legally acceptable". This definition defined digital forensics with regard to its investigation processes. Similar to that of McKemmish (1999) is a definition provided in Daryabar et al. (2013), which defines digital forensics as a process of preparing, acquiring, preserving, examining, analysing and reporting potential digital evidence. Daryabar et al.'s definition does not explicitly touch on the science component of digital forensics. However, both define it with reference to the different digital processes that investigators need to follow as they carry out investigations up to the point of presenting potential evidence. McKemmish's definition touches on the legal aspect of an investigation, whereas that of Daryabar et al. is silent about corporate or legal investigation. The legal component of conducting a digital forensic investigation in the cloud is key to establishing and ensuring that the collected digital evidence is legally admissible in court.

In 2001 a technical committee of the Digital Forensic Research Workshop (DFRWS) formulated a broad and comprehensive definition that defined digital forensics as "the use of scientifically derived and proven methods toward the identification, collection, preservation, validation, documentation, analysis, interpretation and presentation of digital evidence derived from digital sources for the purpose of facilitating or furthering the reconstruction of events suspected to be of a malicious activity or helping to anticipate unauthorised actions shown to be disruptive to planned operations" (DFRWS, 2001; Agarwal et al., 2011; Harrington, 2012; Raghavan, 2013). Reith et al. (2002) adopted the same definition, whereas Carrier (2003) opted for a more narrow focus and referred to digital forensics as the use of scientifically derived and proven methods of identifying digital evidence that verifies or contradicts an existing theory, and shows signs of tampering that can be used to facilitate or further the reconstruction of events in an investigation. This definition only focuses on the identification and analysis processes of digital forensics.

Kent et al. (2006) define digital forensics as a scientific procedure used to identify, classify, collect, evaluate and analyse potential digital evidence while maintaining a high level of integrity throughout the entire investigation process. This definition also defines digital forensics in terms of the investigation processes. However, the key to their definition is that it puts emphasis on basing digital forensics on scientific principles and procedures. Furthermore, it brings in the element of maintaining the integrity of digital evidence. This is a good definition that introduces the element of admissibility of digital evidence.

Zatyko (2007) provides one of the most intriguing analyses of the different definitions of digital forensics. Therein, digital forensics refers to “the application of computer science and investigative procedures for a legal purpose involving the analysis of digital evidence after proper search authority, chain of custody, validation with mathematics, the use of validated tools, repeatability, reporting and possible expert presentation”. This definition scopes the ‘scientific’ aspect to computer science. Despite the call for the digital forensic community to adopt Zatyko’s definition, it seems not to have attracted enough support for it to become a standard definition.

Eric Huber defines digital forensics as “the collection, examination and reporting of digital evidence in order to answer questions related to digital investigation or as an intelligence gathering task” (Kassner, 2011). This definition defines digital forensics based on the investigation procedure and ignores its scientific aspects, as was the case in Daryabar et al. (2013). Ngobeni et al. (2012) define digital forensics as a scientifically proven methodology for investigating digital devices that are suspected to have been involved in malicious activities for potential evidence that could be used in a court of law to help prosecute the perpetrators. This definition highlights the science behind digital forensics. However, it takes a narrow scope to focus only on the legal aspect of digital forensics, i.e. dealing with malicious activities. Digital forensics can also be used in corporate environments to deal with employees who breach policies and procedures. Even though such corporate cases may also lead to criminal cases; they are solely meant for dealing with malicious employees.

Raghavan (2013) defines digital forensics as an application of scientific principles and processes to the investigation of artefacts or digital devices in order to understand and reconstruct the sequence of events that must have transpired in a crime scene. This definition also emphasises the science component of digital forensics and adds the 'event reconstruction' aspect, which deals with 'repeatability' of activities that took place leading to an incident. This means that digital evidence must be able to tell the whole story and replay (repeat) what actually took place in the crime scene. This aspect deals with the comprehensiveness of digital evidence and adds to its admissibility in a court.

Finally, the ISO/IEC 27043:2015 standard came with a standardised definition of digital forensics that is derived from the one proposed by the DFRWS technical committee. The standard definition defines digital forensics as "the use of scientifically derived and proven methods toward the identification, collection, preservation, validation, documentation, analysis, interpretation and presentation of digital evidence derived from digital sources for the purpose of facilitating or furthering the reconstruction of events suspected to be of a criminal or malicious nature or assisting to anticipate unauthorized actions" (ISO/IEC 27043, 2015; Valjarevic, Venter and Petrovic, 2016; Daubner, 2018). This definition seems to be all-encompassing in that it covers most of the aspects that have been raised by other definitions, e.g. 'scientific principles, procedures and processes', 'event reconstruction' and 'proven methodologies'. The aspect of 'proven methodologies' is also emphasised in Shamsi et al. (2016). It would however seem that this definition does not leave room for creativity in sourcing and analysing digital evidence with regard to 'proven methodologies'. It nevertheless emphasises repeatability and reliability, which adds to the admissibility discussions. Since it is standardised and a good starting point, the digital forensic community can hopefully now adopt it widely to have a common understanding. This study therefore also adopts the definition provided in the ISO/IEC 27043:2015.

The next section elucidates the actual implications of cloud computing for the discipline of digital forensics.

3.3 RISE OF CLOUD COMPUTING ELICITED A DIGITAL FORENSIC CRISIS

The rising interest in the use of cloud computing services presents both opportunities for criminal exploitation and a host of challenges for digital forensic investigators (Martini and Choo, 2012). For example, Casey (2012) argues that criminals could exploit cloud storage services like Dropbox to store incriminating evidence, launch cyber-attacks and break the strongest encryptions keys. The LinkedIn case where encrypted user credentials were stolen, decrypted and then sold in the Darkweb might be one case in point for criminals using the cloud computing power to break encryption keys (CSA, 2017). Surely, cloud computing makes it easier for criminals to store illegal and incriminating files (e.g. child pornography videos and pictures) in third party cloud storage service providers. The same cloud setup could also make it extremely difficult for digital forensic investigators to seize data that could be used as potential evidence (Dykstra, 2015). This is more so if the child pornographic data is stored in data centres that span multiple jurisdictions covering areas where it is not even considered illegal. Besides, even if law enforcement agencies could discover such contraband images and videos, it would be difficult for them to terminate such services in a third party's cloud servers that are located in a foreign country. Obtaining search warrants and getting the necessary consent from foreign law enforcers in urgent and extremely time-sensitive situations may prove to be implausible. It is on this premise that the US government working with other qualifying governments passed the "Clarifying Lawful Overseas Use of Data – the CLOUD" Act (Hatch, Coons, Graham and Whitehouse, 2018). The CLOUD Act comes at the right time when the Safe Harbour framework proved insufficient in ensuring high-level data protection as stipulated and mandated in the GDPR (Monteleone and Puccio, 2017; Bu-Pasha, 2017). The Privacy Shield was mandated after the Safe Harbour was invalidated (Monteleone and Puccio, 2017). However, the focus of the 2016 Privacy Shield is on privacy of personal data. The CLOUD Act is aimed to solve the problem of cross-border data requests and to facilitate timely access to electronic data in custody, control or possession of service providers across borders. Global Research has since labelled the CLOUD Act a dangerous piece of legislation (Global Research, 2018; Hickey, 2018; Matsakis, 2018).

Aggravating this issue, is that most cloud clients would consider encrypting their data (more so if it is illegal) before moving it to the cloud to try and mitigate the risk of it being accessed by unauthorised third party users, i.e. criminals, investigators, malicious or prying service providers and co-resident clients. This could lead to greater complexity and long delays in obtaining digital evidence. The investigators, malicious or prying service providers and co-resident clients are included because nowadays it is not only the criminals that are suspect, but everyone. Not even service providers could be trusted with their clients' data. For instance, Google's cloud offerings such as Gmail, YouTube and Google Docs have provisions in their terms of service (i.e. "By submitting, posting or displaying the content you give Google a perpetual, irrevocable, worldwide, royalty-free, and non-exclusive license to reproduce, adapt, modify, translate, publish, publicly perform, publicly display and distribute any Content which you submit, post or display on or through, the Services.") that give them a non-exclusive licence to access and exploit all the contents provided by their clients (Garon, 2012).

In short, Google as a cloud service provider has the right to purposely modify its clients' data and leave the aggrieved clients with no recourse and no legal right to recover any lost or damaged data, regardless of the cause of damage and even if the provider was grossly negligent. Hence, Calloway (2012) argues that ingrained in what they refer to as "clickwrap agreements" of prying cloud service providers are the limitations on liability provisions, absolute safe harbours and hidden icebergs for the multi-millions of unsuspecting cloud clients. Such clients habitually click "I Agree" and in the process relinquish their legal rights.

For the above reason and in order to gain a competitive advantage, most cloud service providers might consider encouraging their clients to encrypt their data before moving it to the cloud (Taylor et al., 2010). Such providers could further encrypt their clients' data that has already been encrypted as mandated by law. For example, Spideroak, a cloud storage service provider – a direct competitor to Dropbox – encrypts its clients' data and then hands the encryption keys back to the client to guarantee that it (Spideroak) cannot decrypt it to inadvertently disclose clients' data to anyone – including investigators and law enforcement agencies (Garon, 2012). Even though this might add an extra layer of protection to the client's data, it creates a challenge and increases complexity for an investigator. For instance, it would not be enough for an investigator to only get the client's decryption keys. The investigator will also be required to obtain the decryption keys of the cloud service provider, who might be unwilling to hand them over (Sibiya et al., 2012). This can significantly increase the cost of an investigation, to such an extent that it would defeat the cost-benefit analysis, which is also known as the "proportionality doctrine" in US law (Harrington, 2012). The process might be easier if both sets of keys are with the client.

The US's "proportionality doctrine" is defined as a cost-benefit analysis of an investigation (Harrington, 2012; Ajoy, 2012). As such it is applied to limit investigators to acquire inaccessible digital evidence only when the benefits of doing so far outweigh the expense. Therefore, within a cloud environment, for a less important case, investigators could only do their investigations on the client's side. Only if the case is unusually important, they would move to the cloud service provider's side. Digital forensic investigators should work with a legal team to define and limit the scope of an investigation to that which is reasonably possible within the given budget and timeframe. This is a very important element of investigating crimes committed in the cloud where there could be massive potential digital evidence with residual data belonging to other tenants. On the other hand, crucial digital evidence might be overwritten as clients switch to other service providers. This raises the challenge of using evidence with gaps. This might still be a problem in non-cloud environments. Hence, it resurfaces again in the cloud and digital forensic investigators still need to deal with presenting evidence with gaps.

As discussed above, cloud computing might not be a new phenomenon, but its implications for acquiring and preserving digital evidence for the resolution of civil disputes and the prosecution of alleged criminals can never be ignored (Mason and George, 2011). Consequently, digital forensics faces one of its greatest challenges with the rise of cloud computing (Martini and Choo, 2012; Park, Kim et al., 2018; Ruan et al., 2011; Zimmerman and Glavach, 2011). In support, Yadav, Ahmad and Shekhar (2011) and Quick and Choo (2018) argue that digital forensics faces an ever-growing crisis with the emergence of cloud computing. Harrington (2012) argues that technological trends such as cloud computing and social media are pushing legal and ethical boundaries to the limits and this necessitates novel approaches to digital forensic investigations. For example, when an investigator goes beyond the formal process of acquiring digital evidence to befriend a suspect under false pretence on social networking media to gain access to the suspect's private information, this could easily lead to an "invasion of privacy, intrusion upon seclusion or unreasonable warrantless search or other tort liability" (Harrington, 2012). This could easily render all such evidence inadmissible in court as reflected in the United States vs Maynard case (Harrington, 2012). In this case, the US Court of Appeal had to decide if digital evidence obtained by the investigators through a warrantless search of a global positioning system (GPS) device was admissible in court. The court opined that "the defendant had a reasonable expectation of privacy in the sum of his movements, even though he had no expectation of privacy in his individual movements exposed to the general public*"* and dismissed the evidence on the grounds that it was obtained using illegal methods that violated the privacy of the defendant (Harrington, 2012).

Furthermore, there is little guidance on how to acquire digital evidence and conduct digital forensics in cloud environments (Martini and Choo, 2012). Martini and Choo (2012) argue that there are no guidelines that are specific to dealing with digital evidence in cloud environments. Hence, researchers are arguing that cloud computing environments are making it more difficult for investigators to acquire and analyse digital evidence to the same standards as are expected in traditional server-based systems. This attributes to a certain extent to the difficulty in establishing what data is stored or processed by what applications on what end-point device (Taylor et al., 2011). However, cloud computing can also facilitate the work of a digital forensic investigator, for example, in collecting and storing potential digital evidence before a lawsuit or policy breach. This thesis indirectly shows how to do this.

The tension between cloud computing and digital forensics is also compounded by the complexity of acquiring evidential data from the cloud. This complexity stems from the fact that traditional tools, processes, procedures and methodologies for obtaining legally defensible digital evidence cannot be superimposed on this environment. Existing tools or methods are simply not scalable or fit-for-purpose (Martini and Choo, 2012) in this environment, and if used, could easily result in digital evidence of questionable quality that could well be deemed inadmissible in the courtroom. For example, digital forensic investigators do not have the ability to physically acquire the suspected digital device for a bit-by-bit imaging in the virtual cloud environment where disks, memory and CPU are virtualised and shared between multiple clients (Harrington, 2012). In this instance, the complexity comes in isolating the suspect's data and ensuring that all data and applications that belong to other co-resident clients are not captured or disrupted. Digital evidence containing any residual data, which is data that belongs to other co-resident clients, could bring unnecessary disrepute to the credibility of the tools used for collection. This could easily nullify a serious case in a court of law failing to convict the accused. This particular problem of acquiring residual data, which may be viewed as intrusive, has been partly addressed by the United Kingdom's Criminal Justice and Police Act of 2001 (Mason and George, 2011). This Act gives legal authorities additional power to seize evidence with its residual, i.e. only if it is practically impossible for the two to be separated. Hence, this study argues that the rise of cloud computing brought along a digital forensics crisis. In agreement is Fred Cohen & Associates' (2009) assertion that cloud computing might yield a great business model with all its benefits (i.e. cost savings, flexibility, scalability, load balancing and agility), but it creates a completely unworkable legal and investigation model for digital forensic investigators and law enforcement agencies. Such complexity can be demonstrated by the case *US Government vs Microsoft* which started in 2014 (Matsakis, 2018) and resulted in the CLOUD Bill – which has huge ramifications for the future of cloud computing (Hatch et al., 2018).

The next section gives a high-level overview of the digital forensic challenges in the cloud. The main aim is to elucidate the specific digital forensic challenges of cloud computing. This section ends by trying to give a holistic picture of these challenges.

3.4 DIGITAL FORENSIC CHALLENGES IN THE CLOUD

This section outlines at a high level the challenges of conducting a digital forensic investigation in a cloud computing environment. A number of researchers have made attempts to identify some of the digital forensic challenges that arise in the cloud computing environment.

3.4.1 Digital Forensic Challenges in the Cloud: A Literature Review

For instance, Ruan et al. (2011) argue that the challenges of digital forensics in a cloud computing environment manifest across three dimensions, i.e. organisational, legal and technical dimensions. Therein, it is argued that addressing the challenges on all three dimensions would facilitate the establishment of a digital forensic capability. The challenges discussed in Ruan et al. (2011) include multi-jurisdiction; multi-tenancy; service level agreements; forensic data acquisition and collection; elastic and live digital forensics; evidence segregation; and working with virtual environments. The work of Ruan et al. (2011) does not acknowledge that different cloud deployment models (e.g. public, private, community and hybrid) along with their different service models (SaaS, IaaS and PaaS) have unique challenges.

The work of Birk (2011) and of Birk and Wegener (2011) focuses on the technical challenges of conducting a digital forensics investigation in the cloud. Birk (2011) covers public and private deployment models and acknowledges that the amount of potential evidence available to the digital forensic investigator differs between the different cloud deployment and service models. Therein, it is argued that digital forensic evidence could either be at rest (in storage disk space), in motion (file transfer over the network) or in execution (in memory) and it could be sourced from either a virtual cloud instance or network layer or client system (Birk, 2011). In support, Sibiya et al. (2012) argue that digital forensic evidence in the cloud is likely to be split or partitioned and stored in geographically distributed jurisdictions, which would make it hard to locate. Consequently, conducting a digital forensic investigation in each case will also differ. This work goes on to show what is possible and not possible with the different service models. For example, Birk (2011) argues that the ability to access virtual instances for gathering digital forensic evidence is limited in SaaS and PaaS. In the SaaS service model, the client can rely on logs provided by the cloud service provider (CSP). This is because only the CSP has control of the underlying infrastructure. Hence, it is concluded that cloud clients do not have the chance to perform an effective digital forensic investigation in the SaaS service model. According to Birk (2011), it is much simpler for cloud clients to conduct a digital forensic investigation on the PaaS and IaaS service models than on the SaaS service model. Ali et al. (2018) support Birk’s findings on the complexity of a digital forensic investigation on SaaS as compared to the other service models. This is generally true for most cases of SaaS where the clients do not have direct access to the underlying infrastructure on which the applications are running.

According to Hare-Brown and Douglas (2011) the potential challenges include the following:

- A lack of well-written contracts that support digital forensic investigation: This work goes on to say that there are currently no standard operating procedures between clients and CSPs that stipulate comprehensive rules of engagement. Although this might have been the case back in 2011, there has been improvement in that the UK is for instance mandating digital forensic readiness for small enterprises dealing with the UK government (Moussa, Ithnin and Zainal, 2018).
- Clients' reliance on CSP goodwill (Ali et al., 2018): The challenge comes when the CSP is reluctant or unwilling to support the digital forensic investigation team. Negotiations in this case could lead to significant time lags, which could easily risk evidential data being corrupted or lost before it could be captured for analysis.
- The discovery and handling of digital forensic artefacts with disclosure clauses as part of contractual obligations (Hook, 2018): The challenge here is that even if most CSPs are legally obligated and duty bound to disclose security incidents to the affected parties, they may choose not to do so because of the fear of losing their customers.
- Cross-jurisdictional issues: The argument therein relates to that fact that it is generally not legally permissible for law enforcement authorities to access cloud-based systems that lie beyond their jurisdiction. Moreover, serving warrants and court orders internationally is regarded as a time-consuming and costly exercise.
- Timescale: Huge time lags are introduced by the bureaucracy involved in establishing international law enforcement cooperation and collaborations. This work asserts that long timelines greatly diminish the hope of discovering and recovering digital evidential data from the CSPs. This is an issue that could be even trickier with CSPs having short data and log retention policies.
- Proportionality issue (Harrington, 2012): This is a focus of today's courts. It is stipulated that digital forensic investigators should acquire or seize or capture ONLY

evidential data that is pertinent or relevant to the investigation and nothing more or less. In terms of capturing snapshots of virtual machines in a multi-tenant cloud environment, how can we ensure that digital forensic investigators capture snapshots with ONLY the relevant evidential data, without any residual data from other clients who might not be involved in the investigation?

- Considerations associated with the presentation of digital forensic evidence (Orton, Alva and Endicott-Popovsky, 2015): It is a known fact that evidential data must conform to certain standards (e.g. secure preservation; a clear and chronological chain of custody; synchronised timestamps across international boundaries and different time-zones) and if not, it could easily become inadmissible in court.

Hare-Brown and Douglas (2011) conclude that the challenges are not insurmountable; they could be solved with solid technical and legal support, as well as thorough cooperation and collaboration with all the relevant stakeholders. The list of legal matters as provided by them is not exhaustive. More work on the legal requirement of conducting a digital forensic in the cloud can be found in Orton et al. (2015). However, the above list illustrates the point that the study in hand must directly or indirectly consider legal matters to improve the admissibility of digital evidence.

Reilly et al. (2011) argue that the major challenges of cloud computing emanate from the fact that CSPs have not yet come up with concrete ways of implementing measures to ensure a digital-forensic-ready cloud computing environment. For example, D’Orazio and Choo (2018) were compelled to use techniques that circumvent security mechanisms to be able to facilitate the collection of digital forensic evidence from iOS cloud applications. However, the iOS cloud ensures the integrity of synced files for digital forensic investigations (Ahmed and Xue, 2018). Furthermore, digital forensic experts have not yet come up with clear principles and procedures that could be effectively used to conduct a digital forensic investigation in the cloud (Reilly et al., 2011). Reilly et al. assert that the main challenge is related to data acquisition. They mentioned that search and seizure procedures used in conventional digital forensic investigation are impractical in the cloud environment. This is because client data resides in remote data centres of the CSPs. This goes hand-in-hand with the challenge of maintaining the chain of custody relating to the acquisition of potential evidence (Reilly et al., 2011). Given these two challenges, it becomes even more difficult to put together the pieces of acquired evidence and events to be able to reconstruct the crime scene and create the timelines.

Furthermore, Reilly et al. (2011) argue that cloud computing environments do not allow for important artefacts that could possibly hold crucial evidence like registry entries, temporary files and memory cache to be captured and preserved for a digital forensic investigation. One other important challenge that has been missed by most of the covered literature is that of presenting technical digital forensic evidence to the jury. Presenting digital forensic evidence to the jury has been hard with traditional digital forensics. It is only going to get worse in the complex cloud computing environment.

The work of Taylor et al. (2011) examined the legal aspect of digital forensic investigation in cloud-based systems. This work focuses on the complex processes of digital forensic evidence acquisition and analysis, which they claim are more complex in public and hybrid cloud models as opposed to private cloud models. Taylor et al. (2011) discuss and compare the challenges according to differences on the deployment models. They also outline the challenges of digital forensic evidence acquisition and analysis as follows:

- Multiple jurisdiction (Perloff-Giles, 2018)
- Encryption of data before it gets migrated to the cloud (Dlamini et al., 2017)
- Lack of established digital forensic guidelines on cloud-based systems (Dykstra, 2015; Farina et al., 2015; Moussa et al., 2018)
- Difficulty of establishing a chain of custody (Garcia, 2014)
- Difficulty in the recovery of digital forensic evidence (Casey, 2011)
- Lack of established method for evidence acquisition (Moussa et al., 2018)
- Complexity of identifying potential digital forensic evidence (Martini and Choo, 2012)
- Imaging data in the cloud possibly not being practical
- Multi-tenancy of clients’ data and applications (Mason and George, 2011)

Taylor et al. (2011) also alluded to the most overlooked but important aspect of providing and presenting digital forensic evidence in court. The issue of tracking malware that originates from cloud-based systems was argued to be one of the most complex tasks for digital forensic investigators. This is presumably because of the transient nature of cloud services. Even more complex is the issue of tracking down the effects of malware with malicious cloud clients claiming ignorance of stealth and evasive malware running on their systems. Presenting evidence to the jury in support of such cases can be difficult. This could easily reduce serious sentences to nothing and, in the process, let criminals off the hook.

From the above discussion, it is clear that there are still many grey areas with regard to digital forensic investigations in the cloud. This is an indication that it will definitely require considerable amount of time, effort and money to overcome these challenges. The next subsection provides a taxonomy depicting a high-level snapshot of the challenges.

3.4.2 Digital Forensic Challenges in the Cloud: A High-Level Taxonomy

Figure 3.1 provides a high-level taxonomy of the challenges of digital forensics in the cloud.

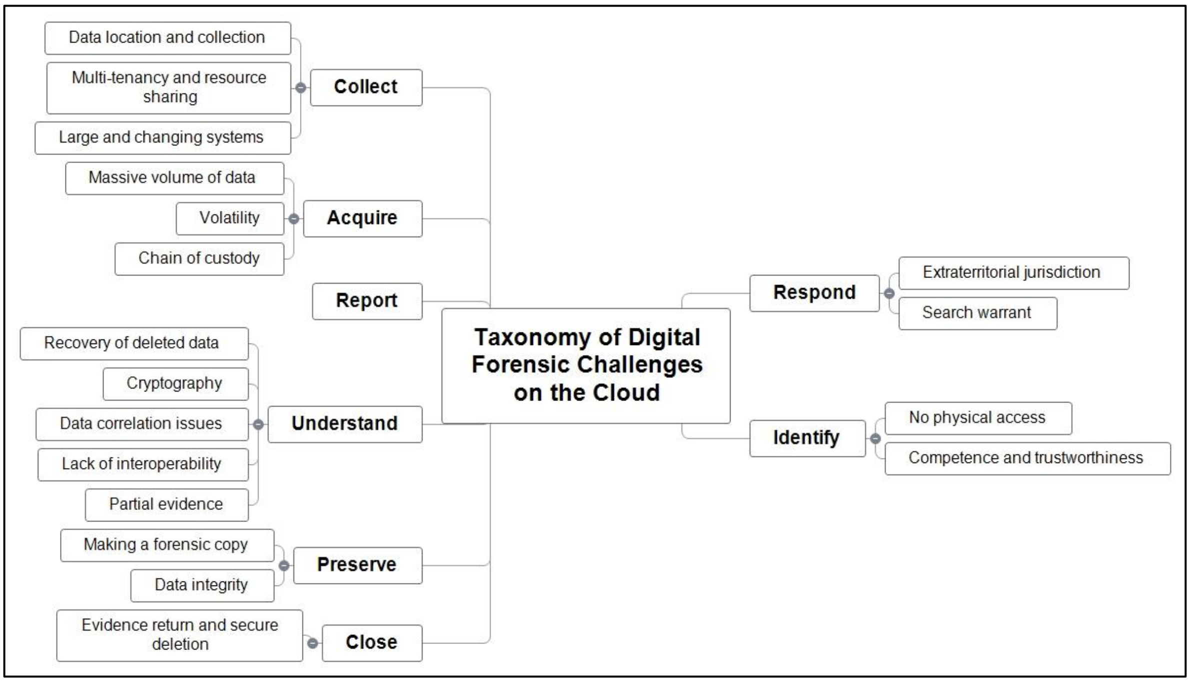

Figure 3.1:- Taxonomy of Digital Forensic Challenges in the Cloud (Lopez et al., 2016)

Lopez et al. (2016) classify the challenges across the digital forensic investigation processes, namely *identify*, *respond*, *collect*, *acquire*, *understand*, *preserve*, *report* and *close*. The *identify* process talks to the challenge of physically acquiring digital evidence with regard to competence of investigators and trustworthiness of the identified data. The *respond* process discusses issues around the difficulty of obtaining and serving search warrants in multi-jurisdictions. The *collect* process focuses on data location, which relates to the *identify* process. This process also touches on multi-tenancy and resource sharing, which may lead to collection of residual data belonging to tenants that are not involved.

The difficulty of selective collection, compared to collecting everything, is key here. The last part refers to dynamic and large systems where there is so much to collect and the rapidly changing environment makes it hard to collect digital evidence. The *acquire* process speaks to the massive volumes and volatility of data that must be collected and the difficulty of maintaining a proper chain of custody. The *understand* process deals with understanding partial evidence, i.e. evidence with gaps from selective acquisition. This process also focuses on recovery of deleted data and its correlation issue. It furthermore focuses on cryptography and lack of interoperability among cloud systems.

The preserve process deals with data integrity and copies of digital evidence. This relates to proper handling of the digital evidence's chain of custody. The *report* process refers to presenting digital evidence in court or in a disciplinary hearing. Finally, the *close* process focuses on challenges of dealing with evidence after a case is closed. This relates to proper sanitisation of media that hold digital evidence for secure deletion and return. Further details on these can be found in Lopez et al. (2016).

Figure 3.2 provides a similar, yet more comprehensive taxonomy of digital forensic challenges in the cloud. Dykstra (2015) divides the cloud forensic challenges into eight categories: *incident first responders; data collection; legal analysis; anti-forensics; architecture; role management; standards*; whereas Lopez et al. (2016) use digital forensic processes for the categorisation. Each of these categories have different subcategories. For example, the *role management* category has an *identity management* subcategory, which further has sub-subcategories, i.e. *difficulty of criminal attrition in the cloud* and *errors in cloud management*. The *legal* category is further divided into four subcategories, i.e. *contract SLA*, *jurisdiction*, *privacy* and *root of trust.*

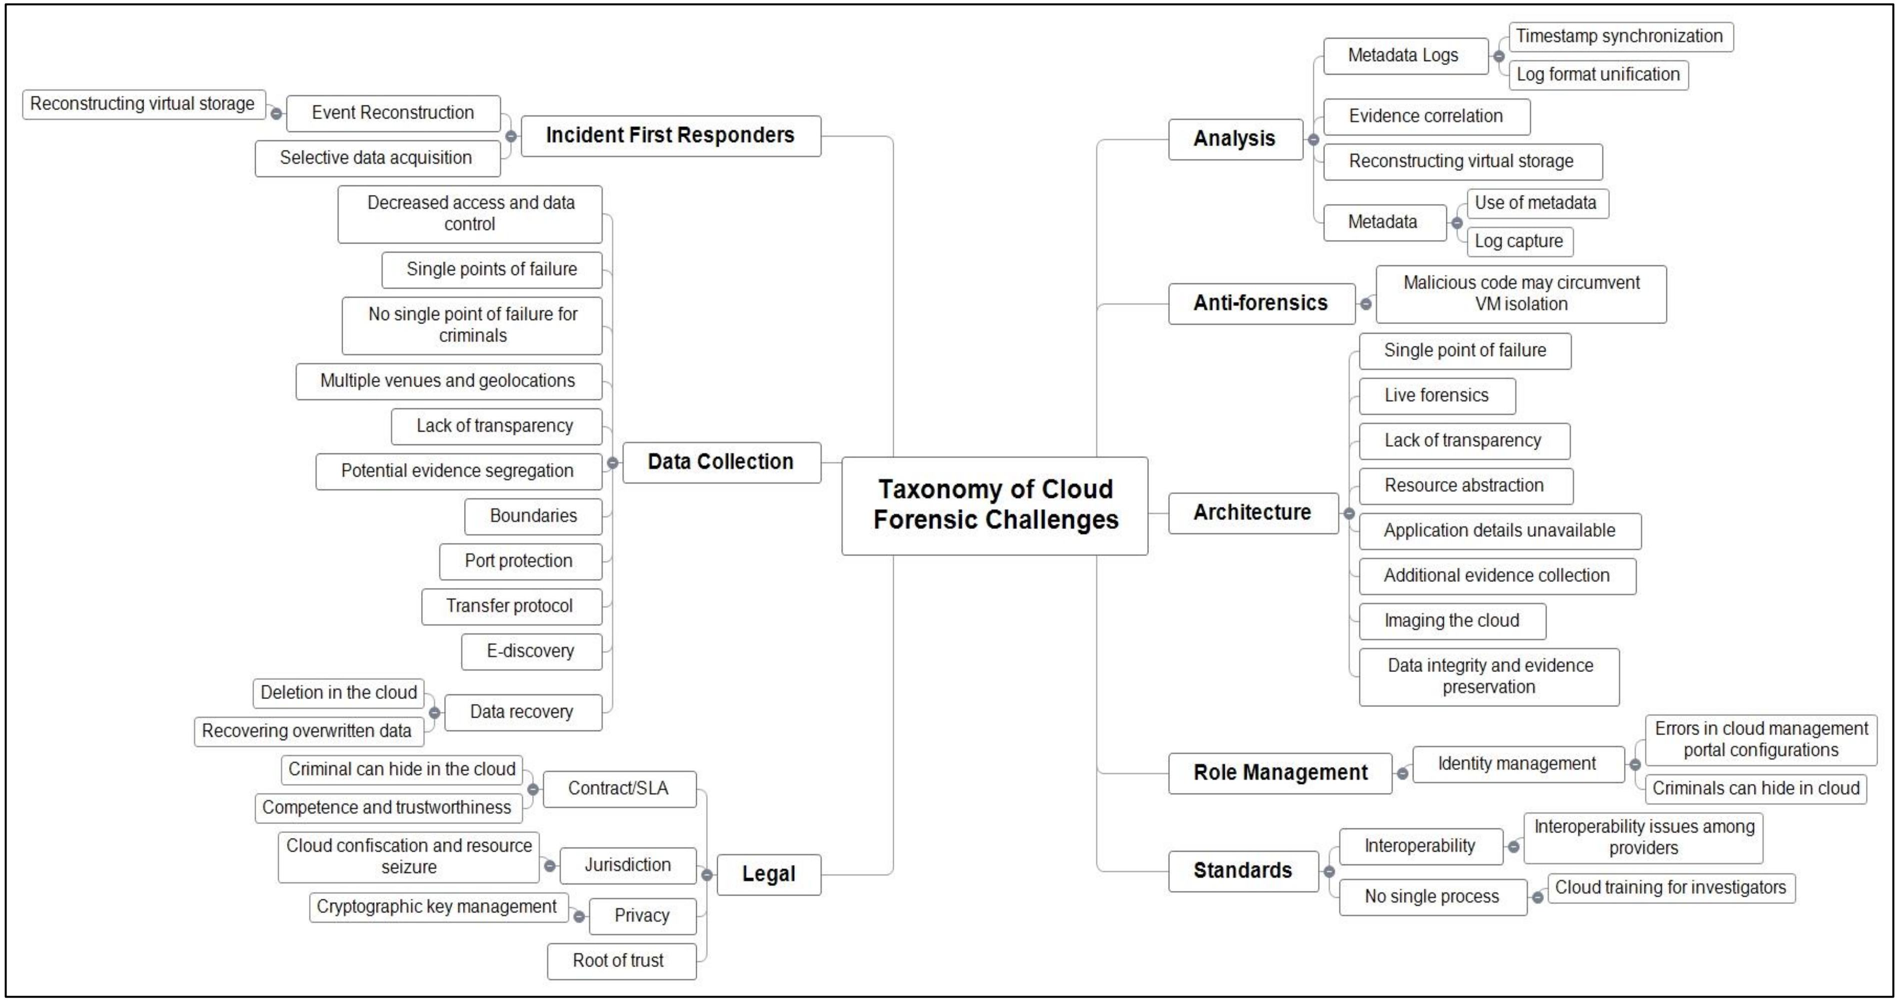
Taxonomy of Cloud Forensic Challenges
Incident First Responders
Event Reconstruction
Reconstructing virtual storage
Selective data acquisition
Data Collection
Decreased access and data control
Single points of failure
No single point of failure for criminals
Multiple venues and geolocations
Lack of transparency
Potential evidence segregation
Boundaries
Port protection
Transfer protocol
E-discovery
Data recovery
Deletion in the cloud
Recovering overwritten data
Legal
Contract/SLA
Criminal can hide in the cloud
Competence and trustworthiness
Jurisdiction
Cloud confiscation and resource seizure
Privacy
Cryptographic key management
Root of trust
Analysis
Metadata Logs
Timestamp synchronization
Log format unification
Evidence correlation
Reconstructing virtual storage
Metadata
Use of metadata
Log capture
Anti-forensics
Malicious code may circumvent VM isolation
Architecture
Single point of failure
Live forensics
Lack of transparency
Resource abstraction
Application details unavailable
Additional evidence collection
Imaging the cloud
Data integrity and evidence preservation
Role Management
Identity management
Errors in cloud management portal configurations
Criminals can hide in cloud
Standards
Interoperability
Interoperability issues among providers
No single process
Cloud training for investigators

Figure 3.2:- Taxonomy of Cloud Forensic Challenges (Dykstra, 2015)

The *contract SLA* category refers to *competence and trustworthiness*, and to the fact that *criminals hide their identity* and cover their tracks in the cloud. The *jurisdiction* category covers challenges around cloud confiscation and digital evidence seizure. The *privacy* category refers to cryptographic key management that ensures the privacy of tenants. The *data collection* category is similar to the *collect* category in Figure 3.1 (Lopez et al., 2016). The *incident first responder* category captures challenges on event reconstruction and selective data acquisition. The *architecture* category deals with challenges around resource abstraction, imaging, live forensics, additional evidence collection, digital evidence integrity and preservation. *Anti-forensics* category focuses on the challenges of having malicious code that circumvent virtual machine isolation. An *analysis* category focuses on metadata and its logs, evidence correlation and virtual storage reconstruction. More details on each of these categories, subcategories and sub-subcategories can be found in Dykstra (2015). Figures 3.1 and 3.2 present a high-level overview of digital forensic challenges in the cloud. Substantial preparations and planning are vital to address these challenges and ensure effective digital forensic investigations in the cloud.

The author believes that this could be best achieved by first preparing the cloud computing environment to become ready for digital forensics. This points to a need for research efforts that focus more on a proactive approach towards digital forensics and it takes our discussion towards digital forensic readiness. The next section argues that digital forensic readiness could be one of the ways to help address most of the identified challenges and prepare the cloud environment for an effective digital forensic investigation.

3.5 DIGITAL FORENSIC READINESS

The rise of critical IT infrastructures and cloud computing raises new digital forensic challenges. Digital forensics has been traditionally viewed as a reactive approach for investigating incidents after they occur (Alharbi et al., 2011; Alharbi, 2014; Kigwana and Venter, 2018). There used to be insufficient research to help organisations proactively plan and prepare on how to respond to incidents when they occur. Hence, nowadays a considerable amount of digital forensic research (Moussa et al., 2018; Park, Kim et al., 2018) focuses on a proactive approach towards the collection and preservation of digital evidence prior to an investigation. This is mainly because traditionally reactive approaches of collecting digital evidence are insufficient and not scalable in a cloud computing environment (Chung et al., 2012).

The characteristics of cloud computing, such as agility, multi-tenancy, shared nature and dynamism have heightened the need for organisations to consider a digital forensic readiness (DFR) capability of collecting potential digital evidence and preserving it in a legally and forensically sound manner. The next section provides a number of definitions of DFR that are found in existing literature.

3.5.1 Definition of DFR

Quite a number of definitions for DFR are found in literature. However, Tan (2001) provides an earlier description that defines DFR with respect to maximising an environment's capability of collecting digital evidence, whilst at the same time minimising the cost of doing so during an incident response. Rowlingson (2004) bases his definition on the one found in Tan (2001). Rowlingson (2004) defines DFR as "the ability of an organisation to maximise its potential to use digital evidence whilst minimising the cost of an investigation". Key to this definition are the two phrases 'maximise the potential use of digital evidence' and 'minimising the cost of an investigation'. These two are to be noted as they become key in newer definitions.

Ten years after the definition of Rowlings (2004), Mouhtaropoulos, Li and Grobler (2014) defined DFR as a "pre-incident plan that deals with an organisation's ability to maximise digital evidence usage and anticipate litigation". Therein, DFR is defined as the pre-incident plan within a digital forensic investigation life cycle that deals with digital evidence identification, preservation, storage, analysis and use, whilst minimising the cost of an investigation (Mouhtaropoulos et al., 2014). The aim is to proactively capture, handle and manage digital evidence in order to provide for a timely and cost-effective investigation. The two key phrases in Rowlingson (2004) come back again. However, this time they are substantiated by the different processes of conducting a digital forensic investigation.

ISO/IEC 27043 (2015) defines DFR as the process that assures that organisations have made the necessary, prudent and strategic preparations for accepting potential events of an evidential nature to be used in a digital forensic investigation, prior to an incident. This standard clearly defines readiness as a process that occurs before an incident (Kebande and Venter, 2016). It also defines pre-incident strategies to ensure proper systems and trained employees to deal with an incident before and when it occurs. However, the standard puts more emphasis on the standardised digital forensic readiness investigation process, without any mention of the environment in which it is to be applied. Similar to Tan (2001) and Rowlingson (2004), this standard outlines the goals of a DFR capability as follows (ISO/IEC 27043, 2015):

- To maximise the potential use of digital evidence
- To minimise direct or indirect costs of digital investigations
- To minimise interference of business operations
- To improve the level of information security

According to ISO/IEC 27043 (2015), and Kebande and Venter (2016), the above goals are accomplished using three processes of a DFR, namely planning, implementation and assessment. The planning process includes scenario definition; source identification; pre-incident collection, storage and handling of potential digital evidence plans; pre-incident analysis plan; and system architecture definition. The implementation process deals with implementation of each of the subprocesses in the planning process. It ends with implementation of an incident detection. Implementation of the analysis subprocess leads to the assessment process, which assesses the implementation in the implementation process to check if it is doing its individual tasks well. The assessment process feeds back to all the previous processes and subprocesses. The ISO/IEC 27043:2015 standard seems comprehensive and shows how these readiness processes link to the actual digital forensic investigation processes. Since the standard resonates with the proposed 3-tier model in this thesis, the thesis uses it (ISO/IEC 27043:2015 standard) as a baseline.

The INFOSEC Institute defines DFR as having an appropriate level of capability to be able to preserve, collect, protect and analyse digital evidence so that it can be used effectively in legal matters, in security investigations, in disciplinary proceedings, in a labour tribunal and/or in a court of law (INFOSEC Institute, 2016). According to the INFOSEC Institute, the main aim of a DFR capability is to optimise the time and cost of conducting a digital forensic investigation in the hope of having good results. Compared to the first two definitions, this definition brings in the ‘where is the digital evidence to be used?’, i.e. “effective use in legal matter, in security investigation, in disciplinary proceedings, in a labour tribunal and/or court of law”. This definition shows that digital evidence collected from a DFR capability can be used in a number of areas.

Therefore, and based on the above definitions, this study defines DFR as a proactive process to identify, collect, acquire, preserve and store potential digital evidence prior to a legal dispute, security investigation, disciplinary proceedings, labour tribunal and/or court of law; in order to optimise on the turnaround time, effort and cost of conducting an effective digital forensic investigation. This is the definition that the author applies throughout the research project and thesis.

Having defined DFR, it becomes essential to discuss why it is necessary and what some of its benefits are. Hence, the next section delves exactly into that. It discusses the benefits of having a DFR capability and is geared towards those organisations that are making or planning a move towards cloud computing.

3.5.2 The Benefits of a DFR Capability

The benefits of having a DFR capability (Endicott-Popovsky et al., 2007; Rowlingson, 2004; Sule, 2014; INFOSE Institute, 2016) can be summarised as follows:

- Prepare for the potential need for digital evidence in advance (Endicott-Popovsky et al., 2007): A DFR capability makes digital evidence readily available when requested. Furthermore, it could help avoid digital evidence being routinely deleted, based on data retention policies.
- Maximise the potential of having admissible digital evidence (Rowlingson, 2004): A DFR capability ensures that digital evidence does not have gaps that occur as a result of the culprits making an attempt to erase their tracks when they discover that they are under investigation. It also avoids improper handling of digital evidence.
- Optimise the cost, time and effort of conducting an investigation (Sule, 2014): A DFR capability cuts down the time it takes to do an investigation because the digital evidence is already available for the investigators when required. This has a direct impact on and significantly reduces the cost and effort of doing so.
- Minimise business disruptions (INFOSE Institute, 2016): A DFR capability ensures that business operations are minimally interrupted (or not at all) as investigators collect digital evidence. Investigations can go on without interfering with business operations and processes as the digital evidence is collected prior to the investigation and ready by the time a lawsuit or investigation is instituted.

- Determine the attack: A DFR capability ensures that multiple probes are strategically placed in an active mode to continuously detect symptoms and dynamics of incidents, preferably before they occur. This would potentially help to proactively stop an incident from happening in the first place, i.e. if the DFR raises insightful alerts at the right time to the right response team. Should it not be able to stop an incident from happening with the alerts, a DFR would help investigators to reconstruct the events leading to an incident.
- Reduce the cost of regulatory or legal requirement for data disclosure (Sule, 2014): A DFR capability will make it easy for organisations to disclose digital evidence for compliance requests in terms of data protection legislation. For example, the General Data Protection Regulation (GDPR) mandates European companies dealing with personal and special data to disclose data breaches within 32 hours of an incident (GDPR, 2016). Without a good DFR capability, this is not feasible.
- Add value to existing business processes (Sule, 2014): A DFR capability provides extra value to incident response; business continuity and disaster recovery; monitoring and logging; data retention; and crime prevention efforts.
- Demonstrate due diligence, good corporate governance and regulatory compliance with legal mandates (INFOSE Institute, 2016; Endicott-Popovsky et al., 2007; Rowlingson, 2004): A DFR capability demonstrates due diligence in the fight against digital crime and illustrates good corporate governance and compliance with mandates that require organisations to collect user logs to monitor their activities.
- Deter malicious insiders from covering their tracks of criminal activity (Sule, 2014): If appropriately relayed to all employees, a DFR capability can act as a deterrent to minimise a malicious-insider business risk. Knowing that user activities are monitored by the DFR capability, malicious users would tend not to do malicious activities for the fear of being caught.
- Provide support for insurance discounts (Sule, 2014): A DFR capability is a reflection that an organisation is actively doing something about managing its risk profile. Hence, a demonstration of the DFR capability might assist organisations to pay smaller insurance premiums.

Organisations could use a DFR capability as part of an overall enterprise risk management strategy. This could help to manage the impact of key business risks by providing digital evidence to detect malicious activities. It is therefore essential that organisations should proactively collect and preserve potential digital evidence in a legal and forensically sound manner. Digital forensic readiness has become a business requirement for most organisations, more so those that are considering a move to cloud computing. The next section positions DFR in the cloud and provides a discussion on some of the main drivers.

3.5.3 DFR Positioning in the Cloud

This section discusses some of the forces that are driving most organisations towards considering a DFR capability, more especially as they move to embrace the cloud. For this study, the main drivers (Dlamini et al, 2014) include the following:

- Corporate governance and legal requirements (Mouhtaropoulos et al., 2011)
- Policy (Park, Akatyev et al., 2018)
- Costly business disruptions
- The duty to gather and preserve digital evidence (Taylor, 2012)
- Strict court obligations to ensure digital evidence admissibility (Casey, 2011)

Each of these drivers is unpacked in the next subsections.

3.5.3.1 Corporate Governance and Legal Requirements

Apart from preparing for a reasonably anticipated legal litigation or dispute, organisations that are planning a move to the cloud could use a DFR capability as part of an all-encompassing incident response procedure to demonstrate good corporate governance and compliance with legal and regulatory mandates. For example, according to the King III report on corporate governance, the effective utilisation of digital forensic tools can enable cloud-bound organisations to prove their due diligence with respect to good governance (Grobler et al., 2010b).

In terms of the legal and regulatory mandates, the ISO 27001/2 standard (for instance) has a provision that makes it essential for organisations, including those that are moving to the cloud, to identify and gather potential digital evidence that is complete, admissible and concrete, prior to a litigation or dispute. This is to determine the root cause of an incident and make means to prosecute the perpetrators in a timely manner (Grobler et al., 2010a). The PCI DSS (Payment Card Industry Data Security Standard) also has an obligation for financial organisations to enable digital forensic processes to help provide for timely forensic investigation in the event of a compromise to any of its cloud service providers (Mouhtaropoulos et al., 2011). Looking to the near future, this trend is only expected to escalate, with more legal and regulatory mandates stipulating DFR obligations.

3.5.3.2 Policy

From being a good corporate governance demonstrator and being part of a legal and regulatory mandate, DFR has now become a policy matter. Hence, some organisations in the United Kingdom (UK) are already implementing digital forensic readiness policies (Irwin, 2012; Mouhtaropoulos et al., 2011; Rowlingson, 2004; Park, Kim et al., 2018). This came about after a legal mandate was announced for all organisations that are dealing directly or indirectly with the UK government to reasonably anticipate and respond in a forensically ready manner to any potential incident that might lead to a dispute or litigation. The aim of such a policy is to provide a systematic, standardised and legal basis for the acceptance and admissibility of potential digital evidence that may be required in a formal dispute or legal litigation process. Similar to the claims made in Tan (2002) (cited in Danielsson and Tjøstheim (2004)) and Rowlingson (2004) (cited in Grobler and Louwrens (2007)), most researchers claim that having such a policy in place will maximise organisations' potential to gather credible digital evidence that could be admissible in court, whilst minimising the cost of conducting a digital forensic investigation.

3.5.3.3 Costly Disruption to Business Operations

If implemented and executed correctly, a DFR capability should also help organisations to avoid business disruptions on the cloud service provider's side during an investigation. Hence, Cobb (2011) asserts that the objectives of a DFR policy are to maximise the usefulness of legally and ethically acquired admissible digital evidence and to minimise any costs of an investigation and disruptions to business operations. A well-structured DFR capability presents the potential to significantly reduce the cost and time of an investigation, while it could also increase the prospects of a quick and successful legal or dispute outcome with minimal disruptions to business activities. For example, in a cloud setting, when all the potential digital evidence is kept at a secure remote site, there is absolutely no need for the investigators to interrupt the cloud service provider's business operations. Investigations could be carried out with minimal business disruption and without raising unnecessary suspicion about the tenants being investigated.

3.5.3.4 Duty to Proactively Gather and Preserve Potential Digital Evidence

A DFR capability should enable cloud tenants and service providers to take proactive measures towards systematic acquisition; tamper-proof preservation; and secure storage of potential digital evidence. This is in anticipation of a potential litigation and/or disputes that may adversely affect and disrupt business operations. The aim is to proactively gather, preserve and store credible and admissible digital evidence to be ready in case of litigation or dispute and in response to "the duty to preserve", as obligated by some courts of law in the US (Cross and Kuwahara, 2010). Taylor (2012) argues that some US courts require litigants to capture, preserve and produce relevant and potential digital evidence well in advance. Failure to do so could result in heavy sanctions, tort liability and could even default judgement against the litigant. In order for such evidence to be accepted and admissible in court, the proactive processes and the toolsets and techniques that might be used to gather, preserve and store the evidence must be legal, ethical and forensically sound. They must also respect users' privacy and may not infringe their basic human rights. It is the duty of every cloud tenant and service provider to ensure that they take proactive measures to preserve potential evidence – more so if they reasonably anticipate an incident that could potentially lead to a litigation or dispute.

3.5.3.5 Inadmissibility of Digital Evidence

The foregoing assertion on the duty to preserve digital evidence raises the need to involve a legal expert to help determine the exact scope of digital evidence to be captured and the legality of the digital forensic tools to be used. It also needs to outline the necessary steps to gather evidence and ensure that it remains admissible in court, as per Daubert's much discussed criteria of determining the reliability and admissibility of scientific evidence (Computer Forensic and Computer Expert Witness Services, 1993; Majmudar, 1993; Mcleod, 2000; Orofino, 1996; Walsh, 1998; Welch, 2006). These criteria state that scientific and/or digital evidence (1) must be grounded on empirically testable theory or technique; (2) the theory must have been deeply scrutinised and peer-reviewed; (3) its potential error rate should be clearly stated; and (4) it must be based on generally accepted scientific principles (Majmudar, 1993). Otherwise, the collected evidence might be easily considered inappropriate or inadmissible in a court of law; or worse still, it might be considered a misconduct or unlawful invasion of privacy by the investigating authority and/or lawyer, as evident in the *United States vs Maynard* case (Harrington, 2012).

In summary, the compelling benefits presented by a DFR capability for cloud computing cannot be ignored. They make a good value proposition for organisations that are moving their data and applications to the cloud. Hence, it can be deduced that DFR is well positioned and suited to make a huge impact in cloud computing infrastructures. Therefore, in one of the next chapters, the author outlines some of the system requirements that are to be considered by organisations in their efforts to implement a DFR capability for the cloud environment.

3.6 CONCLUSION

The field of digital forensics is experiencing many challenges due to the rise of cloud computing. It would seem like cloud computing is making it easier for malicious users to perform their nefarious activities, and yet at the same time it makes it hard for digital forensic investigators to conduct their investigations. A traditional reactive digital forensic approach is surely not suitable for investigations in the cloud environment. Hence, this study and others argue that a DFR capability could assist digital forensic investigators to conduct effective investigations in the cloud in a timeous manner.

The next chapter reviews some of the most related and relevant work that has already been done to address some of the challenges raised.

CHAPTER 4 RELATED WORK

4.1 INTRODUCTION

The previous chapter provided background work to set the scene for the research reported on in this thesis. Chapter 4 provides a review of existing literature on how other researchers have attempted to solve the widespread data leakage threat in cloud computing. Researchers have made their attempts from several perspectives. However, the focus of this research is placed on three aspects, i.e. VM placement; authentication and digital forensic readiness. Below is a discussion on the related work.

The chapter is structured as follows: Section 4.2 discusses related work with regard to VM placement in the cloud. Section 4.3 discusses related work with regard to authentication in the cloud, while Section 4.4 discusses digital forensics in the cloud. This section is divided into three parts. Section 4.4.1 discusses related work with regard to a standardised way of conducting digital forensic investigations in the cloud and Section 4.4.2 discusses related work in terms of digital forensics readiness in general. Section 4.4.3 discusses related work in terms of digital forensics readiness in the cloud. Section 4.5 concludes this chapter and highlights the focus of Chapter 5.

4.2 VM PLACEMENT IN THE CLOUD

The VM placement problem in the context of cloud computing has been extensively studied by different researchers (Kesidis et al., 2018; Ferdaus et al., 2017; Filho et al., 2018; Quan, Wang and Ren, 2017; Levitin, Xing and Dai, 2018; Mashayekhy, Nejad and Grosu, 2014; Bartók and Mann, 2015; Alnajdi, Dogan and Al-Qahtani, 2016) who took different points of view. The latter are captured and summarised in a number of surveys (Filho et al., 2018; Challita, Paraiso and Merle, 2017; Madhusudhan and Satish, 2017; Thulo and Eloff, 2017; Alnajdi et al., 2016; Kaur and Bhardwaj, 2016; Masdari, Nabavi and Ahmadi, 2016; Pires and Baran, 2015; Usmani and Singh, 2016). The next subsection critically reviews each of the existing surveys in literature. This is followed by related work with respect to security-aware VM placement and by a subsection on conflict-aware VM placement. The last past of this subsection contains a brief discussion of VM placement in OpenNebula cloud infrastructure.

OpenNebula is a testbed for this study. The main goal of this chapter is to identify and highlight some of the research gaps in existing VM placement literature to help position this research.

4.2.1 Surveys on VM Placement

For example, Filho et al. (2018) provide a comprehensive review of the state of the art on VM placement in the cloud. They highlight open issues and challenges whilst reflecting on their relevancy in an increasing and demanding market. Filho et al. (2018) also review and classify different approaches that seek to address the VM placement problem in cloud computing in an effort to identify open issues and provide pointers for future solutions. Unfortunately, this work (Filho et al., 2018) does not identify the security issue of VM placement. It also does not mention the cost or risk associated with the physical isolations of VMs that belong to conflicting tenants. Similar to Filho et al. (2018), the work of Alnajdi et al. (2016) provides a critical analysis of existing dynamic VM placement algorithms to outline open research challenges and provide directions for future work. Alnajdi et al. (2016) also disregard the open issue of security associated with VM placements.

Challita et al. (2017) review VM placement literature that focuses on reducing power consumption, maximises resource utilisation and avoids traffic congestion. They also mention security as one area that remains unresolved with respect to VM placement (Challita et al., 2017). Kaur and Bhardwaj (2016) review VM placement algorithms that focus on resource consolidation. Kaur and Bhardwaj’s work is focused on research efforts that attempt to minimise the number of physical nodes to allocate VMs, VM allocation time and power consumption. Similar to the work of Kaur and Bhardwaj (2016) is the work by Usmani and Singh (2016), which also provides a comprehensive literature review of the state-of-the-art VM placement and resource consolidation techniques with the aim to minimise and improve energy consumption that they claim is increasing to unacceptable levels. The work of Madhusudhan and Satish (2017) focuses on reviewing and classifying VM placement algorithms that try to consolidate resources in order to maximise resource utilisation and minimise energy consumption. Three of these surveys (those of Kaur and Bhardwaj (2016), Usmani and Singh (2016), and Madhusudhan and Satish (2017)) fail to point out the impact of VM placement consolidation and its security implications. Only the work of Challita et al. (2017) does so.

Thulo and Eloff (2017) review existing literature on VM placement algorithms. Their findings show that there is gap in research efforts that consider the security aspect of VM placement (Thulo and Eloff, 2017). Therefore, and in order to close the gap, Thulo and Eloff go further to investigate existing optimised VM placement algorithms that at least have a potential to be further augmented with security features. Masdari et al. (2016) review and classify VM placement schemes based on their VM placement algorithm and evaluate their capabilities and objectives. However, and similar to the work of Thulo and Eloff (2017), Masdari et al. (2016) also mention that there is gap with regard to research work that addresses the security aspect of VM placement algorithms. Pires and Baran (2015) review and classify VM placement literature with respect to QoS, energy efficiency, service level agreements and resource consolidation. Their findings show that some researchers are to a lesser extent starting to focus on the security issues of VM placement algorithms (Pires and Baran, 2015). However, their work also points to the fact that this issue is not being addressed at scale, as expected.

In summary, the covered surveys indicate that the VM placement problem in cloud computing has been extensively studied from different viewpoints. However, the covered surveys all point to a lack of research efforts that focus on the security implications of VM placement algorithms. It is good to note that there are some isolated research efforts that have been identified as moving towards covering this research gap. However, these are still insufficient and wide apart. The next subsection discusses some of these research efforts.

4.2.2 Security-aware VM Placement

Besides the insufficient research efforts that investigate security-related implications of VM placement algorithms, there are some isolated and fragmented efforts – such as by Levitin et al. (2018), Thulo and Eloff (2017), Ahamed (2016), and Shetty, Yuchi and Song (2016) – which investigate data leakage threats as a result of security-related vulnerabilities in VM placement algorithms.

Levitin et al. (2018) model data security and survivability requirements to address an inter-VM attack that exploits a co-resident-based data vulnerability to leak or corrupt data held in a co-resident target's VM. They propose a data replication technique that partitions a tenant's data into multiple blocks and randomly distributes them to multiple servers to enhance security (Levitin et al., 2018). They also go further to create multiple replicas for each block to improve data survivability in a cloud that is subject to inter-VM attacks. This is a good approach to deal with data corruption. However, it comes at a high cost in terms of the underlying infrastructure that holds the multiple replica data blocks. The proposed random placement of data blocks to different servers does not guarantee non-co-residence of an attacker and a target VM. Therefore, the proposal by Levitin et al. (2018) cannot be argued to really address the problem of data leakage through a targeted inter-VM attack.

Thulo and Eloff (2017) propose a solution that uses optimised traffic and network-aware VM placement algorithms as a baseline and that incorporates security features with a goal to achieve an optimised security-aware VM placement algorithm. However, their study (Thulo and Eloff, 2017) ends before it can reflect how this augmentation is to be done and without mentioning the actual security features that are to be considered. Ahamed (2016) follows the same approach as Thulo and Eloff (2017) and proposes a solution that adds security features on optimal existing VM placement algorithms. Ahamed (2016) starts from an assumption that VMs have vulnerabilities and then profiles each VM according to its associated vulnerabilities as depicted in Common Vulnerabilities and Exposures (CVE) and Common Vulnerability Scoring System (CVSS) databases. The security profiles are then ranked and used in what they refer to as a security-aware and energy-efficient VM placement solution. This approach places what they consider vulnerable VMs together and also groups those that are not so vulnerable together.

Shetty et al. (2016) follow a similar approach to that of Ahamed (2016) and evaluate the vulnerability of VMs based on the National Vulnerability Database (NVD). The probability of risk for each VM is calculated based on its connections and dependencies with other VMs. The resultant evaluation is used to quantify the overall security risk of physical hosts, which is directly proportional to the total sum of vulnerabilities of all VMs hosted in each physical host. The proposed security-aware VM placement algorithms used in Shetty et al. (2016) ensure that VMs with high risks are placed in low survivability hosts and separate from those with a low risk (which are placed in high survivability hosts). They adopt this argument to somehow eliminate the possibility of placing highly vulnerable VMs or ‘bad neighbours’ on physical hosts that host low-risk VMs (Shetty et al., 2016).

In summary, only a few researchers have been investigating the security implications of VM placements in the cloud. Hence, there is a big need for more research efforts in this space. For example, there is a need to look at the placement of VMs belonging to conflicting tenants.

4.2.3 Conflict-aware VM Placement

Some researchers have already looked at the issue of handling conflicting tenants' requirements in VM placements. For example, the work of Mashayekhy et al. (2014); Si et al. (2014); Narwal, Kumar and Sharma (2016); Kwiat et al. (2015); as well as those of Han et al. (2015), (2013) and (2014) have already made progress on this topic.

Mashayekhy et al. (2014) consider data protection requirements in terms of restricting VM co-residence and co-location in 'trust restrictions' and 'disclosure restrictions'. They propose cryptographic mechanisms to solve the problem (Mashayekhy et al., 2014). However, reliance on cryptographic solutions alone is not the answer. Although cryptographic solutions can provide an added layer of security, they cannot stop a dedicated intruder from stealing and leaking encrypted datasets, which could then be brute-forced at a later stage by using powerful enough machines sourced from the cloud.

Some efforts are directed toward strengthening the physical isolation layer between tenants' VMs by employing the popular Chinese Wall Model (CWM) (Brewer and Nash, 1989). The CWM is a specialised access control policy that addresses issues of confidentiality in domains that are sensitive to conflict-of-interest (CoI); more so in commercial environments. The CWM defines impenetrable walls that segment data and enclose it in mutually disjoint CoI classes to avoid inadvertent access and leakage of confidential data to competitors.

Based on the CWM, Tsai et al. (2011) developed a Chinese Wall Central Management System (CWCMS) that handles the deployment of VMs. However, they did not consider the scalability problems associated with cloud infrastructure. Scalability in cloud computing is associated with optimal use of the underlying physical infrastructure. In the approach adopted by Tsai et al. (2011), the number of conflicts within each conflict class determine the population of physical nodes. Furthermore, one of the limitations of their work is that it considers a dichotomous approach to managing conflict-of-interest (CoI), i.e. tenants are either in conflict or they are not (Tsai et al., 2011).

Si et al. (2014) proposed a security awareness VM placement scheme (SVMPS). This scheme proposes two CoI relations, namely "aggressive conflict of interest relation" and "aggressive in ally with relation". Both relations are based on Brewer and Nash's CWM and are used to enhance isolation between conflicting tenants of the same cloud provider. Si et al. (2014) also propose a security-aware VM management scheme based on the CWM. However, their work only ensures that VMs of conflicting users are placed on different physical nodes. The approach in this thesis and the one in Tsai et al. (2011) partly address the problem; however, they are still insufficient to achieve a more desirable solution, which could ensure an even wider separation of conflicting tenants' VMs. Si et al. (2014) also fail to address the different degrees of CoI which would introduce some flexibility to the proposed model.

Similar to Tsai et al. (2011), Sailer et al. (2005) also do not discuss the issue of optimal utilisation of resources in the cloud infrastructure. Furthermore, the solution proposed by Sailer et al. (2005) does not provide tenants with any assurance that their VMs are not placed on the same physical infrastructure as that of their competitors. Wang et al. (2012) in turn seem to address the problems raised in Tsai et al. (2011) and Sailer et al. (2005), and they provide tenants with a mechanism to determine and verify co-residence of VMs instantiated on the same physical node. Wang et al. (2012) simultaneously access data from conflicting tenants' VMs and then measure the file access latency based on three factors – disk head contention; I/O request blocking; and disk cache and pre-fetch failure. From their experiments, they argue that it takes on average twice as long to read from the same physical node than from two separate ones (Wang et al., 2012). For example, if simultaneous access of data from two separate physical nodes takes two milli-seconds, it would take on average four milli-seconds to simultaneously access the same data from two VMs in the same physical node. The work of Wang et al. (2012) is a step in the right direction in enforcing and verification of the CWM in the cloud environment beyond bucket partition in Amazon S3 (simple cloud storage service). However, the author of this thesis argues that the CWM *as is*, is not well suited for the cloud environments.

Similar to Wang et al. (2012), Barthe et al. (2011) provide tenants with a formal mechanism to verify a CSP's physical VM isolation – an issue that they argue has not yet been fully investigated. The proposed mechanism therein formally establishes the hypervisor to enforce strong isolation properties. Their hypervisor ensures that no operating system can read or write memory that does not belong to it. Even though Barthe et al. (2011) address the issue of resolving conflicts on a virtualised platform, they do not consider it in the IaaS cloud platform and do not make use of the CWM. Wu et al. (2010) adopted and enforced the CWM to address the problem of insecure information flow. Similar to this thesis, Barthe et al. (2011), Sailer et al. (2005), Wang et al. (2010), and Wu et al. (2010) resolve conflict-of-interest (CoI) problems in cloud computing at the IaaS layer. However, they follow a dichotomy approach as do Tsai et al. (2011). This approach dictates that there is either a conflict or there is no conflict, as opposed to considering different degrees of conflict.

Amri, Hamdi and Brahmi (2017) also identify the issue of inter-VM interference in cloud environments. They argue that this issue emerges from the ambitions of server consolidation, with service providers aiming to improve energy efficiency and reap the cost-saving benefits of optimal resource utilisation. Although Amri et al. (2017) assert that the benefits cannot be realised until the inter-VM interference can be minimised, their work does not provide a concrete solution to this problem. They do however raise the necessary awareness of the inter-VM threat in cloud environments. Such awareness demonstrates the importance of the research reported on in the current thesis and shows that it is solving a real problem that both exists and is worth solving. However, and similar to the work of Wang et al. (2012), the work by Amri et al. (2017) is focused on inter-VM interference with regard to performance degradation as a result of resource contention. This is a bit different from the focus of this study, which is basically to address the data leakage threat posed by inter-VM attacks. However, the work by Amri et al. (2017) forms the basis of this thesis.

Narwal et al. (2016) assess the work of Han et al. (2013) which makes use of game theory principles to compare different VM placement policies in order to determine one that minimises an attacker's possibility of co-residence with other conflicting tenants. The main parameters in this game theory model are attackers and defenders. The attackers try to co-locate as many of their malicious VMs with as many target VMs as possible (i.e. to increase efficiency and coverage of their co-residence placement). The defenders use a set of VM placement policies (instead of one policy) to do the placement. This is such that when a VM placement request comes, the defender randomly selects a VM placement policy with a pre-defined probability from the set of all placement policies. This work also maintains a workload balance and minimum power consumption (Narwal et al., 2016).

Kwiat et al. (2015) use game theory principles to demonstrate the interdependency between users of the cloud. This interdependency is such that a vulnerability in a VM of one user affects other co-resident tenants' VMs and the controlling hypervisor. Kwiat et al. (2015) assume that every user is rational and makes decisions that maximises its payoff. The end goal is to minimise the negative effects of the interdependency of users' co-located VMs. Even though Kwiat et al. (2015) to a certain extent address the security implications of VM placement in the cloud, they do not address the prevalent issue of conflict of interest among co-located VMs. Moreover, the proposal in Kwiat et al. (2015) cannot handle inter-VM attacks from other co-resident tenants.

Han et al. (2014) propose a VM placement policy that allocates a new VM to a physical node that already has the highest number of VMs. Han et al. (2015) extend their earlier work (Han et al., 2013 and 2014) with a mathematical formulation of the solution to mitigate the threat of inter-VMs attacks on co-residents whilst satisfying constraints in workload balance and power consumption. In an attempt to prevent an attacker from starting too many VM instances in order to improve an attacker's efficiency and coverage (as discussed in Han et al. (2013)), Han et al. (2015) put all VMs of a user on the same physical host. This approach helps to control inter-VM attacks since at the host level there would be no co-residence of malicious and non-malicious VMs. However, for a normal tenant, having all your VMs instantiated on a single host creates a single point of failure. Han et al. (2016) extended their work on security game theory by introducing a solution that makes it hard for attackers to achieve co-residence with their target co-tenant. Their solution monitors the behaviour of attackers and legitimate tenants. Clustering techniques and semi-supervised learning are used to classify the tenants as either legitimate or attackers. This is a good approach. However, the success of the work done by Han et al. (2016) hinges on the claim that attackers might act differently from legitimate tenants. Should it so happen that the behaviour is similar for both attackers and legitimate tenants, the proposed solution would fail.

4.2.4 VM Placement in OpenNebula Cloud

Bagnasco, Vallero and Zaccolo (2018) propose a fair scheduling service (FaSS) for OpenNebula cloud infrastructure. The FaSS solution prioritises VM placement requests based on an initially assigned weight and historical resource usage. The FaSS is designed similar to the Haiza scheduling algorithm (Caballer et al., 2014) and they both support OpenNebula cloud infrastructure. Both of them interface with the existing scheduler without interfering with its underlying code and logic. However, neither of these algorithms addresses the issue of VM isolation to minimise in a cost-effective manner the risk of confidential data leakage posed by inter-VM attacks.

In summary, the overview of related work with respect to VM placement in the cloud has highlighted a number of research gaps in respect of existing approaches:

- The literature covered shows that there is a need for solutions that consider cost and risk implications of conflict-aware VM placement.
- Existing work takes a rigid approach (i.e. either you are in conflict or not) to managing CoI, without considering varying degrees of conflict.
- Existing work places more emphasis on the placement of a VMs from QoS, performance and resource utilisation viewpoints, which in most cases exclude the security aspect as rightly pointed out in Masdari et al. (2016).

In order to advance the current state of the art and contribute to the body of knowledge, this study attempts to close these research gaps. The next section discusses related work in terms of authentication for cloud computing environments.

4.3 AUTHENTICATION IN THE CLOUD

Cloud computing demands a new way of authenticating users. Compromised and weak credentials leave cloud services like Amazon Web Service, Microsoft Azure, Microsoft Office 365, Google Apps, Dropbox and others wide open to unauthorised access, which has a potential to escalate to a serious data leakage threat. According to Ablon (2018), Tout (2018) and Experian (2018) there is a rising data leakage threat from compromised and weak user credentials. Strong authentication presents a plausible solution and can play a key role as a first line of defence against cyber criminals on shared cloud services. Biometric scanning was initially offered as a better solution beyond the traditional username and password combination. However, the cost of biometric scanners is often very high. Numerous researchers are now looking at cost-effective ways to provide strong authentication mechanisms for cloud services. For example, smartphone-based sensors are currently being exploited to provide cost-effective voiceprints and fingerprints (Strom, 2015). The work of CA Technologies (2014) suggests that strong user authentication must not inconvenience users.

Raphiri et al. (2015) argue that strong authentication as a first line of defence can be used to secure the ‘front door’ to cloud computing services. Hence, they argue that strong authentication hinges on using freely available multiple factors such as MAC addresses and geo-location coordinates to authenticate users, based on their access devices and location access request point. Raphiri et al. (2015) assume that contextual data and user credentials are already secure throughout the entire authentication process.

Approaching strong authentication from a banking point of view, Dlamini et al. (2015) add the use of SIM card serial numbers; IMEI (International Mobile Equipment Identity) number; OTP provided through an SMS or email; and a concept called SurePhrase. This is over and above the MAC addresses and geo-location coordinates proposed by Raphiri et al. (2015). The SurePhrase concept is meant to provide an extra level of authentication on top of one-time password (OTP) tokens. The move towards strong multi-factor authentication (MFA) has seen a number of global corporates such as Google, Apple, Twitter, Facebook and LinkedIn using multiple factors (which include OTPs) to try and strengthen their authentication systems (Strom, 2015).

The research focus has generally moved beyond strong authentication mechanisms based on multiple factors and more towards a risk-based approach. A risk-based approach refers to an authentication system that makes access decisions based on the risk posed to the resources being requested by users. At the heart of a risk-based authentication system is a self-learning risk engine. The risk engine makes use of multiple factors to scale and adapt access decisions based on risk indicators. Kennedy et al. (2013) define a risk engine as an authentication system that continuously mines, monitors, analyses and processes user behavioural and context data. Strom (2015) defines it as risk-appropriate authentication that must consider numerous use cases and be able to evaluate minimum levels of accountability that are in line with the level of risk.

A risk-based authentication system authenticates users based on a combination of attributes (Goode, 2015; EMC Corporation, 2013; Saif, Siebenaler & Mapgaonkar, 2013). Goode (2015) and Saif et al. (2013) argue that a risk-based authentication solution can improve security without burdening users with extra levels of detail. The implication is that authentication must be done with ease of use and hide extra details to enhance the user experience. This will probably ease the security and usability trade-off for authentication systems. Such a solution is supported by CA Technologies (2014), which argues that authentication must scale up and down using a combination of factors to authenticate users in a cost-effective and convenient way. Unfortunately, Goode (2015), EMC Corporation (2013) and Saif et al. (2013) do not discuss how contextual data is to be kept secure and confidential throughout the process of authentication.

Some researchers refer to a risk-based authentication as adaptive authentication (RSA, 2015).

Some refer to it as context-aware authentication (Strom, 2015). RSA's adaptive authentication uses device forensics and user behavioural analysis to balance strong security in terms of authentication and usability. In support, Dlamini et al. (2015) argues that risk-based and strong authentication must be done in a seamless manner so as not to get between users and their core duties. Strom (2015) extends the above work to include user roles and activities. The user activity proposed in Strom (2015) is similar to the behavioural or context data proposed in RSA (2015). The work of Webroot (2014) further extends strong and risk-based authentication discussions to include speed and efficiency, and argues that authentication systems must be strong and scaled up or down, depending on the risk posed. However, strong authentication must be balanced with speed and efficiency, which raises ease and convenience of access. Hence, the work of Dlamini et al. (2016) ensures that authentication processes do not get in the way of users. Furthermore, they introduce a secure way to transmit user credentials from their login devices to the servers that hold user profiles (Dlamini et al., 2016).

Research efforts seem to be moving towards using keystroke dynamics. This approach authenticates users based on their typing patterns. The work of Ru and Eloff (1997) forms part of the foundational work in this area, together with that done by Kumar, Patwari and Sabale (2014); Ali et al. (2015); Haque, Khan and Khatoon (2015); Pisani, Lorena and Carvalho (2015); Gurary et al. (2016) and Dlamini et al. (2017). Similar to the work of Ru and Eloff, Kumar et al. (2014) group users into three classes, i.e. fast, moderate and slow, based on their keystroke dynamics. Authentication decisions take this into consideration when authenticating users. However, such a view results in high false positives and false negatives (Dlamini et al., 2017). This is because a user can fall in all three groups, depending on other external factors such as expertise, emotional state, session, or time of the day. For example, a fast user could be wrongly classified as a moderate user when fatigue creeps in.

Ali et al. (2015) consider external factors such as time, health conditions, emotional state and environment in their approach. Their results are much better. Pisani et al. (2015) extended the work of Ali et al. to propose a solution that uses multiple base classifiers and adaptive algorithms to overcome the external challenge introduced by the time factor. Gurary et al. (2016) make use of a soft touch keyboard on smartphones and add motion events in the place of key press events on a normal keyboard. The challenge with this approach is that it is not as unique as one would expect it to be. Hence, in the approach by Gurary et al. (2016), a login requests from two different users could result in a tie. This challenge might seem like a small glitch. However, it defeats the basic authentication principle, namely that users must be uniquely identified.

Dlamini et al. (2017) take a different approach and use what they refer to as local and global anomaly, and contextualisation. They extract keystroke features to create a user profile and train a neural network based on the typing behaviour of each user (Dlamini et al., 2017). The proposed solution detects how far a login instance is to the trained data, determines if it falls within the local or global anomaly threshold, and then transfers control to a risk engine that assigns a risk level score to the instance. An MFA engine takes over to prompt the user of the corresponding authentication token.

The current state of the art places greater emphasis on adding more authentication credentials and attributes to provide seamless and fast risk-based authentication. The covered work is in agreement on the use of multi-factors in conjunction with a risk-based approach. However, there are some research gaps. Firstly, apart from the work by Barreto et al. (2017) and Dlamini et al. (2017), the covered literature does not focus on how to secure user credentials or context data used for authentication. Secondly, the covered literature also does not discuss how to deal with locking out legitimate users. The study in hand tries to address the identified research gaps to make a significant contribution to the current body of knowledge. The next section discusses digital forensics in the cloud.

4.4 DIGITAL FORENSICS IN THE CLOUD

This section is divided into three parts. The first part discusses related work with regard to a standardised manner of conducting digital forensic investigations in the cloud. The second part discusses related work in terms of digital forensics readiness in general, and the third part discusses related work in terms of digital forensics readiness in the cloud.

4.4.1 Related Work – Standardised Digital Forensics in the Cloud

Conducting an effective digital forensics investigation to help prosecute cybercriminals in the cloud is still in its infancy and an open challenge (Endicott-Popovsky et al., 2007; Birk et al., 2013). Chapter 3 identified and elucidated on the challenges of conducting a digital forensic investigation in the cloud. It also became clear from Chapter 3 that there are still many grey areas with regard to conducting effective digital forensic investigations in the cloud and that research on conducting an effective digital forensic investigation in the cloud is still insufficient (Grispos et al., 2011; Gupta, 2011). This explains the lack or inadequacy of traditional tools and techniques for gathering digital forensic evidence in cloud computing environments.

The complexity of the cloud computing environment adds to the mix. For example, digital evidence cannot be found in one central location. Khan and Ullah (2017) argue that digital evidence exists on both the CSP's side and the client's side. Morioka and Sharbaf (2016) suggest that some digital evidence could be found in the communication channel between the client and CSP (e.g. an Internet Service Provider (ISP)). Going through each of these (CSP, client and communication channel) to collect digital evidence might involve complex and cumbersome legal processes. Moreover, in the cloud infrastructure, a digital artefact could be fragmented and stored in different locations. All of these create more problems for a digital forensic investigator dealing with a case that involves cloud computing.

Therefore, there is a need to improve existing tools and technologies to cater for the new cloud environment. Some researchers are calling for totally new tools and technologies (Sibiya, Venter and Fogwill, 2015) and many have already taken the initiative to move beyond the challenges identified by several researchers as summarised in Chapter 3, to propose solutions that are tailored for the cloud environment. However, the discussion in this section zooms in to focus only on existing solutions that consider a standardised approach for conducting a digital forensic investigation in the cloud. Standardised approaches are of particular relevance to this thesis because the field of digital forensics as a whole is working towards converging on widely accepted criteria of determining the admissibility of digital evidence and the process of conducting an effective investigation. Hence, research efforts that are not based on standards or widely accepted frameworks fall outside the scope of this section.

Martini and Choo (2012) assert that even though cloud computing might have introduced new challenges for digital forensics, the existing standards and key principles of digital forensics should be maintained. Hence, Martini and Choo's work is based on widely used NIST frameworks of McKemmish (1999) and Kent et al. (2006). Based on these frameworks, Martini and Choo (2012) proposed an integrated iterative digital forensic framework that is meant for cloud environments. Martini and Choo's framework is of particular relevance to this research because it requires digital forensic investigations in the cloud to follow an iterative approach and is based on widely accepted frameworks. However, the discussions need to move towards standardisation efforts.

Standardisation is key for ensuring that digital evidence collected from the cloud conforms to standard operating procedures, processes and principles. However, there is still insufficient standardisation research efforts (Grobler, 2010). Two positional papers identify the need for the research community to develop and adopt standardised approaches for digital forensic process models (Garfinkel, 2010; Grobler, 2010). Garfinkel's work draws from personal experience, literature review and round table discussions with digital forensic practitioners to provide a sneak preview of the future digital forensic research direction. Grobler's work reviews research on international digital forensic standards to determine if there is any progress; it thrashes out practical challenges and provides an overview of the future of digital forensic standardisation efforts. Garfinkel and Grobler concede that the fast-developing field of digital forensics presents many standardisation opportunities (Garfinkel, 2010; Grobler, 2010).

Sibiya et al. (2015) argue that standardisation can help improve the admissibility of digital evidence gathered from the cloud. Standards facilitate collaboration and the exchange of potential digital evidence between CSPs and multi-jurisdictional digital forensic investigators (Sibiya et al., 2015). Standardisation can help in ensuring that digital evidence conforms and complies with acceptable industry standards right across geographical and jurisdictional borders. In a court of law, digital evidence obtained by following acceptable industry standards carries more weight than that based on other non-standardised frameworks, best practices and guidelines. Hence, this thesis argues that there is a surging need for more research work to standardise the digital forensic process for cloud computing. Even though Grobler (2010) raised some valid concerns with regard to the development and adoption of digital forensic standards, it is encouraging to see other research efforts being targeted at the standardisation of digital forensic processes (Birk et al., 2013; Sibiya et al., 2015).

Birk et al. (2013) illustrate the applicability of the ISO/IEC 27037 digital forensic standard to the context of cloud computing. The ISO/IEC 27037 standard focuses on only the following digital forensic processes, i.e. identification, collection, acquisition, and preservation of potential digital evidence which can all be applied on a cloud environment. Birk et al. (2013) map and interpret the ISO/IEC 27037 standard in the context of cloud computing.

Sibiya et al. (2015) present and define a standardised digital forensic process model for conducting digital forensic investigation in cloud environments. This process model is implemented based on the ISO/IEC 27043 (Sibiya, Venter and Fogwill, 2012) and it practically demonstrates how to implement a digital forensic process model based on a standard for cloud environments and how to ensure that it is in compliance with the ISO/IEC 27043 standard processes.

Evidently, towards the end of 2017, the United Kingdom government mandated that digital forensic labs serving its criminal justice system must ensure that digital evidence is compliant with the ISO/IEC 17025 (Leyden, 2017). ISO/IEC 17025 specifies requirements for competence of tests and calibration of laboratories. The ISO/IEC 17025 does not necessarily link to the cloud environment. However, for future cloud-based digital forensic analysis a link may be found. Moving towards compliance is a step in the right direction. However, the move has been slapped with criticism from the academic community. For example, the call to mandate ISO/IEC 17025 on digital evidence from digital forensic labs in the UK has been labelled an ‘inappropriate, stupid and expensive’ exercise. The critics argue that mandating the ISO/IEC 17025 standard might compromise the quality of digital evidence available to the criminal justice system (Leyden, 2017). Despite the critics, this initiative will ensure that digital evidence that does not conform to the requirements stipulated in the ISO/IEC 17025 cannot be presented in UK courts. From this point of view, enforcing this standard will definitely reduce the quantity of digital evidence that gets admitted in the courts. Hopefully, it will also improve the quality of the admissible digital evidence.

It is encouraging to see the initial efforts towards the enforcement of digital forensic standards. We only hope that other countries can follow. It would be interesting to see other standards like the ISO/IEC 27043 and ISO/IEC 27037 being enforced worldwide. Mandating and enforcing digital forensic standards move the body of knowledge towards a common ground for digital forensic investigators. This is what some researchers (Garfinkel, 2010; Sibiya et al., 2015) argued could help move the digital forensic community towards ensuring that digital evidence collected from the cloud conforms to industry standards.

The fore-going section sets the scene and provides the current state of the art on some of the existing work that has already been carried out in terms of standardised digital forensics in the cloud. It has been mentioned that research efforts focusing on conducting an effective digital forensic investigation in the cloud is still insufficient (Park, Kim et al., 2018). This is even worse when one zooms in to work that is based on standards. It gets even worse when considering enforcement of existing digital forensic standards. Hence, the author of this thesis argues that there is a glaring gap in research on how to conduct a digital forensic investigation that is based on and informed by existing digital forensic standards. This thesis will therefore attempt to address this research gap.

The next section discusses digital forensic readiness as one way to address the identified research gaps (as discussed in this section) and to prepare the cloud environment for an effective and efficient digital forensic investigation.

4.4.2 Related Work - Digital Forensic Readiness

The turn of the 21st century marked the beginning of a new era in the digital forensics field. Back in 2001 researchers at the Digital Forensics Research Workshop focused primarily on the processes of several tools and techniques for conducting a digital forensic investigation and on the maturity of these processes (Endicott-Popovsky and Frincke, 2007). It would seem that digital forensic practitioners would have done everything possible to develop reliable tools and techniques. However, online criminal activities continued to grow and only a handful of reported cases resulted in successful prosecution. Two successful cases reported in Endicott-Popovsky and Frincke (2007) reflect the ineffectiveness or inefficiency of existing digital forensic investigation tools and techniques. The ineffectiveness or inefficiency comes in terms of the prolonged time and high costs of conducting reactive digital forensic investigations in comparison to the minor consequences suffered by the offenders. Alharbi et al. (2011) argued that taking a reactive approach to hastily collect digital evidence was time consuming and costly. They also argued that a reactive approach could have legal repercussions, which could result in malicious offenders getting away with minor sentences (Alharbi et al., 2011). For example, a denial of service attack which incurred damages worth $400 000, took 417 hours of investigation time and amounted to $27 800 of investigation costs. In the end, the offender was given a mere community service sentence, mainly because at the time of this incident there were no laws making the malicious act of ‘denial of a service’ a criminal offence (Endicott-Popovsky and Frincke, 2007; Endicott-Popovsky et al., 2007). This demonstrates the ineffectiveness or inefficiency of the digital forensic techniques or legal systems of the time. Thus, a call was raised to combat inefficiency and ineffectiveness of the “then current” tools, methods, techniques and frameworks of conducting a digital forensic investigation.

A need was created for a proactive approach that would ensure effective and efficient digital forensic investigations. The works of Tan (2001) and Rowlingson (2004) perhaps lay the most important foundation for the concept of digital forensic readiness (Reddy and Venter, 2012). Since their early work, several researchers (Carrier and Spafford, 2003; Danielsson and Tjøstheim, 2004; Endicotte-Popovsky and Frincke, 2007; Endicott-Popovsky et al., 2007; Alharbi et al., 2011; Pooe and Labuschagne, 2012) have realised the importance of taking a proactive approach towards digital forensic investigations. Some researchers have tackled it from a business perspective in general (Danielsson and Tjøstheim, 2004; Rowlingson, 2004; Pooe and Labuschagne, 2012), while others focused specifically on a particular technical aspect (e.g. network forensic readiness) of it (Endicott-Popovsky et al., 2007; Endicott-Popovsky and Frincke, 2007; Mouton and Venter, 2011; Ammann, 2012). Still others focused on further aspects like sound investigation (Carrier and Spafford, 2003). The divergence of these pieces of work demonstrates that even today there is no single successful methodology or approach towards a standard digital forensic readiness capability (Taylor et al., 2007).

Alharbi et al. (2011) argue that this is mainly because work aimed at adopting a proactive approach towards effective and efficient digital forensic investigations is still insufficient and imprecise. For example, even though several researchers (Danielsson and Tjøstheim, 2004; Rowlingson, 2004; Endicott-Popovsky et al., 2007; Taylor et al., 2007; Barske, Stander and Jordaan, 2010; Duranti and Endicott-Popovsky, 2010; Mouton and Venter 2011) cite the seminal work of Tan (2001), which defined proactive digital forensics in terms of two objectives ("maximizing the ability of an environment to gather credible digital evidence" and "minimizing the cost of an incident response"), there is still no standard definition or agreed-upon process of proactive digital forensics (Taylor et al., 2007). This demonstrates fragmented and ad-hoc research efforts as alluded to in the work of Taylor, Endicott-Popovsky and Frincke (2007) and Endicott-Popovsky et al. (2007), who argue that there is still no comprehensive and enterprise-wide digital forensic readiness capability. Hence, proactive digital forensic investigations remain an open issue that requires even more work.

Endicott-Popovsky and Frincke (2007) and Endicott-Popovsky et al. (2007) proposed a methodology for embedding digital forensic readiness in information systems based on the NIST Information Systems Development Life Cycle, which is informed by their 4R strategy (Resistance, Recognition, Recovery and Redress). The 4R strategy is aimed at making systems resistant to attacks, detect and recognise attacks, and quickly recover from the attack. An additional component is to establish who is responsible for the attacks, and to gather concrete and admissible digital evidence to hold them accountable in a court of law (Endicott-Popovsky et al., 2007). Similar to the proposal in this thesis, the work of Endicott-Popovsky and Frincke (2007) and Endicott-Popovsky et al. (2007) suggest that evidence acquisition and collection should consider legal requirements for compliance and preservation standards that would ensure admissibility in courtrooms. Furthermore, the work of Endicott-Popovsky and Frincke (2007) and Endicott-Popovsky et al. (2007) resonate with the proposed 3-tier (planning, implementation and assessment) model adopted in this thesis. Their work also put more emphasis on chain of custody and proper preservation of digital evidence in their disposition phase. However, the thesis in hand considers these within the implementation tier. Furthermore, instead of taking a waterfall approach as is the case in Endicott-Popovsky and Frincke (2007) and Endicott-Popovsky et al. (2007), this thesis implements an iterative approach similar to that in the work of Martini and Choo (2012).

Danielsson and Tjøstheim (2004) adopt an approach similar to the ten-step processes by Rowlingson (2004) to also outline a structured digital forensic readiness capability, based on five fundamental guidelines. These guidelines are of key significance in the iterative 3-tier proposal in this thesis. For example, Danielsson and Tjøstheim (2004) argue that digital forensic readiness should analyse the organisational need for digital evidence. This resonates with the first process within the planning tier of the proposed 3-tier cloud model in this thesis. Danielsson and Tjøstheim (2004) seemed to consider the global nature of digital forensics, much like this thesis does. Similar to Endicott-Popovsky and Frincke (2007) and Endicott-Popovsky et al. (2007), the work of Danielsson and Tjøstheim (2004) considers a risk assessment and goes further to say this must be combined with a cost-benefit analysis. The assessment tier of the proposed iterative 3-tier model also makes use of the risk assessment and a cost-benefit analysis approach. The main aim is to ensure that the cost of acquiring digital evidence should be evaluated and weighted against the risk of doing so. This is in line with the US' proportionality doctrine discussed in Chapter 3.

More importantly, Danielsson and Tjøstheim (2004) raise one most overlooked issue that they claim might be a major stumbling block to the deployment of digital forensic readiness capabilities, namely privacy concerns. They argue that privacy concerns could be mitigated by incorporating privacy-enhancing technologies into forensic readiness tools and procedures. Danielsson and Tjøstheim (2004) move away from incorporating privacy-enhancing technologies, but argue that by considering applicable privacy legislation, the privacy concerns could be addressed. They conclude by asserting that considering their guidelines would enable organisations to proactively configure their information systems to collect and preserve potential digital evidence according to applicable international or national legislation and be cognisant of users' privacy concerns.

Grispos et al. (2017) provide a good snapshot analysis of digital forensic readiness solutions. They classify digital forensic readiness solutions across five dimensions (policy; process, system and frameworks; people; forensic-by-design) and concluded that previous work focused on the first for dimensions and overlooked forensic-by-design. Forensic-by-design is part of a DFR strategy that includes incident response, forensic capabilities and best practices that are integrated into the design and development of systems (Pandya, Homayoun and Dehghantanha, 2018).

Forensic-by-design is not equivalent to DFR, but it can be used to enhance DFR. Mink et al. (2016) allude to the fact that forensic-by-design should be incorporated into today's DFR systems. Ab Rahman et al. (2016) and Ab Rahman, Cahyani and Choo (2017) propose a forensic-by-design framework that integrates tools and best practices in the design and development of cloud systems. Grispos et al. (2017) conclude that forensic-by-design should start by considering digital forensic requirements in all software development processes. Therefore, this thesis considers and incorporates digital forensic readiness requirements from the onset. Chapter 5 looks at the system requirements, which include those that are specific to DFR in the context of cloud computing. The next section discusses related work with regard to DFR in the cloud.

4.4.3 Related Work – Digital Forensic Readiness in the Cloud

Thus far, the above work in section 4.4.2 has discussed digital forensic readiness in general to set the scene for its application in the cloud environment. There is a need for work that discusses digital forensic readiness in the context of cloud computing. For example, Kebande and Venter (2016) identified a lack of DFR's standard operating procedures in the cloud. The question on "how to prepare cloud environments to become ready for an effective and efficient digital forensic investigation", still needs an answer (Chung et al., 2012; Dlamini et al., 2014). Most of the work covered in Chapter 3 focuses specifically on the challenges that cloud computing poses for digital forensic investigators. There is a need to move beyond the challenges towards providing tangible solutions. Several researchers have already taken the initiative to thrash out possible solutions.

Zawoad, Dutta and Hasan (2013) and Baykara, Das and Tuna (2017) propose a secure solution to store and provide logs as digital evidence for digital forensic purposes. The proposed solution in Zawoad et al. (2013) is specific to cloud service providers. It attempts to proactively store logs, whilst preserving their confidentiality, the accuracy and integrity of proofs of past logs, and the privacy of cloud tenants. The aim is to increase the auditability of the cloud environment, especially for regulatory compliance issues (Zawoad et al., 2013). However, the focus of Zawoad et al. (2013) is placed on logs collected from the service provider. This is not the case for Baykara et al. (2017), who collect logs from general traffic flow – not necessarily from cloud service providers.

Patrascu and Patriciu (2013) present a digital forensic-enabled cloud computing architecture that monitors, gathers and reproduces user activity from guest virtual machines. The focus is on the hypervisor as the kernel of virtualisation. Unlike the work of Zawoad et al. (2013), Patrascu and Patriciu (2013) do not touch on regulatory compliance and they also overlook the security aspects of the collected logs. This issue is key to this thesis.

However, depending on cloud service providers, digital evidence has proven to be problematic, more especially after the revelations that law enforcement agencies are collecting data from cloud service providers without the consent of the data owners (Teing, Dehghantanha and Choo, 2017). Therefore, Teing et al. (2017) argue that dependency on compromised cloud service providers might cause problems for digital investigators. They then concede that client edge devices may be the only viable source of digital evidence. Other researchers extend Martini and Choo’s framework to cater for potential digital evidence remnants extracted from client edge devices after using BitTorrent Sync applications and CloudMe (Teing et al., 2017). Mendoza et al. (2015) take advantage of persistent data storage of web browsers to extract, collect and aggregate potential digital evidence at the client side and to support a digital forensic investigation. Quick and Choo identify some of the potential digital evidence that remains on clients’ edge devices after accessing Dropbox, Microsoft SkyDrive and Google Drive – all of which an investigator could use for forensic analysis (Quick and Choo, 2013a, 2013b, 2013c, 2014).

Khan and Ullah (2017) argue that collecting potential digital artefacts from either a client or cloud service provider yields an incomplete picture that might result in evidence with gaps. They propose a log aggregation digital forensic analysis framework for cloud computing environments that gathers logs of virtual sessions. The logs are gathered from both the client edge node accessing a service and from the servers of the cloud service provider. These logs are then integrated, pre-processed and stored in what they refer to as the evidentiary log repository for investigators. The focus is placed on indexing, normalisation, integration, correlation and sequencing of the logs for ease of fast retrieval when required. However, Khan and Ullah (2017) overlook the issues around securing the logs in transit (as they are transmitted from the point of collection to the remote repository) and at rest (in the repository).

Kumar Raju, Gosala and Geethakumari (2017) tackle the issue of event reconstruction in cloud forensics. Their work modifies the Leader-Follower algorithm to achieve what they claim is proper log aggregation. They argue that their approach reduces the number of events in the target digital evidence found in the cloud and makes it possible to do proper and effective event reconstruction of cloud systems. Event reconstruction is a process that aims to achieve a proper chain of custody to illustrate the process of how digital evidence reached its current state (Kumar Raju et al., 2018). Kumar Raju et al. (2017) use a framework that they call Cloud Service-based Event Reconstruction (CLOSER) which focuses on cloud service logs. The main challenge with this approach involves the aggregation and reduction of the events as this may produce evidence with gaps, which may result in a loss of credibility or accuracy. However, Kumar Raju et al. (2017) claim that their solution achieves high reduction without ‘much’ loss of data, which implies no loss of digital evidence credibility or accuracy. However, they do not quantify what ‘much’ signifies.

Morioka and Sharbaf (2016) argue that browser persistent potential digital evidence as mentioned in Mendoza et al. (2015) can be easily deleted after browsing. They argue that anyone with basic computer literacy can erase all traces of communication between client’s and service provider’s devices and in the process erase all potential digital evidence (Morioka and Sharbaf, 2016). They also assert that investigators can make use of digital evidence that could be found in ISPs – the communication channel between the client and CSP (Morioka and Sharbaf, 2016) – to create a complete picture of the crime scene. Although this is a good approach, the challenge is getting hold of the client devices before clients can tamper with the evidence.

Dykstra and Sherman (2013) propose a solution called forensic Openstack tools, which overcome the challenge of remote digital evidence integrity by storing logs in hash trees and returning digital forensic evidence with cryptographic hashes. Their proposed solution allows investigators to conduct their forensically sound and trustworthy investigations without any interaction with the cloud service provider. Their work is similar to the proposal in this thesis in that there is a remote repository for digital forensic evidence from where investigators are to conduct their investigations, without necessarily tampering with the live cloud systems (Dykstra and Sherman, 2013). The design by Dykstra and Sherman (2013) reduces the cost of conducting an investigation and it does not disrupt business operations of the cloud service provider. Furthermore, it provides the integrity of the digital evidence. However, none of the covered related work has looked at DFR in the cloud from a standardisation point of view.

Kebande and Venter (2016) adopted the digital forensic readiness processes described in the ISO/IEC 27043: 2015 standard and proposed a cloud forensic readiness model. Their proposed model uses a botnet to proactively harvest potential digital evidence and preserve it in preparation for a digital forensic investigation as outlined in the ISO/IEC 27043: 2015. Similar to Dykstra and Sherman's work, Kebande and Venter hash digital evidence for integrity before it is accessed by investigators. It is interesting to note that the forensic-by-design framework proposed in Ab Rahman et al. (2017) is also based on and adopts the ISO/IEC 27043: 2015.

The work of Marshall and Paige (2018), though not necessarily related to digital forensic readiness, has identified incorrectness of digital evidence in digital forensic tools and methods as a challenge. This is because of the incomplete validation and verification of digital evidence. Marshall and Paige (2018) support the use of the ISO 17025 and ISO/IEC 27041 accreditation standards to improve the validity of digital evidence. Hence, the author of this thesis argues that the move towards standardisation efforts is positive.

The foregoing two subsections have set the scene regarding the current state of the art with regard to DFR in the cloud. The covered related work has highlighted and made the initial move towards incorporating forensic-by-design in cloud solutions. However, more work is still required on how this could be done in a more practical manner. The related work has also highlighted the problem of over dependency on cloud service providers for potential digital evidence. Despite the fact that researchers are now focusing on other avenues to extract digital evidence, a research gap still exists on how to reduce this dependency. There remains a need to explore remote and centralised log repositories. Furthermore, there is a need for research that explores ways to secure potential digital evidence in transit and at rest in remote centralised repositories. Finally, and most importantly, there is a need for research that explores ways of conducting a digital forensic readiness investigation in the cloud, based on existing standards. This thesis attempts to address these research gaps. The next section concludes this chapter.

4.5 CONCLUSION

Cloud computing has come with an ever-rising and prevalent data leakage threat. Several researchers have made attempts to address this threat from different viewpoints. However, this thesis has focused on discussing related work regarding three relevant aspects, i.e. VM placement; authentication; and digital forensics and readiness. The critical analysis of the related work has helped us to arrive at the following conclusions:

- From a VM placement perspective – the covered literature seems to be focusing on other aspects such as QoS, resource utilisation, performance, energy efficiency and costs, but not on security. Only a few research efforts have been made to investigate security-aware VM placements. The few that do cover the security aspects of VM placement are insufficient and they do not consider the cost and risk implications of security-aware VM placements. Therefore, this thesis pose the research question; *how can we improve and provide VM placement that prevents the co-location of conflicting VMs on cloud computing*. As a contribution to the body of knowledge, the author aims to provide logical and physical separation of conflicting VMs based on the risk and cost implications of conflict-aware VM placement.
- From an authentication perspective – the related work in this area seems to be leaning towards the use of multi-factor and risk-based authentication approaches. Recent research adds contextual and behavioural data to the multi-factor and risk-based authentication mechanisms. However, there are some research gaps. For example, the covered literature in this area does not discuss how to deal with unauthorised users armed with stolen or compromised legitimate user credentials. In addressing the gaps identified in literature with respect to authentication in the cloud, this thesis pose the research question; *how can we improve and provide appropriate authentication that would be suitable for cloud computing?* In answering this research question, the author makes a bold assumption that attackers already have stolen credentials in their

possession and as a contribution to the body of knowledge proposes a risk-based multifactor authentication solution to stop them.

- From a digital forensic perspective – the covered literature indicates a move towards standardisation in the field of digital forensics. Even though this move is not widely adopted, it is a good first step in the right direction. The covered literature in this area also reflects a growing interest in the use of digital forensic readiness for the agile cloud computing environments. However, there is still not enough research in these areas. Hence, we can deduce that conducting an effective digital forensics investigation to help prosecute cybercriminals in the cloud is still quite immature and an open challenge. Considering that I did not even touch the issue of attribution, makes it exciting to work in this field. In an attempt to address the shortcoming in the digital forensic literature, the author posed the research question; *how can we prepare the cloud to become ready for e-discovery of digital evidence?* The author posits that digital forensic readiness can be used to prepare the cloud environment for an investigation and to make plausible attempts to prevent attacks from happening in the first place.

The next chapter outlines the system requirements of solution that this thesis proposes. The requirements listed are gleaned from the research gaps identified in this chapter.

CHAPTER 5 A USE CASE DRIVEN APPROACH FOR REQUIREMENTS ELICITATION

5.1 INTRODUCTION

The previous chapter provided a review of related work, with a specific focus on VM placement; authentication; digital forensics and readiness in cloud computing. The purpose was to understand the current state of the art in all four aspects. A clear understanding of the current state of the art aided the process of identifying research gaps that have not yet been addressed. The author's understanding of the current state of the art also helps to position this thesis in the current body of knowledge. This is to show that the solution that the author proposes in this thesis does not exist in isolation but is based on and extends existing related work.

The identified research gaps help to scope and focus the research efforts on the areas that stand to make a significant contribution. Furthermore, the gaps help to clearly show how this thesis advances the current state of the art and moves beyond what other researchers have already done. Therefore, Chapter 5 takes the research gaps identified in Chapter 4 as a baseline input for the system requirements elicitation process.

The system requirements elicitation process investigates and identifies functional requirements that may not have been captured by the research gaps. These other requirements are generated from a use case driven approach. Use cases present an abstract, readable and understandable view of a system from an end user's perspective. Moreover, use cases provide a systemic and intuitive way of keeping all requirements focused on providing value for the end user. Hence, use cases are considered one of the most effective and widely used techniques for eliciting and documenting system requirements (Odeh et al. 2004; Wang, Anokhi and Anderl, 2017). Therefore, Chapter 5 adopts a use case approach to elicit, formulate and document functional requirements. The final output of this chapter is a list of all-encompassing system requirements. The system requirements are a necessary pre-requisite input for the design and development of the solution.

The rest of Chapter 5 is structured as follows: Section 5.2 elicits, formulates and discusses requirements as derived from the use cases and existing research gaps. Section 5.3 presents an overall list of requirements for the design of our solution. These requirements are geared toward controlling the prevalent threat of data leakage in cloud computing infrastructures. Section 5.4 concludes this chapter and provides a high-level overview of what is to be discussed in the next chapter.

5.2 SYSTEM REQUIREMENTS ELICITATION

This section focuses on eliciting and discussing the system requirements. It is divided into three subsections, namely system requirements for VM placement, authentication, and digital forensics readiness in cloud computing. Each subsection starts off with a use case for requirements elicitation and ends with a brief description of the requirements concerned.

5.2.1 System Requirements for VM Placement in the Cloud

This section discusses the system requirements for VM placement. It starts with a use case and

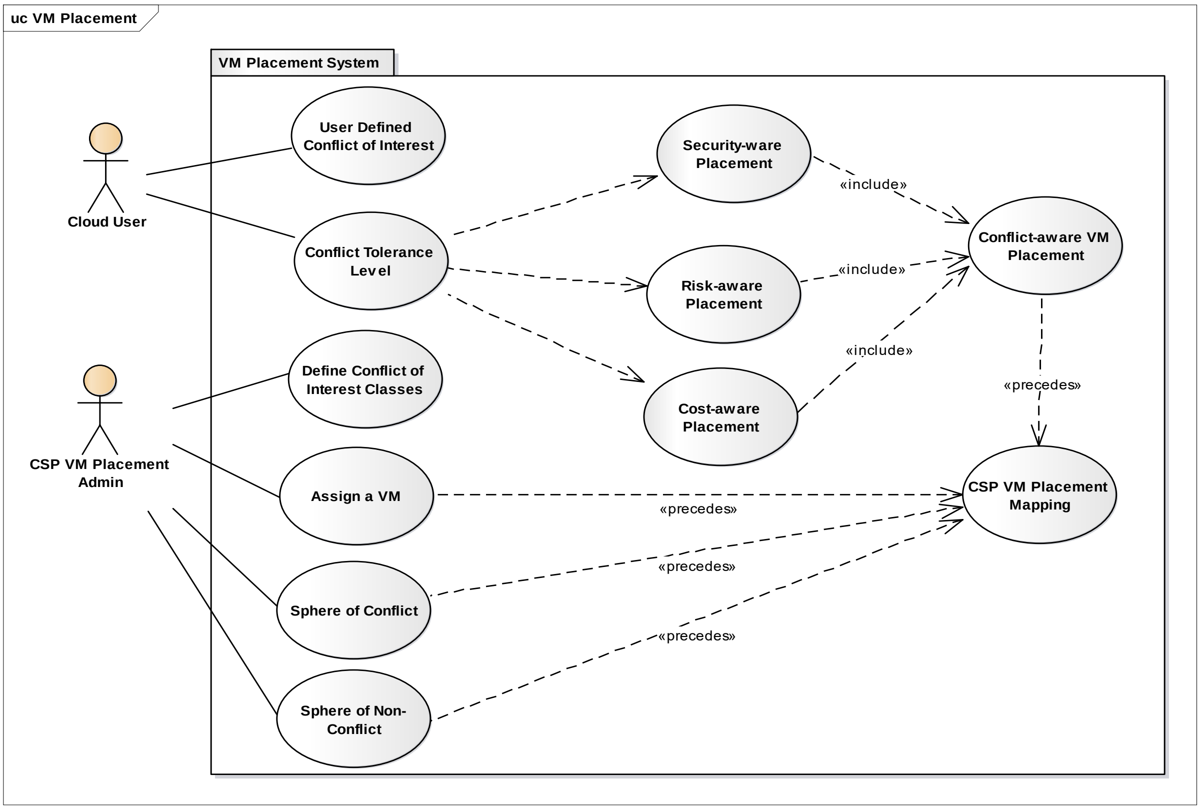

Figure 5.1:- Use Case Diagram for VM Placement in the Cloud

ends with the system requirements of the VM placement component. Figure 5.1 depicts a use case diagram for VM placement in the cloud. The use case has two actors – a cloud user and a cloud service provider's administrator who manages the VM placement. A user defines their own conflict of interest and conflict tolerance level (CTL).

A user's conflict tolerance level is based on the three independent variables: security, risk and cost. These three are combined to achieve a conflict-aware VM placement. The admin defines the general conflict of interest and then assigns a VM to host a user's data. The CSP also determines the sphere and non-sphere of conflict. A sphere of conflict consists of those users that are in conflict with a particular user. A non-sphere of conflict consists of those users that are not in conflict. A CSP's VM mapping gives the administrator an overall view of all placements on the CSP's infrastructure. However, this is only visible to the admin and not to the users. From the above use case, the author identified the system requirements as shown in Figure 5.2.

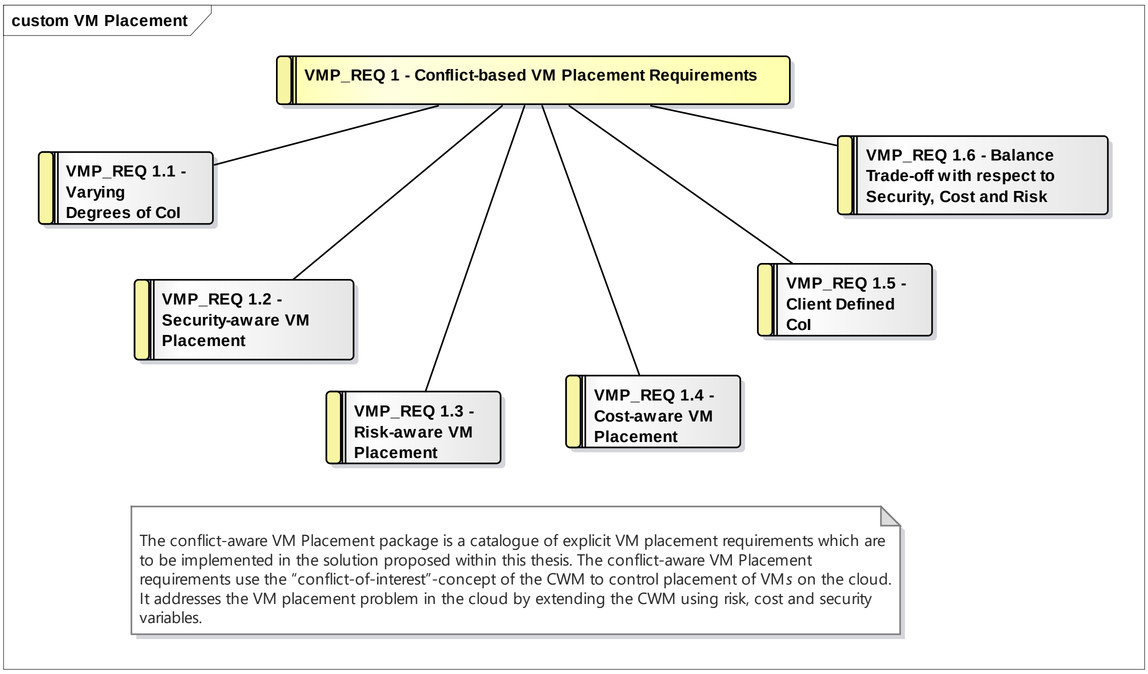
custom VM Placement
VMP_REQ 1 - Conflict-based VM Placement Requirements
VMP_REQ 1.1 - Varying Degrees of CoI
VMP_REQ 1.2 - Security-aware VM Placement
VMP_REQ 1.3 - Risk-aware VM Placement
VMP_REQ 1.4 - Cost-aware VM Placement
VMP_REQ 1.5 - Client Defined CoI
VMP_REQ 1.6 - Balance Trade-off with respect to Security, Cost and Risk

The conflict-aware VM Placement package is a catalogue of explicit VM placement requirements which are to be implemented in the solution proposed within this thesis. The conflict-aware VM Placement requirements use the "conflict-of-interest"-concept of the CWM to control placement of VMs on the cloud. It addresses the VM placement problem in the cloud by extending the CWM using risk, cost and security variables.

Figure 5.2:- Requirements for Conflict-based VM Placement in Cloud Computing

VMP_REQ 1.1:- Must provide for varying degrees of CoI

This requirement abstracts and models conflict-aware VM placements based on different degrees of CoI classes. It moves away from the binary approach (i.e. in conflict or non-conflict) that is prevalent in existing research work. This thesis considers different classes of CoI which are discussed later on.

VMP_REQ 1.2:- Must provide security-aware VM placement

This requirement ensures that users consider the security implications of their VM placement. Based on this requirement, and as a contribution, this study aims to improve on the currently insufficient security-aware placement research for cloud computing. However, security in this case is to be achieved by physically separating VMs that are in conflict VM based on their CTL.

VMP_REQ 1.3:- Must provide risk-aware VM placement

This requirement ensures that users are aware of the potential risk of confidential data leakage from inter-VM attacks that emanate inside the virtualisation layer of a cloud infrastructure when malicious co-resident VMs could siphon confidential data from other non-suspecting users that share the same infrastructure. This is also directly linked to a specific user's CTL.

VMP_REQ 1.4:- Must provide for cost-aware VM placement

This requirement indicates a figurative cost (without any loss of generality – not in monetary terms such as US $ or UK ₤ or Euro €) of co-hosting users' VMs. There is a high cost associated with keeping conflicting VMs as far apart as possible, for instance in different locations; more especially for users that are in direct conflict. A lower cost is required to co-host a user's VM with that of a conflicting tenant in the same physical node.

VMP_REQ 1.5:- Must allow clients to define own sphere of CoI

This requirement enables a user to define potential conflict of interest with other tenants based on the user's own "view of the world". In other words, a user might know of other conflicting tenants over and above those identified by a CSP. This requirement caters for such special cases.

VMP_REQ 1.6:- Must balance the trade-off for conflict-aware VM placements with respect to security, cost and risk

> This requirement ensures that the final placement of VMs strikes the right balance on the three independent variables (i.e. security, risk and cost) derived from a user's initial CTL and CoI class.

5.2.2 System Requirements for Authentication in the Cloud

Figure 5.3 depicts a use case diagram for authentication in the cloud. There are four actors: a cloud user, access device, a digital forensic investigator and a law enforcement agent. Each of these actors must be authenticated to access the system. The system was initially designed to authenticate users and their access devices. However, after the first round of use case development, it became essential to also authenticate investigators and law enforcement agents accessing our solution with their access devices.

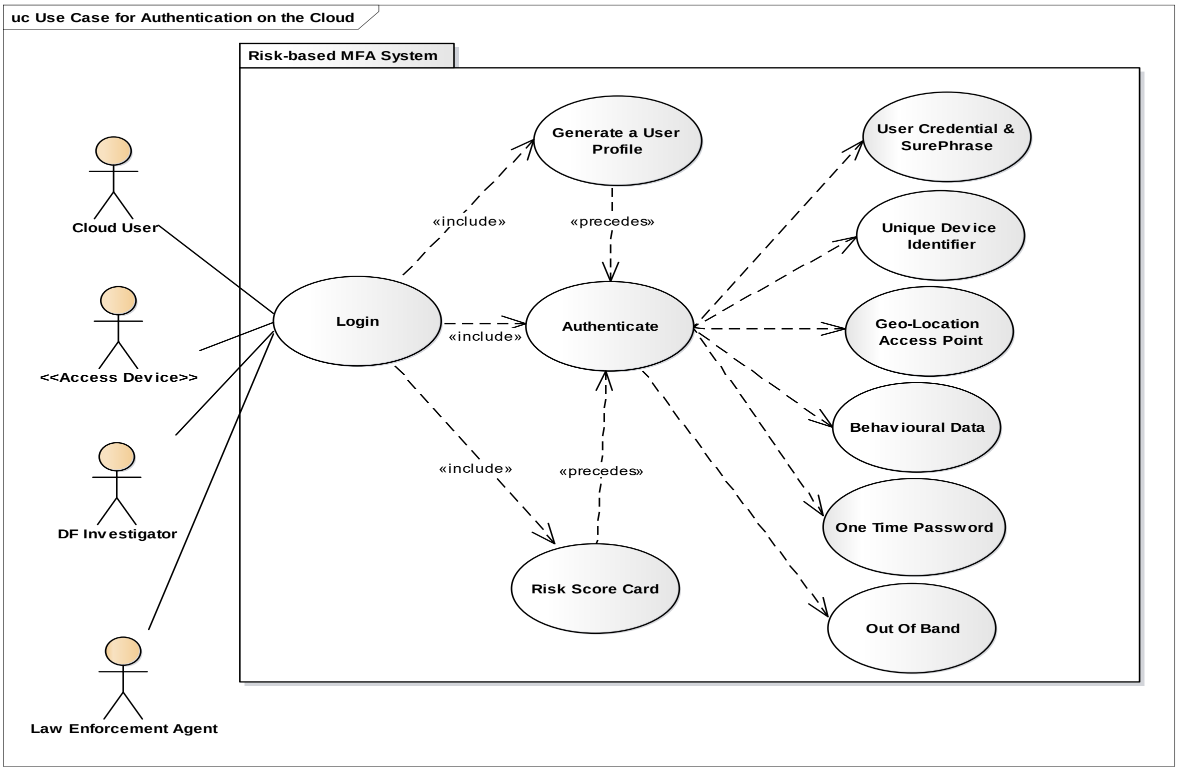

Figure 5.3:- Use Case Diagram for Authentication in the Cloud

The proposed solution dictates that all users accessing it must be authenticated based on their risk levels. Each login instance is captured and stored for reference at a later stage. The captured login instances help to create a user profile. Each new login instance is compared to the historical data in the user profile. Should there be deviations, the user would be prompted for extra login information. At the very extreme, a user will be sent an out-of-band token that requires the user to authenticate with another registered device. This use case reflects the multiple factors that a user might be prompted to use for access. Instead of granting or denying access based on correct user credentials, the solution requires other factors to properly authenticate users based on their risk level. Figure 5.4 depicts the requirements for the approach towards authentication in the cloud. These are derived from the use case in Figure 5.3 and the research gaps identified in Chapter 4.

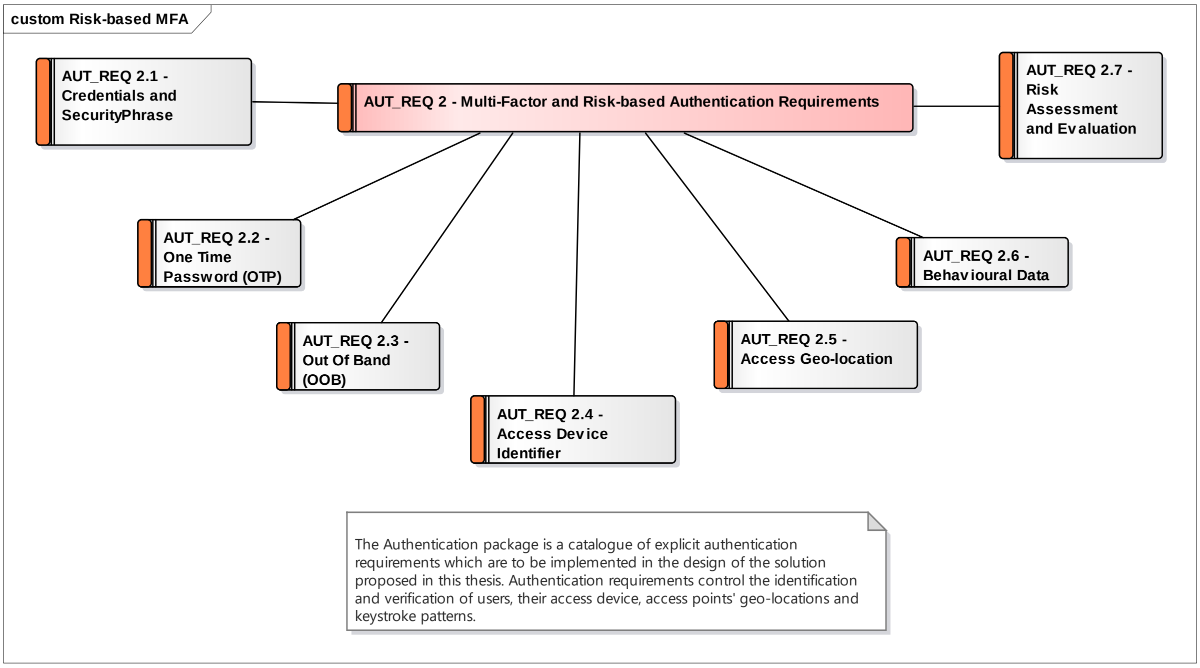
custom Risk-based MFA
AUT_REQ 2.1 - Credentials and SecurityPhrase
AUT_REQ 2 - Multi-Factor and Risk-based Authentication Requirements
AUT_REQ 2.7 - Risk Assessment and Evaluation
AUT_REQ 2.2 - One Time Password (OTP)
AUT_REQ 2.6 - Behavioural Data
AUT_REQ 2.3 - Out Of Band (OOB)
AUT_REQ 2.5 - Access Geo-location
AUT_REQ 2.4 - Access Device Identifier

The Authentication package is a catalogue of explicit authentication requirements which are to be implemented in the design of the solution proposed in this thesis. Authentication requirements control the identification and verification of users, their access device, access points' geo-locations and keystroke patterns.

Figure 5.4:- Requirements for Risk-based Multi-Factor Authentication in Cloud Computing

Below is a brief discussion of each of the requirements.

AUT_REQ 2.1:- Must provide access based on user credentials and SecurityPhrase

This requirement extends existing research work beyond user credentials and adds the concept of SecurityPhrase, which is a combination of characters that are normally required from a strong password and that are generated from a phrase. This factor is basically added to offset often-weak passwords and provide an extra layer of protection.

AUT_REQ 2.2:- Must provide OTP authentication channels

This requirement ensures that users are prompted for a randomly generated token that gets sent either to their email or cell-phone. The user will have to retrieve this and use it to authenticate. This requirement caters for a user whose risk profile has been evaluated to high.

AUT_REQ 2.3:- Must provide OOB authentication channels

This requirement ensures that a very high-risk user is prompted to authenticate using a different device. Once authenticated, the session is transferred back to the original device. A user with such a risk profile is only given two chances. Should they fail with the two tries, the requirement mandates that the user be locked out of the system and unlocking their profile requires an administrator.

AUT_REQ 2.4:- Must provide access based on device identifier

This requirement ensures that users are authenticated along with their access devices. This part captures the unique device identifier and compares it with one that is in a user profile. Similar to AUT_REQ 2.3, this is done in the background without the user's knowledge. If the user credentials and SecurityPhrase pass the authentication, but the device fails; the user will be prompted for an OTP. Successful entry of the OTP would trigger the system to register the device under the user's profile for future login attempts.

AUT_REQ 2.5:- Must provide access based on geo-location

This requirement ensures that user profiles are enriched with geo-location data from access login points. The system captures geo-location information of each login in the background and makes use of it for subsequent access requests. These are captured and used with an acceptable marginal error, because GPS systems are not accurate. However, it is used to strengthen authentication.

AUT_REQ 2.6:- Must provide access based on behavioural data

This requirement captures unique keystroke dynamics of a user and adds them to a user profile for login. Keystroke dynamics have a great potential in uniquely identifying users. However, it requires that the solution be trained repeatedly until it can correctly classify users with minimal false positives. This requirement moves the work in hand beyond the current state of the art. We hope this will eventually replace all other factors so that the system can be able to identify users solely based on their keystroke patterns.

AUT_REQ 2.7:- Must be able to assess and evaluate risk and choose the appropriate authentication combination

This requirement facilitates seamless and on-the-fly risk assessment and evaluation of login attempts. It then helps to choose an appropriate factor combination for authentication that matches the risk profile. This requirement ensures that the authentication process does not interfere with a user's duties. Interference is one problem that most users have been complaining about when it comes to security tools that are meant to enable a safe working environment, but eventually they become a stumbling block.

5.2.3 System Requirements for Digital Forensic Readiness in the Cloud

Due to the complexity of the use case for digital forensic readiness in the cloud, the author has decided to split it into three (see Figures 5.5, 5.6 and 5.7). Figure 5.5 is a use case that focuses on the scoping exercise. This use case requires a CSP and legal expert to determine the scope of digital evidence to be collected, so as to ensure that digital evidence is in compliance with industry standards and regulatory frameworks. It also ensures that digital evidence achieves business compliance goals and is admissible in court (Dlamini et al., 2014).

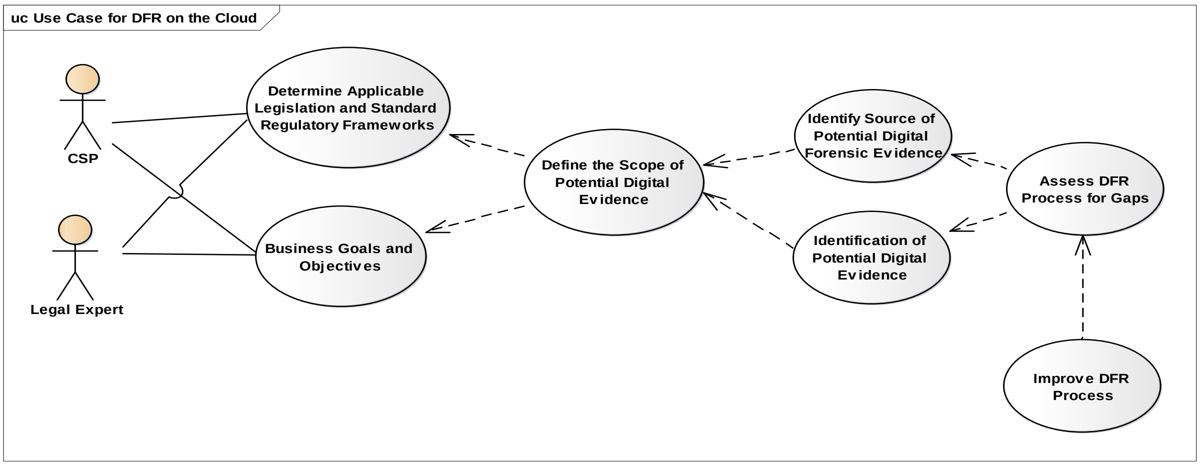
uc Use Case for DFR on the Cloud
CSP
Legal Expert
Determine Applicable Legislation and Standard Regulatory Frameworks
Business Goals and Objectives
Define the Scope of Potential Digital Evidence
Identify Source of Potential Digital Forensic Evidence
Identification of Potential Digital Evidence
Assess DFR Process for Gaps
Improve DFR Process

Figure 5.5:- Use Case Diagram for DFR in the Cloud

Figure 5.6 depicts a use case of a DF Readiness module that does the actual acquisitioning of digital evidence from the cloud to a remote repository. The module starts by verifying sources of potential digital evidence. It also does incident detection and sends an alert to an investigator for immediate action to contain detected incidents. The module then does the acquisition from the verified sources, collects and classifies the evidence before it tags and hashes it for easy retrieval and integrity preservations. After this, the evidence is encrypted before the remote repository is authenticated for verification purposes. Once the repository has been authenticated, a secure channel is opened for transmission. The digital evidence is then transmitted over a secure channel for storage in the remote repository.

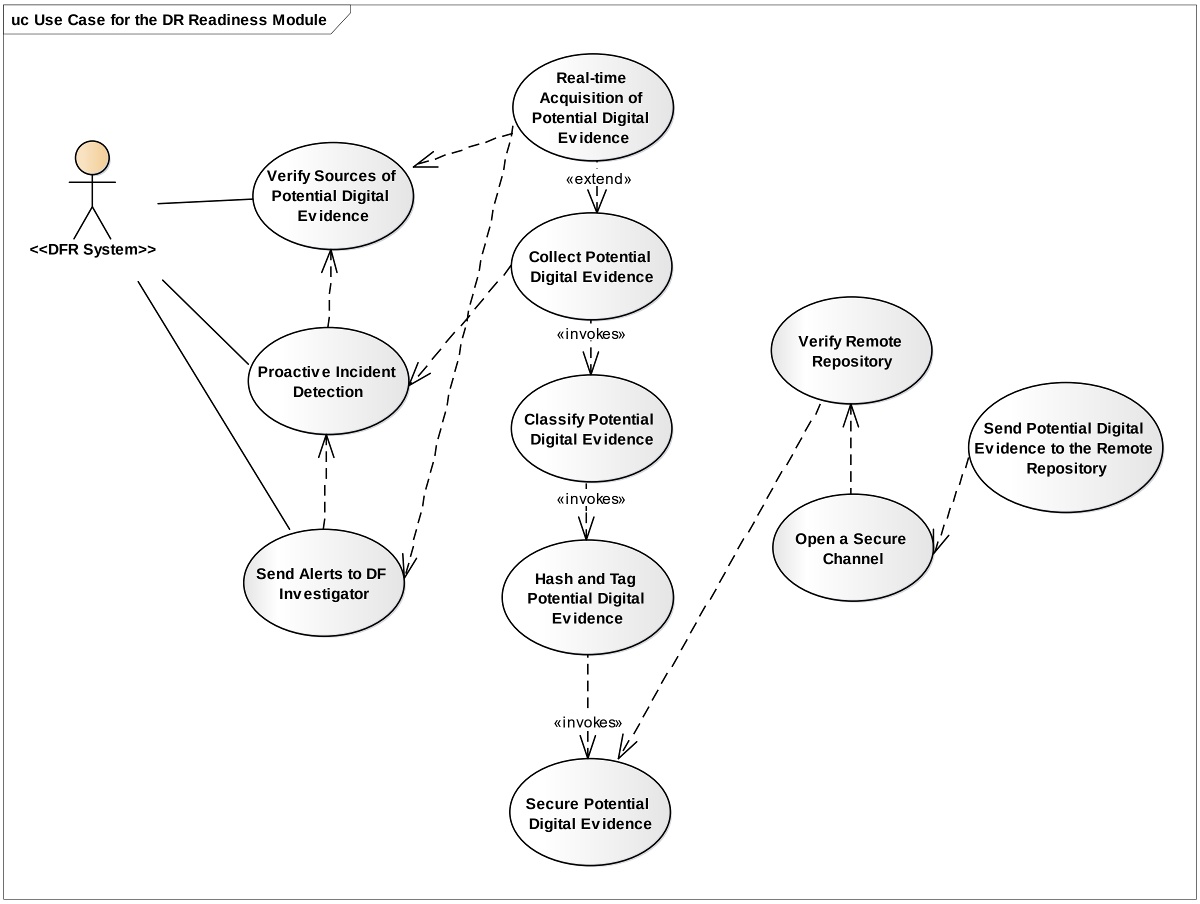
uc Use Case for the DR Readiness Module
<<DFR System>>
Real-time Acquisition of Potential Digital Evidence
Verify Sources of Potential Digital Evidence
«extend»
Collect Potential Digital Evidence
«invokes»
Proactive Incident Detection
Verify Remote Repository
Classify Potential Digital Evidence
Send Potential Digital Evidence to the Remote Repository
«invokes»
Open a Secure Channel
Send Alerts to DF Investigator
Hash and Tag Potential Digital Evidence
«invokes»
Secure Potential Digital Evidence

Figure 5.6:- Use Case Diagram for DR Readiness Module

Figure 5.7 depicts a use case for digital forensic investigators, law enforcement agents and compliance regulators. This use case depicts the goals of each actor when accessing the remote repository where the digital evidence is stored. An investigator would receive incident alerts from the DFR module, take precautions to prevent or stop an incident, and then escalate the incident for investigation. Since digital evidence has already been collected, there is no need to disrupt the services of the CSP. After proper authentication, as outlined in Figure 5.3, an investigator and law enforcer would just go into the process of creating a snapshot of the potential digital evidence and move straight to the analysis and report findings processes. On the other hand, a compliance regulatory officer would go straight to doing a gap analysis in order to determine compliance and due diligence, before they can make recommendations, issue a certificate of compliance or due diligence, and then proceed to report the findings.

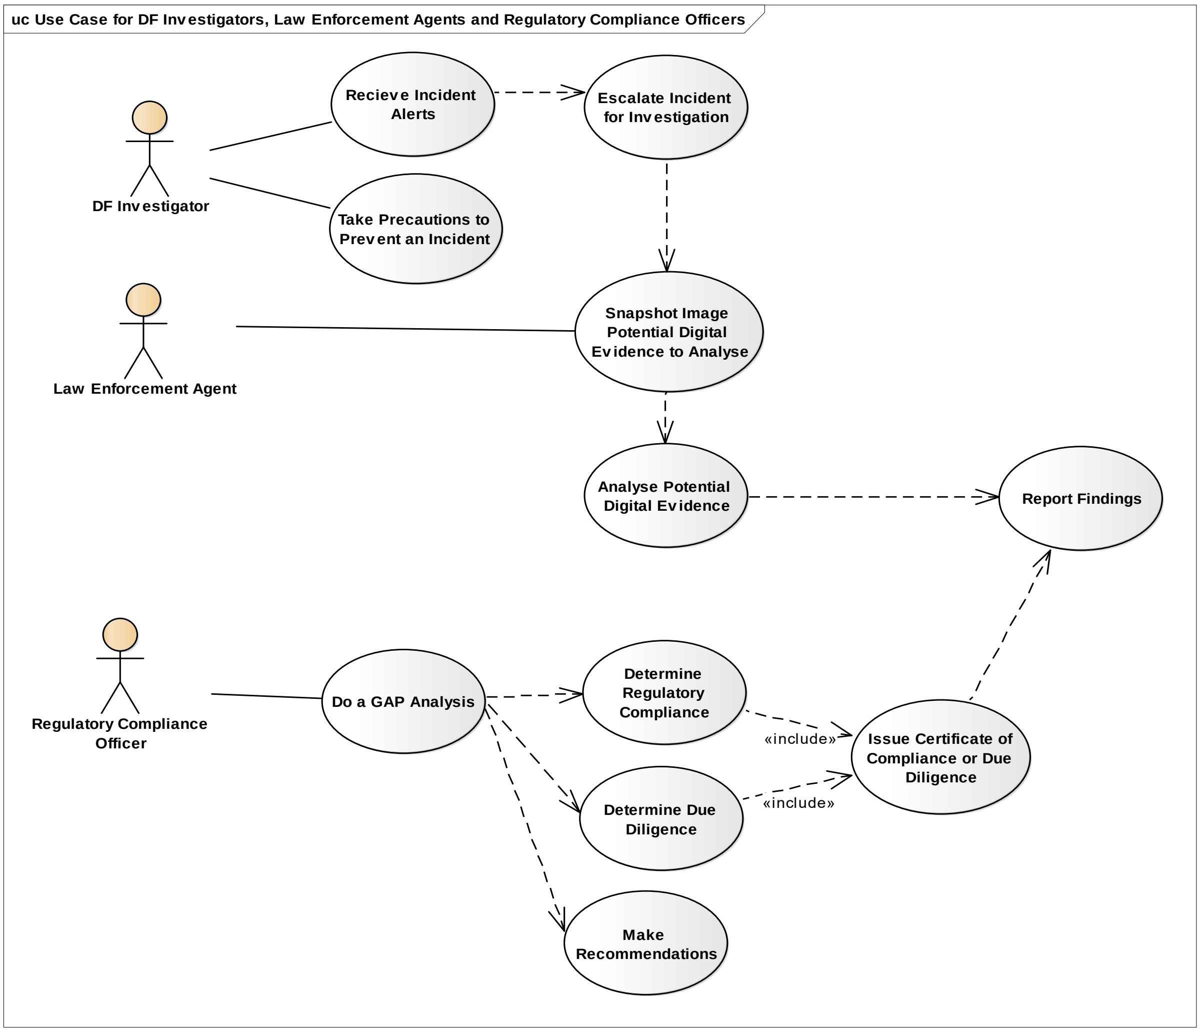
uc Use Case for DF Investigators, Law Enforcement Agents and Regulatory Compliance Officers
DF Investigator
Recieve Incident Alerts
Escalate Incident for Investigation
Take Precautions to Prevent an Incident
Law Enforcement Agent
Snapshot Image Potential Digital Evidence to Analyse
Analyse Potential Digital Evidence
Report Findings
Regulatory Compliance Officer
Do a GAP Analysis
Determine Regulatory Compliance
«include»
Issue Certificate of Compliance or Due Diligence
Determine Due Diligence
«include»
Make Recommendations

Figure 5.7:- Use Case Diagram for DF Investigators, Law Enforcement Agents (LEAs) and Regulatory Compliance Officers

Below are the requirements that arise from the above use cases (i.e. Figures 5.5, 5.6 and 5.7.) and the research gaps identified in Chapter 4. Figure 5.8 captures the following requirements derived from the research gaps (Dlamini et al., 2014):

- Securely and selectively gather, store and preserve legally admissible digital evidence
- Demonstrate due diligence and compliance with legal and regulatory mandates
- Minimise business disruptions
- Minimise the time or cost of acquiring digital evidence in the cloud

Figure 5.8 depicts a holistic picture that includes the requirements in Dlamini et al. (2014) and those that emanate from the above use cases. Below, follows a brief discussion of each of these requirements.

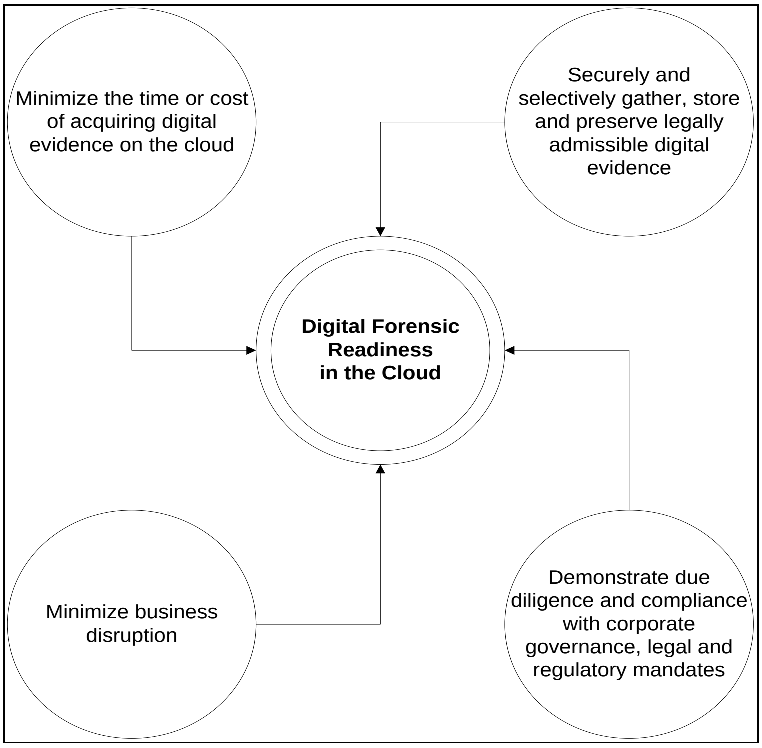
Minimize the time or cost of acquiring digital evidence on the cloud
Securely and selectively gather, store and preserve legally admissible digital evidence
Digital Forensic Readiness in the Cloud
Minimize business disruption
Demonstrate due diligence and compliance with corporate governance, legal and regulatory mandates

Figure 5.8:- Requirements for DFR in the Cloud (Dlamini et al., 2014)

DFR_REQ 3.1:- Must be based on industry standards

This requirement ensures that acquired potential digital evidence is based on and conforms to industry standards. This is an attempt to maximise the usefulness, acceptability and admissibility of potential digital evidence, especially in courts. It somehow emphasises the push towards enforcing standardisation efforts in the field of digital forensics.

DFR_REQ 3.2:- Must minimise business disruptions

This requirement takes cognisance of the fact that multiple users share the same cloud infrastructure. It ensures that an investigation can be carried out without any major disruptions to business activities of the service provider or other co-resident users of the shared cloud infrastructure. Business disruptions might also raise alarms that would prompt suspects under investigation to try and delete crucial evidence before it could be captured by investigators.

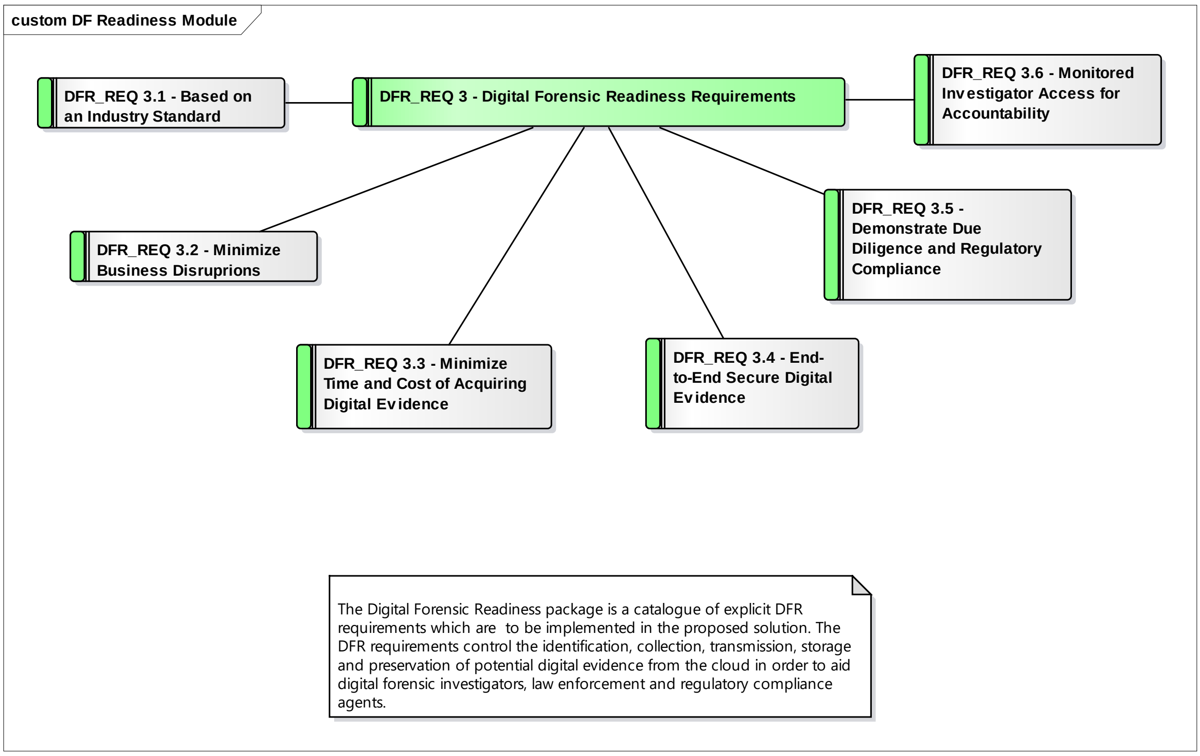
custom DF Readiness Module
DFR_REQ 3.1 - Based on an Industry Standard
DFR_REQ 3 - Digital Forensic Readiness Requirements
DFR_REQ 3.6 - Monitored Investigator Access for Accountability
DFR_REQ 3.5 - Demonstrate Due Diligence and Regulatory Compliance
DFR_REQ 3.2 - Minimize Business Disruprions
DFR_REQ 3.3 - Minimize Time and Cost of Acquiring Digital Evidence
DFR_REQ 3.4 - End-to-End Secure Digital Evidence
The Digital Forensic Readiness package is a catalogue of explicit DFR requirements which are to be implemented in the proposed solution. The DFR requirements control the identification, collection, transmission, storage and preservation of potential digital evidence from the cloud in order to aid digital forensic investigators, law enforcement and regulatory compliance agents.

Figure 5.9:- Requirements for Digital Forensic Readiness in Cloud Computing

DFR_REQ 3.3:- Must minimise time and cost of acquiring potential digital evidence

This requirement necessitates that potential digital evidence be acquired in the shortest time and with the smallest budget possible (Dlamini et al., 2014). This is mainly because significant time lags during evidence acquisition could easily risk potential digital evidence being corrupted or lost before it could be captured for analysis. The proposed solution must be able to significantly minimise the time and/or cost it takes digital forensic investigators to legally acquire potential digital evidence and conduct an investigation in the cloud.

DFR_REQ 3.4:- Must provide end-to-end secure digital evidence processing

This requirement ensures that potential digital evidence is secure from the point of acquisition until it is presented in courts or before a disciplinary committee. There should be no gaps that might expose the evidence to alterations, deletion or leakage, because this might bring it to disrepute when presented. The security of the potential digital evidence must be tight end-to-end, i.e. from acquisition right through to presentation.

DFR_REQ 3.5:- Must demonstrate due diligence and regulatory compliance

This requirement ensures that the solution must strive to demonstrate due diligence and compliance with corporate governance, legal and regulatory requirements. It was recommended based on the high penalties for non-compliance with legal and regulatory mandates that go beyond simple monetary fines to include criminal and/or civil litigation. For example, the Payment Card Industry Data Security Standard (PCI-DSS) states that a level one merchant that fails to comply with its requirements could be fined an amount in the range of $500,000 - $1,000,000 (Dlamini et al., 2014). In an extreme case, failure to comply could potentially lead to a merchant losing its right to accept credit card charges (Centrify Corporation, 2012).

DFR_REQ 3.6:- Must provide monitored access for accountability

This requirement emphasises that access to the acquired potential evidence that is stored in a secure repository must be tightly controlled and monitored for accountability. This requirement would also ensure that unauthorised and authorised users cannot tamper with potential digital evidence.

In summary, the proposed list of system requirements to be considered in the design of the conceptual architecture is depicted in Figure 5.10. The next section concludes this chapter and points to the work to be covered in Chapter 6.

5.3 CONCLUSION

Chapter 5 took the research gaps identified in Chapter 4 as a baseline input for the system requirement elicitation process. Furthermore, this chapter adopted a use case driven approach to elicit and document other requirements on top of those derived from the research gaps. Using these two approaches, Chapter 5's contribution is a list of all-encompassing system requirements. The system requirements are a necessary pre-requisite input for the design and development of a high-level conceptual model.

The next chapter focuses on the design and development of a high-level conceptual solution, based on the list of all-encompassing system requirements. This solution demonstrates how each of the identified requirements come together to form the main modules and expected relationships among them.

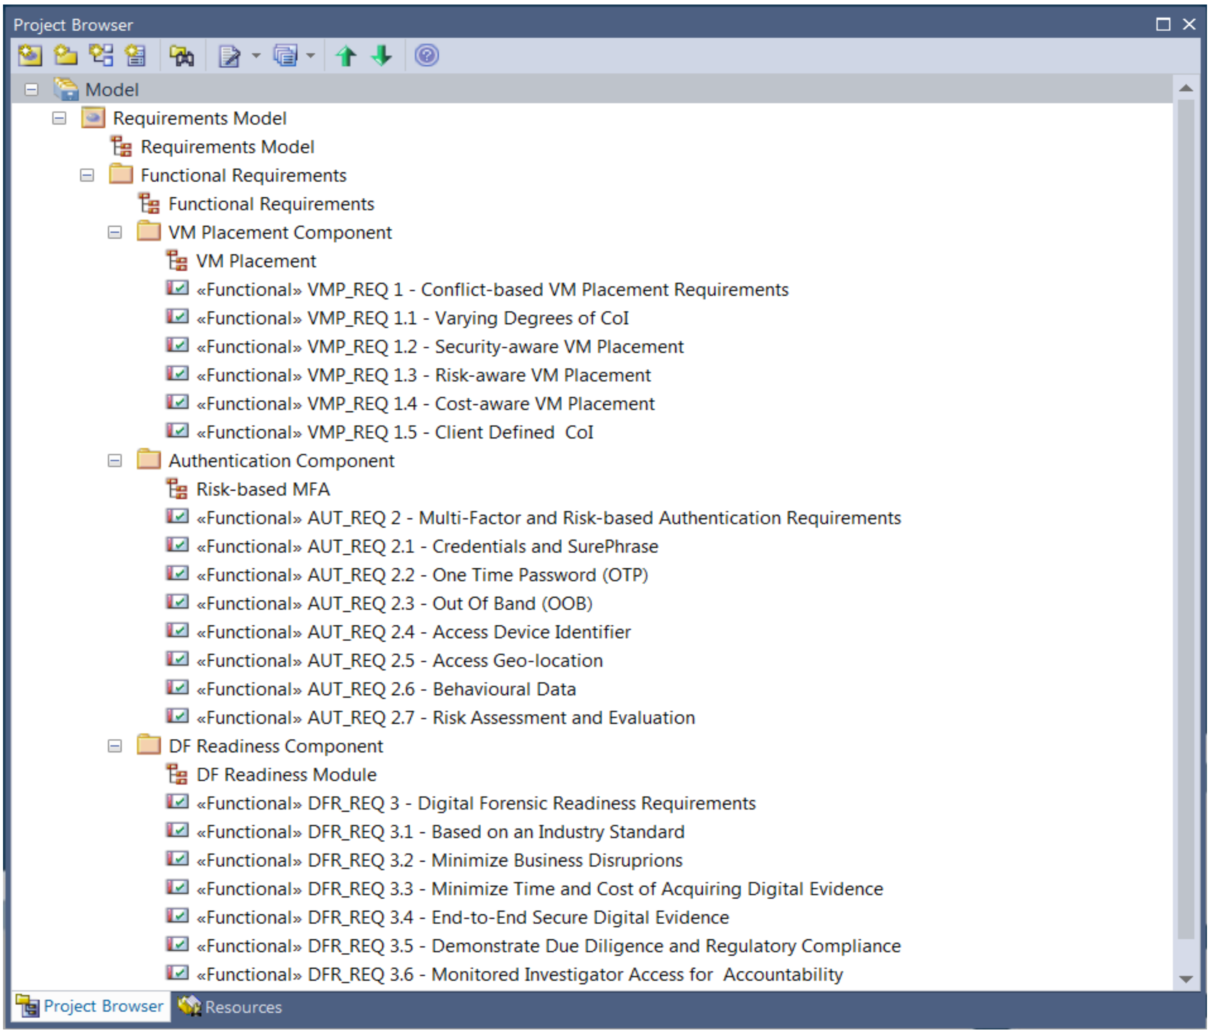

- Model
  - Requirements Model
    - Requirements Model
    - Functional Requirements
      - Functional Requirements
      - VM Placement Component
        - VM Placement
        - «Functional» VMP_REQ 1 - Conflict-based VM Placement Requirements
        - «Functional» VMP_REQ 1.1 - Varying Degrees of CoI
        - «Functional» VMP_REQ 1.2 - Security-aware VM Placement
        - «Functional» VMP_REQ 1.3 - Risk-aware VM Placement
        - «Functional» VMP_REQ 1.4 - Cost-aware VM Placement
        - «Functional» VMP_REQ 1.5 - Client Defined CoI
      - Authentication Component
        - Risk-based MFA
        - «Functional» AUT_REQ 2 - Multi-Factor and Risk-based Authentication Requirements
        - «Functional» AUT_REQ 2.1 - Credentials and SurePhrase
        - «Functional» AUT_REQ 2.2 - One Time Password (OTP)
        - «Functional» AUT_REQ 2.3 - Out Of Band (OOB)
        - «Functional» AUT_REQ 2.4 - Access Device Identifier
        - «Functional» AUT_REQ 2.5 - Access Geo-location
        - «Functional» AUT_REQ 2.6 - Behavioural Data
        - «Functional» AUT_REQ 2.7 - Risk Assessment and Evaluation
      - DF Readiness Component
        - DF Readiness Module
        - «Functional» DFR_REQ 3 - Digital Forensic Readiness Requirements
        - «Functional» DFR_REQ 3.1 - Based on an Industry Standard
        - «Functional» DFR_REQ 3.2 - Minimize Business Disruprions
        - «Functional» DFR_REQ 3.3 - Minimize Time and Cost of Acquiring Digital Evidence
        - «Functional» DFR_REQ 3.4 - End-to-End Secure Digital Evidence
        - «Functional» DFR_REQ 3.5 - Demonstrate Due Diligence and Regulatory Compliance
        - «Functional» DFR_REQ 3.6 - Monitored Investigator Access for Accountability

Figure 5.10:- List of All-encompassing Functional Requirements

CHAPTER 6 A HIGH-LEVEL CONCEPTUAL ARCHITECTURE

6.1 INTRODUCTION

Chapter 5 focused on the system requirements elicitation process. This process derived a list of system requirements from a use case-driven approach and existing research gaps. The list of derived system requirements from Chapter 5 are taken as a pre-requisite input to guide the design and development of a high-level conceptual architecture.

Chapter 6 presents and discusses the high-level conceptual architecture along with all its components in Section 6.2. This architecture is an attempt to curb prevalent data leakage threats in the cloud. The author introduces and describes each of the architecture's main components and highlights how they work. Section 6.3 concludes this chapter and provides a high-level overview of what is to be discussed in the next chapter.

6.2 A HIGH-LEVEL CONCEPTUAL ARCHITECTURE FOR DATA LEAKAGE CONTROL IN THE CLOUD

This section presents and discusses the design of the high-level conceptual architecture for data leakage control in the cloud. The proposed high-level conceptual architecture consists of three core components (C): VM placement (C1); authentication (C2); digital forensic readiness (C3). As noted in the literature review in Chapter 3, existing work seems to be placing more focus on each of these components in isolation. Hence, a unique selling point of this thesis involves integrating all three components for a comprehensive solution. Figure 6.1 is an illustration that depicts what each component consists of and how the three components (i.e. C1, C2 and C3) are integrated into one. The integrated solution has an added value beyond that which is independently contributed by each of the components. Therefore, this thesis argues that a synergy of C1, C2 and C3 is better than the sum of its individual parts. Beyond C1, C2 and C3, the architecture consists of an infrastructure as a service (bottom), user (top) and CIA (confidentiality, integrity and availability) (right) layers.

The next subsections discuss each of the components in detail, starting with Infrastructure as a Service (IaaS) and followed by the users (C1, C2, C3), and CIA components.

6.2.1 Infrastructure as a Service (IaaS)

The proposed conceptual architecture focuses on the IaaS layer of cloud computing for the provision of virtual machines (VMs) to host users' data. The focus is placed on IaaS mainly because it is the most fundamental service model for cloud computing. IaaS is the foundation of all other service models in the cloud and all other service models are built on top of IaaS. For example, a CSP that specialises in Platform as a Service (PaaS) to provide a pay-per-use development platform requires the underlying virtual infrastructure of IaaS. The same is true for Software as a Service (SaaS). For instance, Google App Engine (a PaaS) requires the virtualised infrastructure of Amazon Cloud Formation (an IaaS) to work. Microsoft Office 360 (SaaS) requires the virtualised infrastructure of Microsoft Azure (IaaS) for it to work and have the necessary scalability and agility. Hence, it makes more sense for this thesis to tackle the data leakage threat at the lowest layer of the cloud computing stack, i.e. IaaS. Addressing this problem at the lowest layer of the cloud stack might help prevent it cascading to the PaaS and SaaS layers.

6.2.2 Users

The proposed architecture caters for two distinct user categories. The first category is for general users. General users are the users who make requests for VMs from the CSP to host their data. The CSP allocates a VM or a set of VMs to these users. Once allocated, the CSP then determines where to place the VM or set thereof on their infrastructure. This is done in a manner that is conflict-aware – avoiding placement with conflicting users. The users' access permissions are only restricted to their own VM instances or to those that belong to their organisations. For example, it could be that an organisation has its data hosted in a CSP and requires its users to have access to the VM instance holding the data.

The other category of users consists of digital forensic investigators, law enforcement agents and regulatory compliance officers. Since this category of users requires special access permissions for monitoring, investigation or compliance purposes, the architecture is so designed that they cannot tamper with the general users' VMs on the CSP's infrastructure. This is to avoid business disruptions and to improve the turnaround time of investigations. The digital forensic investigators, law enforcement agents and regulatory compliance officers operate in a somewhat covert channel (in the background). This is to ensure that their operations cannot raise alarms that would alert general users who are under investigation to delete their digital footprints or key digital evidence.

6.2.3 C1: VM Placement

C1 provides decision support for CSPs to facilitate the placement of users' VMs in the cloud infrastructure to avoid inter-VM attacks. An inter-VM attack is normally carried out by a malicious guest VM to compromise the confidentiality of virtualised resources of co-resident guest VMs (Someswar and Kalaskar, 2016; Amri et al., 2017). An inter-VM attack exploits vulnerabilities in the logical separation (hypervisor) layer between co-resident guest VMs that share the same cloud infrastructure. This attack uses virtualisation, which makes it possible for VMs of conflicting users to be instantiated on the same physical infrastructure (Elsayed and Zulkernine, 2015). For example, Perez-Botero et al. (2013) demonstrate a malicious attacker exploiting vulnerabilities in a memory management unit in order to compromise the confidentiality of data on co-resident guest VMs. Other researchers have already demonstrated how a malicious user's VM could exploit vulnerabilities in the hypervisor to trick it to issue a command that could destroy another co-resident's VM (Ristenpart et al., 2009, cited in Gonzales et al. (2017); Amri et al., 2017).

The challenge of dealing with inter-VM attacks puts the discussions on VM placement at the heart of CSPs' decision-making process. Hence, this thesis argues that a proper placement of VMs, taking cognisance of a user's conflict tolerance level with respect to security, risk exposure and cost implications, can help to minimise the impact of inter-VM attacks. For example, VMs belonging to two directly conflicting tenants must not be co-located on the same physical host. The main research goal of C1 is to improve VM placement and facilitate physical separation in order to minimise data leakage threats posed by inter-VM attacks.

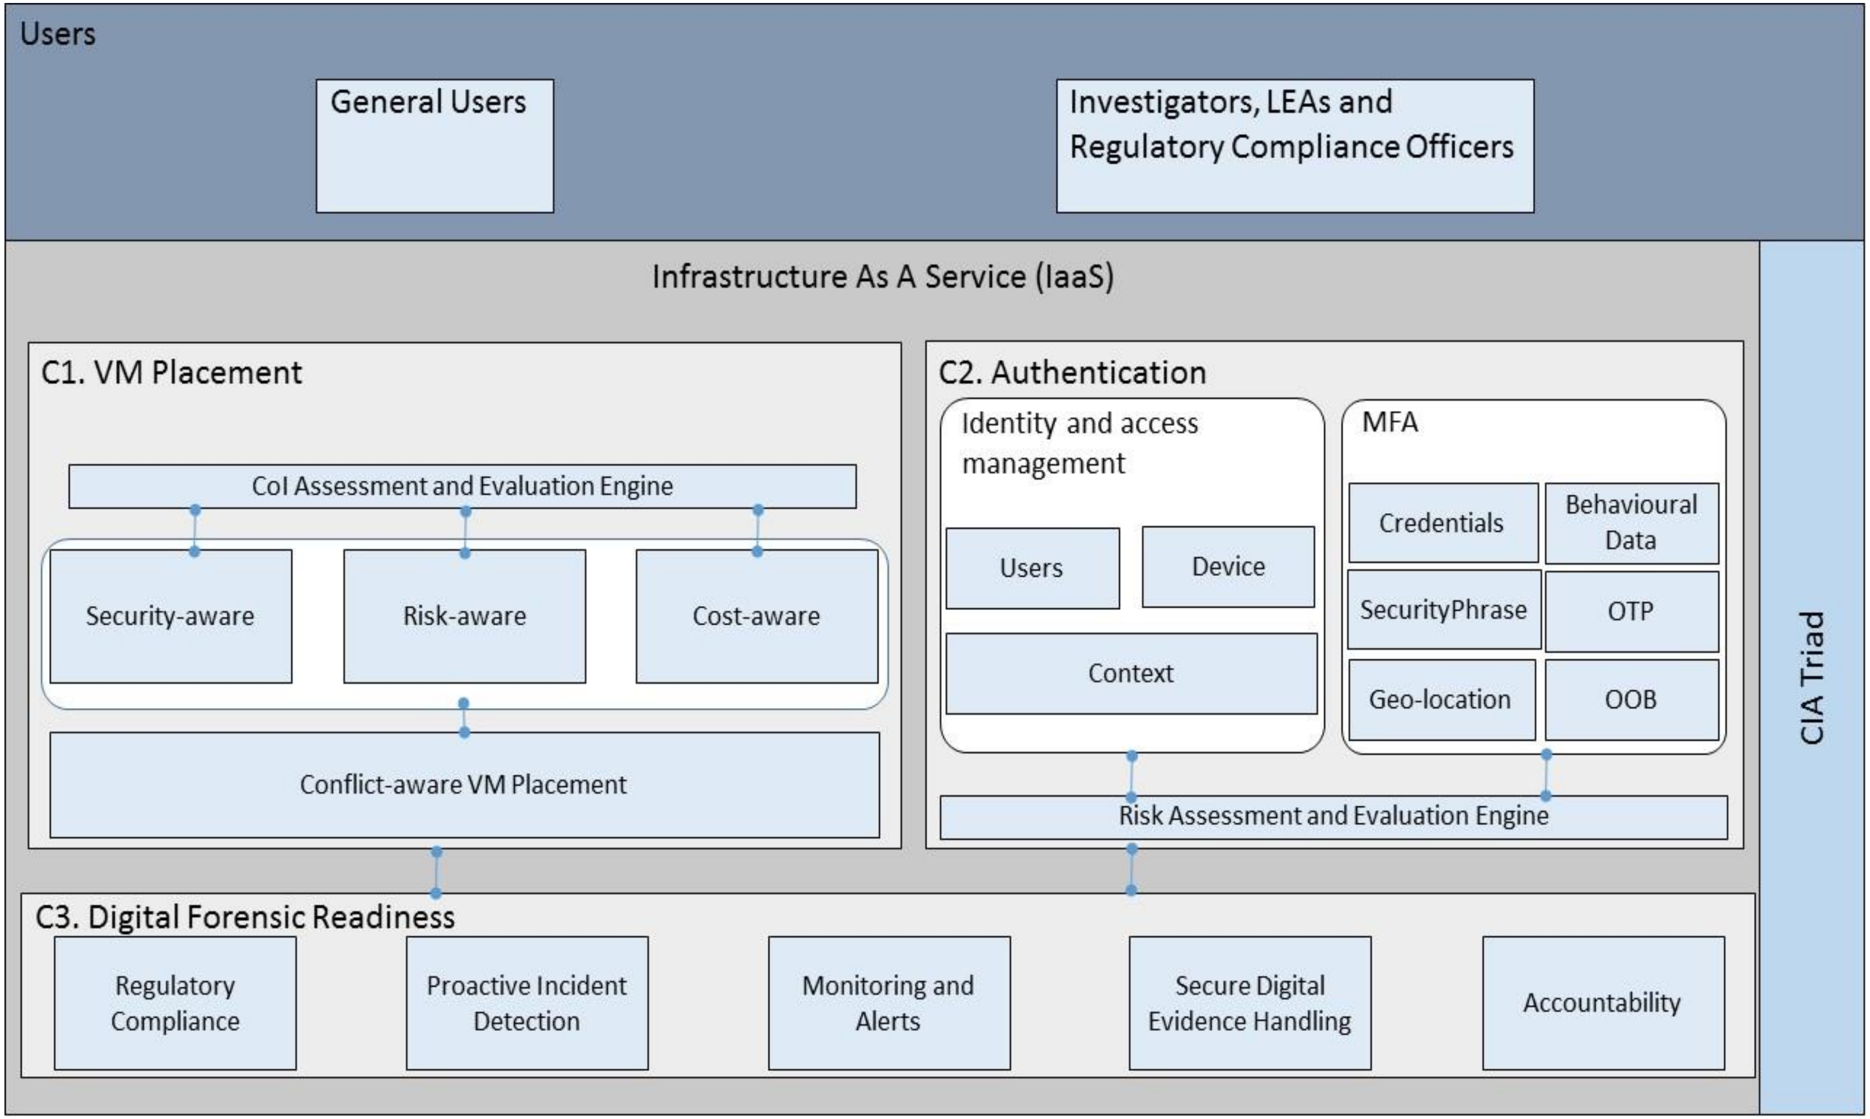
Users
General Users
Investigators, LEAs and Regulatory Compliance Officers
Infrastructure As A Service (IaaS)
C1. VM Placement
CoI Assessment and Evaluation Engine
Security-aware
Risk-aware
Cost-aware
Conflict-aware VM Placement
C2. Authentication
Identity and access management
Users
Device
Context
MFA
Credentials
Behavioural Data
SecurityPhrase
OTP
Geo-location
OOB
Risk Assessment and Evaluation Engine
C3. Digital Forensic Readiness
Regulatory Compliance
Proactive Incident Detection
Monitoring and Alerts
Secure Digital Evidence Handling
Accountability
CIA Triad

Figure 6.1:- A High-Level Conceptual Architecture for Data Leakage Prevention in the Cloud

C1 starts by collecting information from the user for assessing and evaluating their conflict tolerance level, determining with whom (among existing tenants) they have conflict and how far the conflict goes. This aggregated information is accessible only to the cloud service provider and not to the individual tenants, and it is required to determine the security, potential risk and cost of VM placements. These three – security, risk and cost – are then taken as input in the final conflict-aware VM placement.

6.2.4 C2: Authentication

C2 provides support for identity and access management, and for risk-based multifactor authentication (MFA). Hence, C2 consists of three parts: identity and access management; risk assessment; MFA. The idea is to take a risk-based approach to appropriately identify and verify users' identity and access based on their context and access locations. The identity and access management collect data on the users, their devices and context. User data includes credentials, security phrase, and keystroke behavioural patterns, while device data consists of device identifier and frequency of use. Context data includes data on geo-location of access points, place and time stamps of normal access. A decision to grant access takes user, device and context data to a risk engine that assesses the risk and provides a risk score. The risk score is taken as input in decision making in the MFA, which decides what factors to consider in authenticating a user based on the risk score.

For example, a low risk score is obtained when a known user makes an access attempt from their normal access location using a device that is already registered, and when their keystroke behavioural patterns match those found in their profile. Such a scenario and risk score would prompt the MFA to require a prospective user to make use of minimum access details, namely credentials (i.e. username and password) only. However, if the risk is evaluated as high, the MFA would prompt the user for multiple factors, i.e. their credentials, SecurityPhrase, an OTP and an OOB, whilst in the background the system verifies access geo-location and user behavioural data. Hence, C2 considers a risk-based MFA that uses multiple authentication factors to scale up and down in order to adapt access decisions based on the risk scores of users, their devices and contexts.

It must be noted though, that some of the access decisions happen in the background and are not visible to the user. For example, device authentication occurs without the user's knowledge.

The same goes for users' contextual, geo-location and behavioural data. This is to make sure that the proposed solution does not burden users with extra levels of detail, but enhances user experience without compromising the security of their resources. This stance somehow eases the security and usability trade-off mentioned in Martim et al. (2009).

6.2.5 C3: Digital Forensic Readiness

C3 provides a digital forensic readiness capability that proactively captures and preserves potential digital evidence. This is done in a legally and forensically sound manner in anticipation of a potential legal or corporate investigation or lawsuit. C3 consists of the following parts: regulatory compliance; proactive incident detection; monitoring and alerts; secure digital evidence handling; and accountability. These are discussed below.

Regarding the regulatory compliance component, a digital forensic readiness capability is required to proactively capture and preserve potential digital evidence according to applicable regulatory compliance frameworks. It is important to identify applicable regulatory frameworks that govern CSPs' hosting of their users' data. A legal expert is required to facilitate this component and to ensure that capturing and preserving potential digital evidence is done according to the governing laws. Involving a legal expert in the regulatory compliance helps to ensure that all employees and third parties are aware of how to handle potential digital evidence; how to preserve the digital evidence's chain of custody; how not to contaminate the digital evidence at collection, preservation, storage, transportation and manipulation; and for how long they should preserve digital evidence in accordance with legal mandates and internal retention policies. This helps to improve admissibility of digital evidence in courts or in corporate investigations. Furthermore, it also demonstrates due diligence for accountability with regard to regulatory compliance.

Proactive incident detection helps to detect incidents before they could occur. This component collects potential digital evidence and does minor processing to correlate digital forensic evidence with the sole purpose of identifying malicious events that could potentially lead to an incident. However, the processing does not necessarily make use of machine learning algorithms. Incorporating machine learning into this component of the study is left as future work. If event correlation somehow picks up that an incident is about to happen, the system immediately alerts a digital forensic investigator to look further into the suspicious incident.

This takes us to the monitoring and alerts part, which monitors the activities of the event correlation process for alerts on suspicious activity. Serious alerts are pushed to investigators in order to prevent incidents from happening in the first place, rather than dealing with the aftermath.

Secure digital evidence handling ensures that all collected potential digital evidence is secure from the point of acquisition up to its presentation in court. This component of C3 ensures that potential digital evidence is not exposed to unauthorised users, to unauthorised alterations, deletion or leakage, which might bring it into disrepute and render it inadmissible in court. End-to-end secure handling of potential digital evidence also ensures a proper chain of custody. A chain of custody is very crucial to demonstrate that potential digital evidence has been properly handled and there are no loopholes. This is normally used to prove the integrity of potential digital evidence along with timestamps. It helps to show the chronological order of how potential digital evidence has moved from one point to another, who had access to it when, as well as what actions were performed on it.

Accountability ensures that access to potential digital evidence in the secure remote repository is closely monitored for all investigators, law enforcement agents and compliance officers. This somehow facilitates keeping a proper chain of custody in order to preserve the integrity of evidence. Furthermore, this component is meant to detect and report unauthorised access and tampering with evidence by authorised users.

Bringing it all together, C3 (the entire digital forensic readiness capability) is intended to proactively acquire potential digital evidence prior to an investigation. The main idea is to facilitate event reconstruction, guarantee proper chain of custody and ensure admissibility of potential digital evidence when required. Moreover, it helps anticipate, identify, detect and respond to unauthorised or malicious user activity. It also monitors access by authorities for accountability and non-repudiation.

6.2.6 CIA Triad

The proposed architecture is basically meant to ensure that user data and potential digital evidence are secure against attacks intended to compromise their **c**onfidentiality, **i**ntegrity and **a**vailability. This is referred to as the CIA triad (Mosenia and Jha, 2017) in the cybersecurity space. Confidentiality refers to ensuring that data is accessed by authorised users in an authorised manner. For example, an authorised user can access a read-only object for reading purposes only. If it so happens that such a user manipulates their access permissions and ends up with the ability to write on the object, such an action would compromise the object's confidentiality and integrity. **C**onfidentiality is compromised because even though the user is authorised to access the object, their action of writing is not authorised. Furthermore, modifying an object in an unauthorised manner compromises its integrity. **I**ntegrity refers to ensuring that an object's originality is preserved from unauthorised modification. **A**vailability on the other hand refers to ensuring that objects are availed to authorised users when they are required.

The architecture preserves this triad in different aspects. For example, C1 and C2 are to a greater extent meant to guarantee confidentiality and integrity of data as well as potential digital evidence. The IaaS layer is to a larger extent responsible for availability. Since this thesis is meant to address data leakage threats in the cloud, its focus is more on confidentiality than on the other two (integrity and availability). C3 in turn addresses the integrity aspect of the CIA triad.

The next section concludes this chapter and highlights the contents to be dealt with in the next chapter.

6.3 CONCLUSION

Chapter 6 took the requirements raised in Chapter 5 as input to outline and design the high-level architecture for addressing data leakage threats on IaaS cloud infrastructure. The main goal was basically to present and discuss each of the components of the architecture and their building blocks. The roles of each component were discussed, as well as the relationships between the components. A detailed discussion of the relationships will be provided in later chapters.

The next three chapters, i.e. 7, 8 and 9 will discuss the approaches adopted to dissect C1, C2 and C3 respectively. These three chapters, along with Chapter 6, constitute the solution – and therefore the main contribution – as proposed in this thesis.

CHAPTER 7 CONFLICT-BASED VM PLACEMENT (CBVMP) MODEL

7.1 INTRODUCTION

Chapter 6 presented and discussed a high-level architecture for addressing data leakage threats on an IaaS cloud infrastructure. It discussed each of the components of the proposed architecture and their respective building blocks on a high level.

Chapter 7 now focuses on the VM placement component and in particular, it introduces the so-called Conflict-based VM Placement, in short CBVMP, model. This chapter discusses in detail the intricacies of how this component is modelled. The structure of chapter 7 is as follows: Section 7.2 discusses a theoretical formulation of the VM placement problem. Section 7.3 derives and formulates the mathematical VM placement problem. Section 7.4 provides a process flow diagram and presents the VM placement algorithms for implementing the CBVMP model. Section 7.5 presents the screenshots of the implementation of the model. Finally, Section 7.6 concludes this chapter and highlights the focus of Chapter 8.

7.2 A THEORETICAL FORMULATION OF THE VM PLACEMENT PROBLEM

An IaaS CSP uses a conflict-based VM Placement (CBVMP) model to place VMs. In this thesis, the VM placement problem is modelled based on the following statement: a virtual machine (*VM*) is contained in a physical node (*PN*), which is contained in a cluster (*Clu*) that is contained in a data centre (*DC*), and which is at a specific geographical location (*Loc*). However, the relationships between each of these are not necessarily of a one-to-one type. For example, one *PN* may hold more than one *VM* and a *Clu* may hold more than one *PN* (as depicted in Figure 7.1).

$$VM \subseteq PN \subseteq Clu \subseteq DC \subseteq Loc$$

The relationship between a tenant (*te*), a functional business domain (*FBD*) and a conflict of interest (*CoI*) is as follows: a *te* is a member of a *FBD*. The *FBD* is associated with a specific *CoI* class.

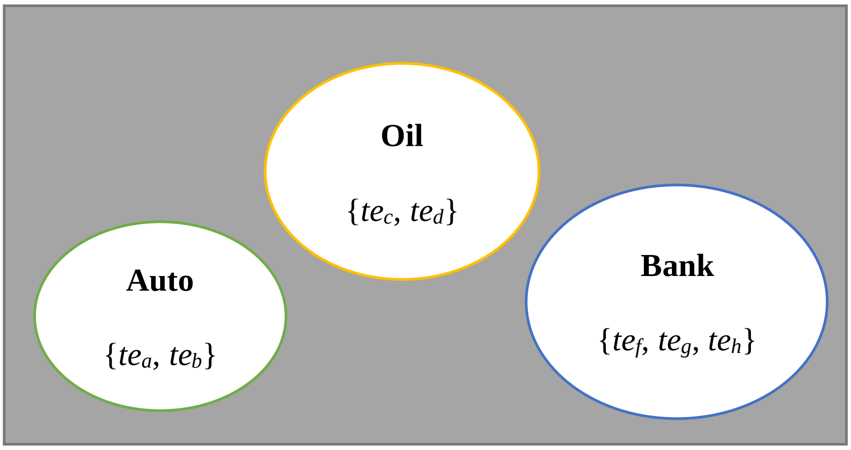

Figure 7.1:- The Relationships Between Tenants, *FBDs* and CoI classes

For example, consider three *FBDs* (i.e. Auto with te_a and te_b; Oil with te_c and te_d, and Bank with te_f, te_g and te_h) as shown in Figure 7.1. In this example, te_c and its direct competitor te_d belong to the same *CoI* class, i.e. Oil. This means that these two are in direct conflict with one another. Consequently, requests by te_c and te_d for *VM* placement in a CSP must be handled in such a way that the data leakage threats posed by inter-VM attacks are minimised. For example, the hosting CSP must ensure that VMs belonging to these two directly conflicting tenants are adequately isolated and not allowed to share the same *PN* or *Clu* infrastructure, so as to avoid inter-VM attacks. The data leakage threat is more pronounced if VMs that belong to conflicting tenants happen to co-reside. Hence, this part of the thesis aims to address the problem of inter-VM attacks by ensuring that conflicting tenants' VMs are isolated and spread further apart. Figure 7.2 shows the structure of a CSP's IaaS infrastructure as proposed in this thesis. In this structure, it is a CSP that spans multiple geographical *Locs*. Each *Loc* comprises of multiple *DCs*. Each *DC* comprises of multiple *Clus*. Each *Clu* comprises of multiple *PNs*. Each *PN* contains multiple *VMs* as depicted in Figure 7.2.

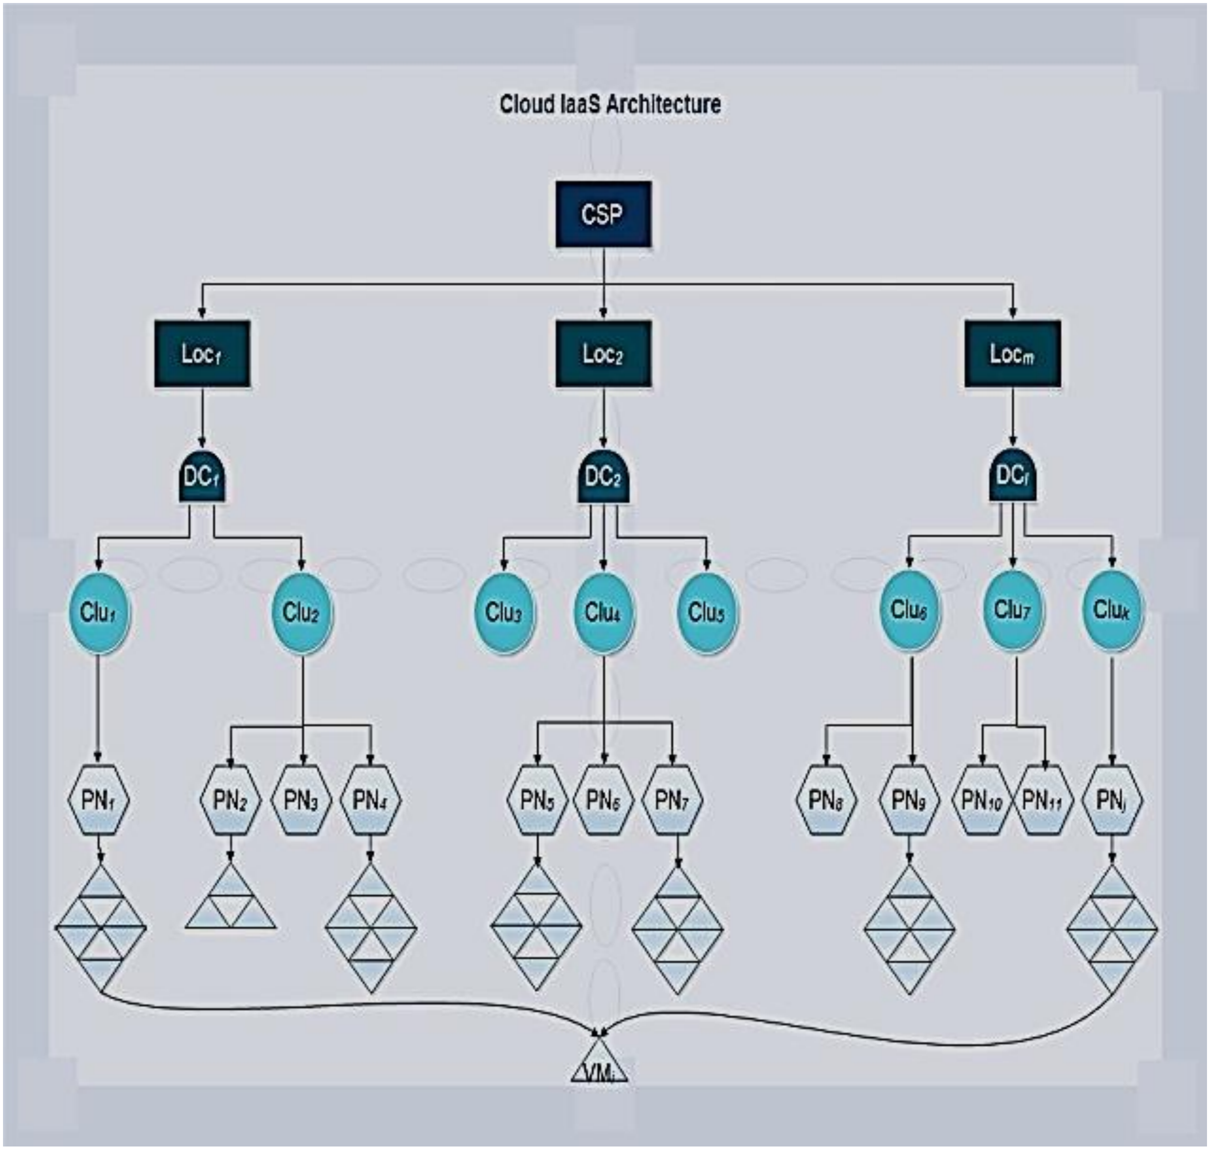

Figure 7.2:- Cloud IaaS Architecture

7.3 MATHEMATICAL FORMULATION OF THE VM PLACEMENT PROBLEM

For a new *VM* placement, a tenant (te_i) provides four inputs:

(1) A tenant identification (te_{i_ID})

(2) A Conflict Tolerance Level (*CTL*)

(3) A Functional Business Domain (*FBD*), which determines a *CoI* class

(4) Size in terms of gigabytes (GB) of a requested *VM* (*sizeOfVM*) that is directly proportional to the resource capacity constraint.

The tenant identification (te_{i_ID}) is a unique identifier for each tenant. A CTL determines how much conflict a tenant te_{i_A} can tolerate for being hosted on the same infrastructure with a conflicting tenant te_{i_B}. The size of a requested VM refers to the storage resource capacity. Each *VM*, *PN*, *Clu*, *DC* and *Loc* at time *t* is subject to resource capacity constraints, i.e. $c_i(t)$, $c_j(t)$, $c_k(t)$, $c_l(t)$ and $c_m(t)$ respectively.

It is also important to note that even though, by default, a tenant would belong to a specific CoI class, they may (over and above this) decide to add more potential competitors to their list. This list then is defined as a sphere of conflict and associated with a sphere of non-conflict. These are defined in the next section. The concepts of a sphere of conflict and of non-conflict are taken from the work of Loock and Eloff (2005).

7.3.1 Sphere of Conflict and Sphere of Non-Conflict

Each tenant te_{i_A} has its own Sphere-of-Conflict (SoC_a) and Sphere-of-non-Conflict ($SonC_a$) sets (Loock and Eloff, 2005). The SoC_a and $SonC_a$ sets are identified by the CSP after obtaining the CTL from a potential tenant. The Sphere-of-Conflict set for a tenant te_{i_A}, is denoted as:

$SoC_a = \{te_{i\_1}, \ldots.\ te_{i\_n}\}$

where:

te_{i_1}, …..te_{i_n} is a set of conflicting tenants to te_{i_A} that uses the same CSP.

For each $VM_a \in te_{i\_A}$ in the set

$SoC_a = \{te_{\ i\_1}, \ldots.\ te_{\ i\_n}\}$ is associated with an address $Addr_a$ in the following form:

$Loc_m\,.\,DC_l\,.\,Clu_k\,.\,PN_j\,.\,VM_i$. This is written similar to an IP address separated by dots.

A non-placement matrix is the list of all addresses for all *VMs* in the SoC_a set that are hosted in the CSP. This is a list of all addresses where a VM cannot be placed because there is a conflicting tenant. For example,

$$Loc_1\,.\,DC_1\,.\,Clu_1\,.\,PN_1\,.\,VM_1$$
$$Loc_1\,.\,DC_1\,.\,Clu_1\,.\,PN_1\,.\,VM_2$$
$$Loc_2\,.\,DC_1\,.\,Clu_3\,.\,PN_3\,.\,VM_4$$
$$Loc_2\,.\,DC_1\,.\,Clu_3\,.\,PN_3\,.\,VM_5\,.$$

This set is only visible to the CSP. For a new VM placement, this set could be reduced by applying a CTL to eliminate some addresses from the union set SoC_a. Therefore, a CTL is linked to each VM_a of tenant te_{i_A}.

The proposed CBVMP model allows CSPs to co-host a tenant's VMs with those of other conflicting tenants. This is referred to as non-optimal placement. Such placement is considered non-optimal because it does not guarantee adequate physical separation between the conflicting tenants' VMs. Tenants could, however, opt for non-optimal placement of their VMs for various reasons, including lower cost and hosting of public data.

Tenants could also opt not to host their VMs with conflicting tenants. This is called optimal VM placement. Optimal VM placement implies that tenants can only co-host their VMs with non-conflicting tenants to ensure adequate physical separation. The set of all non-conflicting tenants is called a Sphere-of-non-Conflict.

The Sphere-of-non-Conflict set is denoted as:

$SonC_a=\{te_{i\_o}, \ldots, te_{i\_z}\}$

where:

$te_{i\_o}, \ldots, te_{i\_z}$ is a set of non-conflicting tenants to te_a that uses the same CSP.

For each $VM_a \in te_{i\_A}$ in the set

$SonC_a = \{te_{i\_o}, \ldots.\ te_{i\_z}\}$ is associated with an address $Addr_a$ in the following form: $\mathrm{Loc}_m\ .\ \mathrm{DC}_l\ .\ Clu_k\ .\ PN_j\ .\ VM_i$. This address is written similar to an IP address separated by dots.

A placement matrix is a list of all addresses for all VMs in the set $SonC_a$ that are hosted in the CSP where placement is possible. Placement is possible only because there are no conflicting tenants on the addresses. An example would be as,

$$Loc_2\ .\ DC_2\ .\ Clu_1\ .\ PN_1\ .\ VM_1$$
$$Loc_1\ .\ DC_1.\ Clu_1\ .\ PN_1\ .\ VM_2$$
$$Loc_3\ .\ DC_4\ .\ Clu_3\ .\ PN_3\ .\ VM_4$$
$$Loc_3\ .\ DC_2\ .\ Clu_1\ .\ PN_1\ .\ VM_3\ .$$

Following the definitions of *SoC* and *SonC*, it is essential to show the relationship of *CoI* as proposed in this thesis. This is to ensure that the reader is aware of how the thesis models the CoI construct.

7.3.2 Conflict-of-Interest Relationship

A *CoI* relation is defined as neither an equivalence nor a binary relationship:

- Not reflexive: if ($te_a\ \neg CoI\ \ te_a$) $\Rightarrow$ te_a cannot be in conflict with itself
- Symmetric: if (te_a *CoI* te_b) $\Rightarrow$ (te_b *CoI* te_a)

  $\Rightarrow$ the *CoI* relation is mutually inclusive. If te_a is in conflict with te_b then te_b is in conflict with te_a.
- Not transitive: if (te_a *CoI* te_b *CoI* te_f)

  $\Rightarrow \neg$ (te_a *CoI* te_f) $\Rightarrow$ if te_a is in conflict with te_b who is in conflict with te_f, it does not imply that te_a is in conflict with te_f.

Given the relationship of the *CoI* construct, the next section defines the three basic variables of the model.

7.3.3 Risk, Physical Security and Cost

The proposed model is cognisant of the fact that tenants can experience different degrees of conflict with other tenants. Hence, the risk (*R*) of confidential data leakage posed by inter-VM attacks also varies, depending on the degrees of conflict. For example, co-residence of *VM*s that belong to two directly competing tenants that are in direct conflict of interest with one another, poses the highest *R*.

Furthermore, minimal *CoI* is regarded as no conflict. This defines another sphere for $te_{\ i\_A}$, which is referred to as a Sphere-of-non-Conflict ($SonC_a$) (as reflected above). In the instance $SonC_a$={te_{i_o},…,te_{i_z}}, where $te_{\ i\_o}$,…,$te_{\ i\_z}$ are identified by the CSP as non-conflicting tenants or tenants with insignificant *CoI* to that of $te_{\ i\_A}$. Ideally, these could be the only tenants that $te_{\ i\_A}$ would prefer to co-reside with. However, CSPs are likely to charge a high cost (*C)* for non-conflict placement of tenants' *VM*s and relative low *C* for flexible co-resident placement with conflicting tenants.

The physical separation (*S*) of *VMs* is determined by a *CTL*. *S* indicates how much physical separation (i.e. co-reside in the same *PN* or reside in a different *PN* within the same *Clu*; in the same *Clu* but different *DC*; in the same *DC* but different *Loc*) a tenant requires between its own *VM*s and that of its competitors. This is followed by determining a cost value *C* for implementing a *CTL*. A risk exposure level *R* is associated with a *CTL*. The next section reflects on how the three variables, *S, C* and *R*, come together to formulate a utility function.

7.3.4 Utility Function *U*

CBVMP introduces a utility function *U* for building some type of functional dependency between the variables *S*, *R* and *C*. *U* computes the overall physical separation of conflict-aware VM placement, based on *S, C* and *R* that a tenant requires. The model uses *SoC* and *SonC* sets as derived by the CSP from a client's inputs of *CTL* and *FBD*. The output is a set (i.e. either *SoC* or *SonC*) of all possible *VM* placement options. These two sets, *SoC* or *SonC*, are revised by taking into consideration the variables S, R and C for the CBVMP model to deliver a reduced set of *VM* placements options. The result is used by the CSP to identify all potential nodes where a tenant's *VM*s could be placed. The reduced set of *VM* placement options is then presented to the tenant for it to make an informed decision on where it would like its *VM*s to be placed by the CSP and at what cost, risk and security. However, the final placement decisions are computed using an objective function that maximises the utility function that takes *S, C* and *R* as input. The objective function is discussed in the next section.

7.3.5 The Objective Function

The main objective from a potential tenant's perspective is to minimise *R* and *C* whilst maximising *S* and eventually *U* of physically separating clients' VMs on the CSP's infrastructure. This results in a multi-objective optimisation approach that, if applied effectively, will address the data leakage threats posed by inter-VM attacks in the cloud. Furthermore, this multi-objective optimisation is subject to a number of constraints, i.e. resource capacity, conflict of interest, tenants' *CTL*, *S*, *C* and *R*.

Hence, the VM placement problem herein is framed as a multi-dimensional bin-packing problem (Hatzopoulos et al., 2013; Wu et al., 2014). A multi-dimensional bin-packing problem is formulated as a vector of *Loc* to *DC* to *Clu* to *PN* and *VM* bins.

The proposed solution has the following binary decision variables: x_{ij}(t), y_{jk} (t), z_{kl} (t) and w_{lm}(t). The variable t indicates the time at when a *VM* placement is done. The binary decision variables in this case take a value of zero or one. These variables are equal to one when placement has been done and otherwise to zero. $x_{ij}(t)$ denotes that vm_i is placed on $PN_j(t)$; $y_{jk}(t)$ denotes that PN_j is placed in $Clu_k(t)$, $z_{kl}(t)$ denotes that Clu_k is placed in a $DC_l(t)$; and w_{lm}(t) denotes that $DC_l(t)$ is placed in $Loc_m(t)$ at a particular time t. Each binary decision variable is associated with three coefficient weight vectors, i.e.

α_{i1}={ α_{11}, α_{21}, α_{31}, α_{41}} $\forall i$ = {1,..,4} relates to the *S* vector.

β_{j2}={ β_{12}, β_{22}, β_{32}, β_{42}} $\forall j$ = {1,..,4} relates to the *C* vector.

γ_{k3}={ γ_{13}, γ_{23}, γ_{33}, γ_{43}} $\forall k$ = {1,..,4} relates to the *R* vector.

α_{i1}={ α_{11}, α_{21}, α_{31}, α_{41}}, β_{j2}={ β_{12}, β_{22}, β_{32}, β_{42}}, γ_{k3}={ γ_{13}, γ_{23}, γ_{33}, γ_{43}} are the coefficient weights for each binary decision variable x_{ij}(t), y_{jk}(t), z_{kl}(t) and w_{lm}(t) respectively. For example, x_{ij}(t) has a coefficient weight vector {α_{11}, β_{12}, γ_{13}} where α_{11} is the coefficient weight in terms of physical separation security; β_{12} is the coefficient weight in terms of cost, and γ_{13} is the coefficient weight in terms of risk posed on a VM_i being placed in PN_j. y_{jk}(t) has a coefficient weight vector {α_{21}, β_{22}, γ_{23}} where α_{21} is the coefficient weight in terms of physical separation security; β_{22} is the coefficient weight in terms of cost and γ_{23} is the coefficient weight in terms of risk posed on a PN_j being placed in Clu_k. Furthermore, z_{kl}(t) has a coefficient weight vector {α_{31}, β_{32}, γ_{33}} where α_{31} is the coefficient weight in terms of physical separation security. β_{32} is the coefficient weight in terms of cost and γ_{33} is the coefficient weight in terms of risk posed on a Clu_k being placed in DC_l. The same goes for w_{lm}(t) with coefficient weight vector { α_{41}, β_{42}, γ_{43}} in terms of risk posed on a DC_l being placed in Loc_m.

This marks the end of the process of deriving the building blocks (elements) of the model.

The problem formulation of the above can be modelled mathematically as an optimisation problem as follows:

$$Max U(S,C,R) = \sum_{i=1}^{m} \begin{pmatrix} \alpha_{11} \\ \beta_{12} \\ \gamma_{13} \end{pmatrix} x_{i,j} + \sum_{j=1}^{n} \begin{pmatrix} \alpha_{21} \\ \beta_{22} \\ \gamma_{23} \end{pmatrix} y_{j,k} + \sum_{k=1}^{o} \begin{pmatrix} \alpha_{31} \\ \beta_{32} \\ \gamma_{33} \end{pmatrix} z_{k,l} + \sum_{l=1}^{p} \begin{pmatrix} \alpha_{11} \\ \beta_{12} \\ \gamma_{13} \end{pmatrix} w_{l,m} \qquad (7.1)$$

The problem formulation is subject to a number of constraints, in other words CTL is expressed in terms of *S* - α_{ij}, *C* - β_{j2} and *R* - γ_{k3}, size (*s*) of a *VM, PN, Clu, DC* and *Loc* and capacity *c(t)* of *PN, Clu, DC* and *Loc*. Note that the capacity of the actual VM is excluded in these constraints. This is because *sizeOfVM* is taken as input and considered from the onset when the CSP is determining potential placement addresses. The constraints are modelled as follows:

$$\sum_{i=1}^{m}\sum_{j=1}^{n}\alpha_{ij}.x_{ij}(t)\leq CTL_i(t) \qquad (7.2)$$

$$\sum_{i=1}^{m}\sum_{j=1}^{n}\beta_{ij}.x_{ij}(t)\leq CTL_i(t) \qquad (7.3)$$

$$\sum_{i=1}^{m}\sum_{j=1}^{n}\gamma_{ij}.x_{ij}(t)\leq CTL_i(t) \qquad (7.4)$$

$$\sum_{i=1}^{m}\sum_{j=1}^{n}s_{ij}.x_{ij}(t)\leq c_j(t)$$

$$\sum_{i=1}^{m}x_{ij}=1 \qquad \forall i \in \{1,\ldots,m\} \qquad (7.5)$$

$$\sum_{j=1}^{n}\sum_{k=1}^{o}s_{jk}.y_{jk}(t)\leq c_k(t)$$

$$\sum_{j=1}^{n}y_{jk}=1 \qquad \forall j \in \{1,\ldots,n\} \qquad (7.6)$$

$$\sum_{k=1}^{o}\sum_{l=1}^{p}s_{kl}z_{kl}(t)\leq c_l(t)$$

$$\sum_{k=1}^{o}z_{kl}=1 \qquad \forall k \in \{1,\ldots,o\} \qquad (7.7)$$

$$\sum_{k=1}^{p}\sum_{l=1}^{q}s_{lm}w_{lm}(t)\leq c_m(t)$$

$$\sum_{k=1}^{p}w_{lm}=1 \qquad \forall l \in \{1,\ldots p\} \qquad (7.8)$$

$$x_{ij}(\mathrm{t})\in\{0,1\} \qquad \forall i \in \{1,\ldots,m\},\ \forall j \in \{1,\ldots,n\} \qquad (7.9)$$

$$y_{jk}(\mathrm{t})\in\{0,1\} \qquad \forall j \in \{1,\ldots,n\},\ \forall k \in \{1,\ldots,o\} \qquad (7.10)$$

$$z_{kl}(\mathrm{t})\in\{0,1\} \qquad \forall k \in \{1,\ldots,o\},\ \forall l \in \{1,\ldots,p\} \qquad (7.11)$$

$$w_{lm}(\mathrm{t})\in\{0,1\} \qquad \forall l \in \{1,\ldots,p\},\ \forall m \in \{1,\ldots,q\} \qquad (7.12)$$

The multi-objective function (7.1) ensures a maximum utility (*U)* of placing each *VM* in a physical node *PN* (x_{ij}); a PN is placed in a cluster *Clu* (y_{jk}); a *Clu* is placed in a data centre *DC* (z_{kl}); and a *DC* is placed in a location *Loc (*w_{lm}) that belongs to a CSP. This multi-objective function ensures that placement of tenants' *VM*s is done based on their cost (i.e. constraint (7.3)); conflict tolerance levels (i.e. constraints (7.2) and (7.4)); risk appetite (i.e. (7.2) and (7.4)); and resource capacity constraints (i.e. (7.5) – (7.8)). Resource capacity constraints (7.5) – (7.8) ensure that vm_i must not exceed the capacity $c_j(t)$ of PN_j, $c_k(t)$ of Clu_k, $c_l(t)$ of DC_l and $c_m(t)$ of Loc_m. The binary decision variables $x_{ij}(t)$, $y_{jk}(t)$, $z_{kl}(t)$ and w_{lm}(t) denote that vm_i is placed on PN_j, which is placed in Clu_k within a DC_l that is also placed at Loc_{m} at time *t*.

Constraints (7.9) – (7.12) ensure that the decision variables $x_{ij}(t)$, $y_{jk}(t)$, $z_{kl}(t)$ and w_{lm}(t) can either take the value one, which indicates placement, and otherwise zero for all *VM*s to be placed in *PN*s, for all *PN*s to be placed in *Clu*s, for all *Clu*s to be placed in *DC*s and for all *DC*s to be placed in *Loc*s. This marks the end of the modelling.

Based on the mathematical model framed as an optimisation problem, the next section introduces a process model. The process model illustrates all the processes that are required to implement the solution to the problem.

7.4 CBVMP PROCESS FLOW MODEL AND CONFLICT-AWARE VM PLACEMENT ALGORITHMS

Figure 7.3 illustrates a process flow diagram that demonstrates the proposed conflict-aware VM placement.

It starts off with prospective tenants requesting to host their data in a VM from a CSP. Depending on the size of the dataset, the CSP takes the request, creates a VM and allocates the necessary resources as specified by the prospective tenant. The CSP then uses the CBVMP conflict-aware VM placement model to select an appropriate *PN* within an appropriate *Clu* held in an appropriate *DC* that ensures optimal physical separation of VMs from conflicting tenants.

The conflict-aware VM placement model makes use of the Best Fit heuristic algorithm (Kernaghan and Lin, 1970; Burke, Kendall and Whitwell, 2009) to choose placement of a particular VM between available *PNs, Clus, DCs* and *Locs.* A Best Fit heuristic searches a reduced list of potential placement options as identified by the model for the best-fitting address. The reduced list of potential placement addresses would already have excluded all conflicting addresses and those with insufficient space. The candidate VM is thus placed at the best-fitting address. Should there be more than one best-fitting address, the algorithm chooses the first best-fitting address. The placement model manages the placement matrix and creates a GPS point to mark the geographical location (physical) of the newly created *VM* on a map. The placement model then issues a certificate of placement with the full address and allocated resources of the *VM* and sends it back to the tenant. The results are not necessarily communicated to the rest of the tenants that share a *PM*. The map on the tenant's side gets updated automatically with the new location of its *VM*. The CSP records the address of final placement. It is possible for other co-resident tenants to decide for example to request a migration because of a new tenant's *VM* sharing the same PN with them.

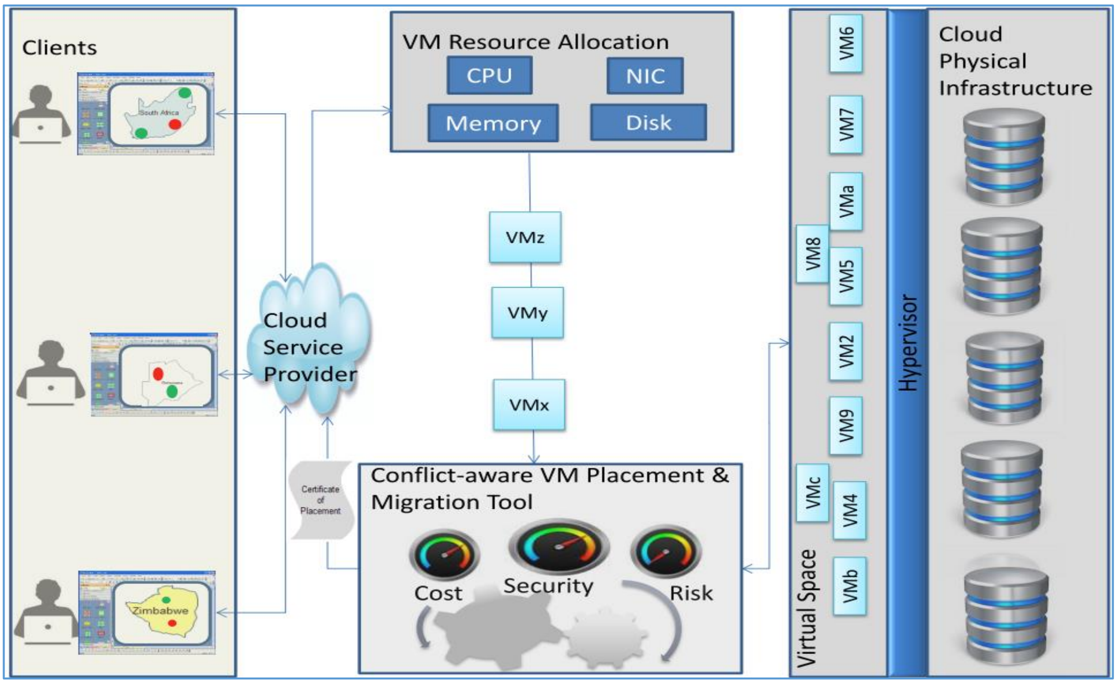

Figure 7.3:- CBVMP Process Flow Diagram

Figure 7.3 illustrates the process flow of the proposed CBVMP model. The actual CBVMP model is implemented using four algorithms after the VMs have already been created. Algorithms 1 and 2 focus on placing a new VM. Algorithms 3 and 4 focus on migrating a previously allocated VM from one address to another. This may be necessitated by a CSP conducting a routine load-balancing exercise. It could also be necessitated by a tenant that has for instance re-classified their previously allocated VM with a high *CTL* to a lower one. Migration may also be required when co-resident tenants do not like to share the same hardware with other tenants. Furthermore, since placement of VMs with a higher *CTL* comes at a higher cost, tenants might exercise their right to reduce the *CTL* in order to pay less. The CSP must take means to control this to avoid excessive VM hopping. However, this issue falls outside the scope of this thesis. The next four subsections discuss each of these algorithms in detail.

Algorithm 1: Optimal Placement of a New VM

Algorithm 1 considers a new tenant whose VM is to be placed for the first time and that requires to be placed only with non-conflicting tenants.

Algorithm 1 Optimal VM Placement Algorithm

```
Require:
      $te_{id}$, $FBD$, $CTL$, $sizeOfVM$, $p\_Type$
Ensure:
      $SonC\{te_o,...te_z\}$
if ($p\_Type = optimal$) then                                  ▷ Optimal vm placement
    $SonC\{\} \leftarrow \{te_o,..te_z\}$
    for all $te_i \in SonC\{te_o,...te_z\}$ do
        allNCAddresses$\{\} \leftarrow \{te_o,..te_z\}$
        $i \leftarrow o$
        for all $addr_i \in$ allNCAddresses$\{te_o,..te_z\}$ do
            if $(addr_i) \geq sizeOfVM$ then
                sufNCAddresses$\{\} \leftarrow addr_i$
            end if
            $i \leftarrow i + 1$
            for all $addr_i \in$ sufNCAddresses$\{te_o,..te_{z-n}\}$ do
                Find bestFitAddress($addr_i$)
                Place $vm_{ix}$ in bestFitAddress($addr_i$)
                Return $vm_{ix}.addr_i$
                Update AllocPlaceMatrix($vm_{ix}.addr_i$)
                Update visibility.Map($vm_{ix}.addr_i$)
            end for
        end for
    end for
end if
```

Algorithm 1: Placing a new VM in an optimal manner

A tenant te_{i_A} makes a request to host its data on a newly allocated VM. A CSP would first check the tenant's *ID* (te_{i_A}), and get the *CoI*, *CTL* and size of the VM. Based on the *CTL*, a CSP would then determine if the placement is optimal or non-optimal. An optimal placement refers to placement that cannot tolerate any conflict. A popular view is that zero tolerance means something cannot be tolerated. However, this study's approach is a complete opposite to the popular view, in that a *CTL* of zero means insignificant conflict or no conflict at all. This is for the lowest tolerance level. A non-optimal placement refers to a somewhat flexible placement that can tolerate certain degrees of conflict as specified in the *CTL*. Algorithm 1 considers an optimal placement. Hence, the CSP would provide a list of all existing tenants that are not in conflict with the prospective tenant – referred to as a *SonC* set. As mentioned before, this set is only visible to the CSP and not to the tenant, for the sake of privacy of existing tenants and to ensure that tenants are not aware of what VMs are placed where. It is not advisable for a tenant to be able to map out placements of other tenants. This is because such a mapping would make it easier for malicious tenants to identify their target and then choose their preferred address of placement to share a *PN* with that of their target. Exposed mapping would increase the chances of data leakage through an inter-VM attack. Such a scenario would nullify all the efforts of isolating VMs that belong to conflicting tenants. The next step is to locate and list all addresses of *VM*s owned by each of these tenants. From this list of addresses (i.e. Loc_m . DC_l . Clu_k . PN_j), Algorithm 1 determines if the host (i.e. *PN*) has sufficient space to host the tenant's VM. From this improved set, the algorithm further determines the addresses of the potential hosts to best place the new VM. Once the best potential host is found, the algorithm places the VM. This then sends the address of the new VM placement to the tenant and updates the placement matrix and geographical map.

Algorithm 2: Non-optimal Placement of a New VM

Algorithm 2 considers a new tenant whose VM is to be placed for the first time – with a requirement to co-reside with conflicting tenants of a specified CTL. Tenant te_{i_A} requests placement of its new VM and specifies a certain CTL. The specified CTL is less restrictive and allows some level of flexibility for the tenant to co-host with some conflicting tenants. The CSP checks the tenant's *ID* (te_{i_A}), and gets the FBD, CTL and size of the VM of the tenant.

Based on the conflict tolerance level (the type of placement that is not optimal in this case), the CSP then provides a list of all existing tenants that are within the specified CTL with the prospective tenant. This is called a SoC_a set; a set containing all tenants that are in conflict with te_{i_A}. From this set, Algorithm 2 then chooses all addresses of tenants that are within the CTL range specified by *te*te_{i_A} and subsequently reduces the list to only those addresses with sufficient storage space.

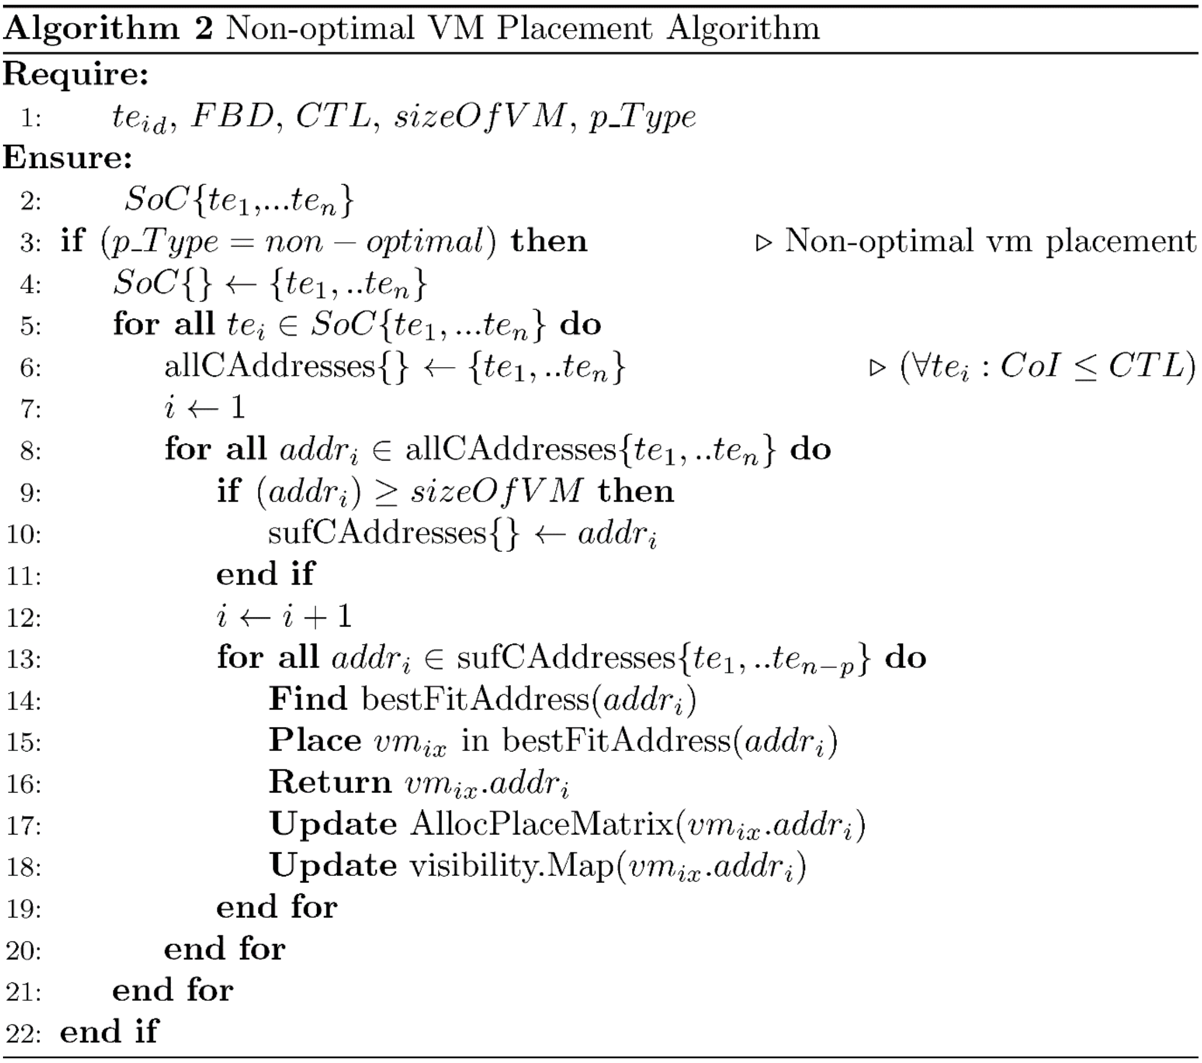

```
Algorithm 2 Non-optimal VM Placement Algorithm
Require:
    $te_{id}$, $FBD$, $CTL$, $sizeOfVM$, $p\_Type$
Ensure:
    $SoC\{te_1,...te_n\}$
if $(p\_Type = non - optimal)$ then                    ▷ Non-optimal vm placement
    $SoC\{\} \leftarrow \{te_1,..te_n\}$
    for all $te_i \in SoC\{te_1,...te_n\}$ do
        allCAddresses$\{\} \leftarrow \{te_1,..te_n\}$         ▷ $(\forall te_i : CoI \leq CTL)$
        $i \leftarrow 1$
        for all $addr_i \in$ allCAddresses$\{te_1,..te_n\}$ do
            if $(addr_i) \geq sizeOfVM$ then
                sufCAddresses$\{\} \leftarrow addr_i$
            end if
            $i \leftarrow i + 1$
            for all $addr_i \in$ sufCAddresses$\{te_1,..te_{n-p}\}$ do
                Find bestFitAddress$(addr_i)$
                Place $vm_{ix}$ in bestFitAddress$(addr_i)$
                Return $vm_{ix}.addr_i$
                Update AllocPlaceMatrix$(vm_{ix}.addr_i)$
                Update visibility.Map$(vm_{ix}.addr_i)$
            end for
        end for
    end for
end if
```

Algorithm 2: Placing a new VM in a non-optimal manner

Algorithm 2 goes on to choose the one address that best fits the new VM. It then places the new VM on the best potential host, sends the address of the new VM placement to the tenant, and updates the placement matrix and geographical map.

The next two algorithms are used by a CSP to migrate an existing tenant's VM from one host to another. They are also divided into two parts; one algorithm for migrating a VM to co-reside with non-conflicting tenants and the other one for migrating an existing tenant to co-reside with conflicting tenants. The scenarios raised by these two algorithms demonstrate the dynamic nature of today's organisations and the potential conflict involved. For example, a conflicting tenant today could be a non-conflicting tenant tomorrow. Furthermore, it could be that the sensitivity of the data in the VM has lost its sensitivity and has become public. For example, a VM that holds blueprints of new products are mostly confidential and highly sensitive prior to the products' release date. Once the products have been released, such blueprints somehow lose their sensitivity and their risk changes to a lower level. Hence, a tenant could opt to exercise some flexibility and reduce their CTL in order to co-host VMs that contain such data with conflicting tenants at a reduced price.

The price gets reduces mainly because the cost of VM placements is directly proportional to the CTL, i.e. directly conflicting tenants that require total isolation from their competitors pay a higher cost for their optimal placements. On the other hand, non-optimal placement comes at a lower cost. This is because non-optimal placement does not have to ensure total isolation from conflicting tenants and poses a certain risk of data leakage that the tenant is willing to accept as a trade-off for the lower cost.

Algorithm 3: Optimal Migration of an Existing VM

Algorithm 3 considers an existing tenant whose VM is to be migrated from one *PN* to another. The VM is to be migrated to a new *PN* where it can co-reside with non-conflicting tenants only. This algorithm is similar to Algorithm 1 for placing a new VM. However, Algorithm 3 can be used by a CSP to migrate an existing VM from one address to another. There are numerous reasons that could cause this, such as load balancing, changing tenant requirements and reducing operational costs.

Similar to Algorithm 1, Algorithm 3 first lists all non-conflicting tenants and then lists all addresses with sufficient storage space. Algorithm 3 next finds the best potential address that could fit in the best possible manner the VM that is being migrated. Algorithm 3 now places the existing VM on the best potential host, sends the address of the new VM placement to the tenant, and updates the placement matrix and cartographic map, while removing the old entry from where the VM was moved.

```
Algorithm 3 Optimal VM Migration Algorithm
Require:
    $te_{id}$, $FBD$, $vm_{ix}.addr_i$, $CTL_i$, $CTL_{i+1}$, $sizeOfVM$, $p\_Type$
Ensure:
    $SonC\{te_o,...te_z\}$
if $(p\_Type = optimal)$ then                                  ▷ Optimal vm migration
    $SonC\{\} \leftarrow \{te_o, ..te_z\}$
    for all $te_i \in SonC\{te_o, ...te_z\}$ do
        allNCAddresses$\{\} \leftarrow \{te_o, ..te_z\}$              ▷ $(\forall te_i : CoI \leq CTL_{i+1})$
        $i \leftarrow o$
        for all $addr_i \in$ allNCAddresses$\{te_o, ..te_z\}$ do
            if $(addr_i) \geq sizeOfVM$ then
                sufNCAddresses$\{\} \leftarrow addr_i$
            end if
            $i \leftarrow i + 1$
            for all $addr_i \in$ sufNCAddresses$\{te_o, ..te_{z-n}\}$ do
                Find bestFitAddress$(addr_i)$
                Place $vm_{ix}$ in bestFitAddress$(addr_i)$
                Return $vm_{ix}.addr_i$
                Update AllocPlaceMatrix$(vm_{ix}.addr_i)$
                Update visibility.Map$(vm_{ix}.addr_i)$
            end for
        end for
    end for
end if
```

Algorithm 3: Migrating an existing VM in an optimal manner

Algorithm 4: Non-Optimal Migration of an Existing VM

Algorithm 4 considers an existing tenant whose VM is to be migrated from one *PN* to another – with a requirement to co-reside with conflicting tenants of a specified *CTL_{i+1}* that is different from the initial *CTL_i* in the new *PN*. Each of these scenarios is handled differently by the CBVMP architecture. Algorithm 4 is similar to Algorithm 2 for placing a new VM in a non-optimal manner.

```
Algorithm 4 Non-optimal VM Migration Algorithm
Require:
    $te_{id}$, $FBD$, $CTL_i$, $CTL_{i+1}$, $sizeOfVM$, $p\_Type$
Ensure:
    $SoC\{te_1,...te_n\}$
if $(p\_Type = non - optimal)$ then                    ▷ Non-optimal vm migration
    $SoC\{\} \leftarrow \{te_1, ..te_n\}$
    for all $te_i \in SoC\{te_1, ...te_n\}$ do
        allCAddresses$\{\} \leftarrow \{te_1, ..te_n\}$        ▷ $(\forall te_i : CoI \leq CTL_{i+1})$
        $i \leftarrow 1$
        for all $addr_i \in$ allCAddresses$\{te_1, ..te_n\}$ do
            if $(addr_i) \geq sizeOfVM$ then
                sufCAddresses$\{\} \leftarrow addr_i$
            end if
            $i \leftarrow i + 1$
            for all $addr_i \in$ sufCAddresses$\{te_1, ..te_{n-p}\}$ do
                Find bestFitAddress$(addr_i)$
                Place $vm_{ix}$ in bestFitAddress$(addr_i)$
                Return $vm_{ix}.addr_i$
                Update AllocPlaceMatrix$(vm_{ix}.addr_i)$
                Update visibility.Map$(vm_{ix}.addr_i)$
            end for
        end for
    end for
end if
```

Algorithm 4: Migrating an existing VM in a non-optimal manner

However, Algorithm 4 can be used by a CSP to migrate an existing VM from one address to another. This algorithm starts by listing all conflicting tenants in the *SoC* set and then chooses only the addresses of tenants with sufficient storage space. The algorithm subsequently searches for the best potential host that could fit in the best possible way the VM that is to be migrated. Once placement is done, the CSP sends the address of the new location of the existing VM back to the tenant, and updates the allocation and placement matrix and geographic map. These algorithms are used for making initial VM placement and migration decisions. It must be noted though that placement and migration of VMs could also be initiated by the CSP for load-balancing purposes and other reasons. However, this must be transparent and visible to the implicated tenants but should still conform to the requirements as spelled out by the respective tenants. The next section discusses a proof-of-concept prototype as an implementation of the CBVMP model's algorithms.

7.5 PROOF-OF-CONCEPT PROTOTYPE IMPLEMENTATION

This section discusses the CBVMP model's proof-of-concept prototype implementation. The model is implemented in Java using three packages: CSP, Tenant and Framework (as depicted in Figure 7.3).

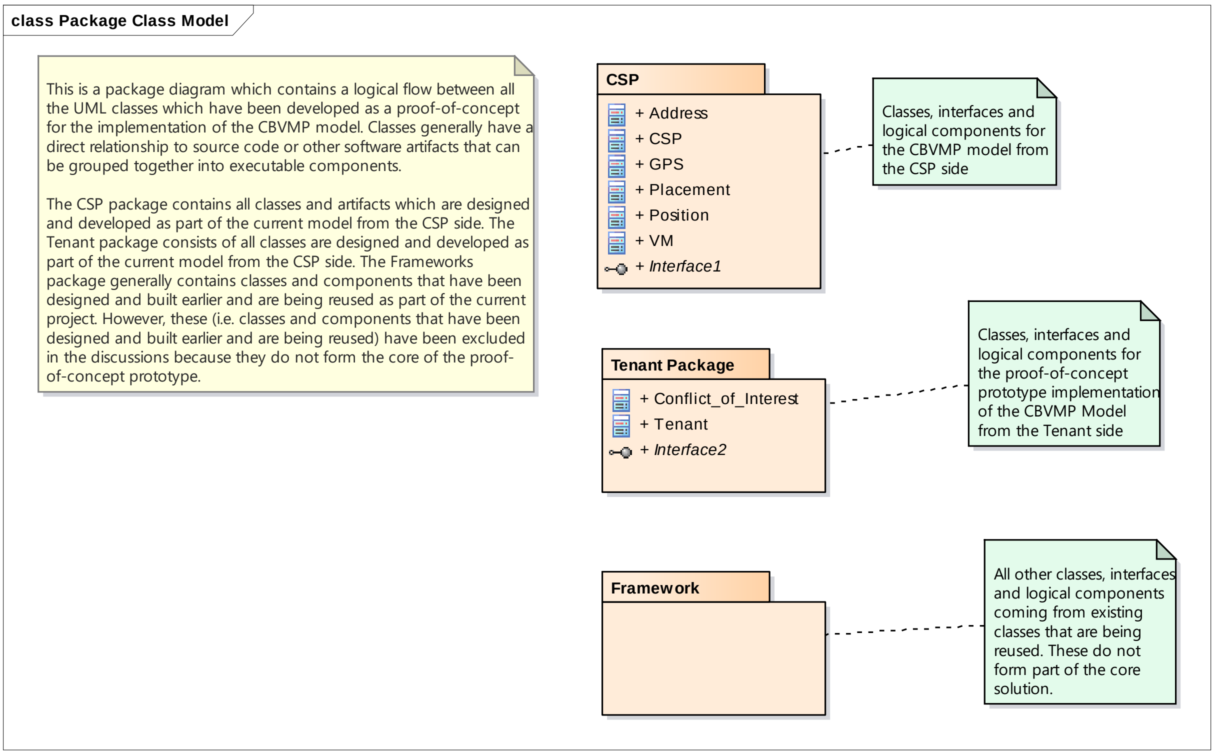
class Package Class Model
This is a package diagram which contains a logical flow between all the UML classes which have been developed as a proof-of-concept for the implementation of the CBVMP model. Classes generally have a direct relationship to source code or other software artifacts that can be grouped together into executable components.
The CSP package contains all classes and artifacts which are designed and developed as part of the current model from the CSP side. The Tenant package consists of all classes are designed and developed as part of the current model from the CSP side. The Frameworks package generally contains classes and components that have been designed and built earlier and are being reused as part of the current project. However, these (i.e. classes and components that have been designed and built earlier and are being reused) have been excluded in the discussions because they do not form the core of the proof-of-concept prototype.
CSP
+ Address
+ CSP
+ GPS
+ Placement
+ Position
+ VM
+ Interface1
Classes, interfaces and logical components for the CBVMP model from the CSP side
Tenant Package
+ Conflict_of_Interest
+ Tenant
+ Interface2
Classes, interfaces and logical components for the proof-of-concept prototype implementation of the CBVMP Model from the Tenant side
Framework
All other classes, interfaces and logical components coming from existing classes that are being reused. These do not form part of the core solution.

Figure 7.4:- UML Package Class Diagram for Implementing the CBVMP Model

A package is a container of a group of classes and interfaces that are related. Packages are used to organise classes and interfaces for better reusability and to resolve name conflicts. Therefore, this thesis also makes use of packages. The focus is mainly on the first two – CSP and Tenant packages. The CSP package consists of those classes and an interface that are related to the CPS's activities. The Tenant package consists of those classes and interface that are related to the tenants. The framework package consists of all other built-in classes such as lang, awt, swing, and javax (among others). Furthermore, the framework package includes other user-defined classes that have been designed and developed earlier to support the current project. Since these classes do not form the core of the solution but only provide support, it was decided not to include the framework package in the discussions below. However, it is worth mentioning that all other classes that support the implementation of the CBVMP model are found in the framework package. There are two main packages for implementing the CBVMP model, namely the *CSP* and *Tenant* packages. The *CSP* package consists of the following: *Address*, *CSP*, *GPS*, *Placement*, *Position* and *VM* classes and an *interface* to the CSP. The *Tenan*t package consists of the following: *Conflict_of_Interest* and *Tenant* classes and an *interface* to the CSP. Figure 7.5 illustrates the actual contents (i.e. attributes and operations) and relationships of each of the classes within the packages. Of note once again is that this figure does not include all other classes from the *framework* package. They fall outside the scope of this discussion and do not form a key part of the model. Figure 7.5 is followed by a brief discussion of the classes.

It must be noted though that the classes use an association-type of relationship. The *CSP* class provides a list of all tenants. The *CSP* class is associated with the *Conflict_of_Interest* class, which provides the different conflicts of interest derived from the functional business domains. The *Tenant* class depends on the *Conflict_of_Interest* class. This class sets the *CTL* for each tenant. The *Placement* class is the most important class as it does the placement of tenants' VMs and generates the cost thereof. This is the class that calculates the U and eventually lists potential allocations. The other classes such as *Position* are used to show the location of tenants' VMs in the geographical map. The *Address* class provides the listing of all placements. It works more like a Domain Name Service (DNS), as it provides the links to the actual VMs of each tenant in the form of an address or URL that is given as follows:

$$Loc_m \,.\, DC_l \,.\, Clu_k \,.\, PN_j \,.\, VM_i.$$

The *GPS* class is basically for providing GPS coordinates of the different DCs of the CSP. The *Placement* class returns an address of placement or migration that shows the global positioning of a tenant VM. A tenant can use this geographical map to trace and track their VM placement at any particular point in time.

The following screenshots illustrate the implementation of the CBVMP using the UML class model in Figure 7.5. Access to the solution is granted by a risk-based multi-factor authentication system. The risk-based multi-factor authentication system is covered in Chapter 8.

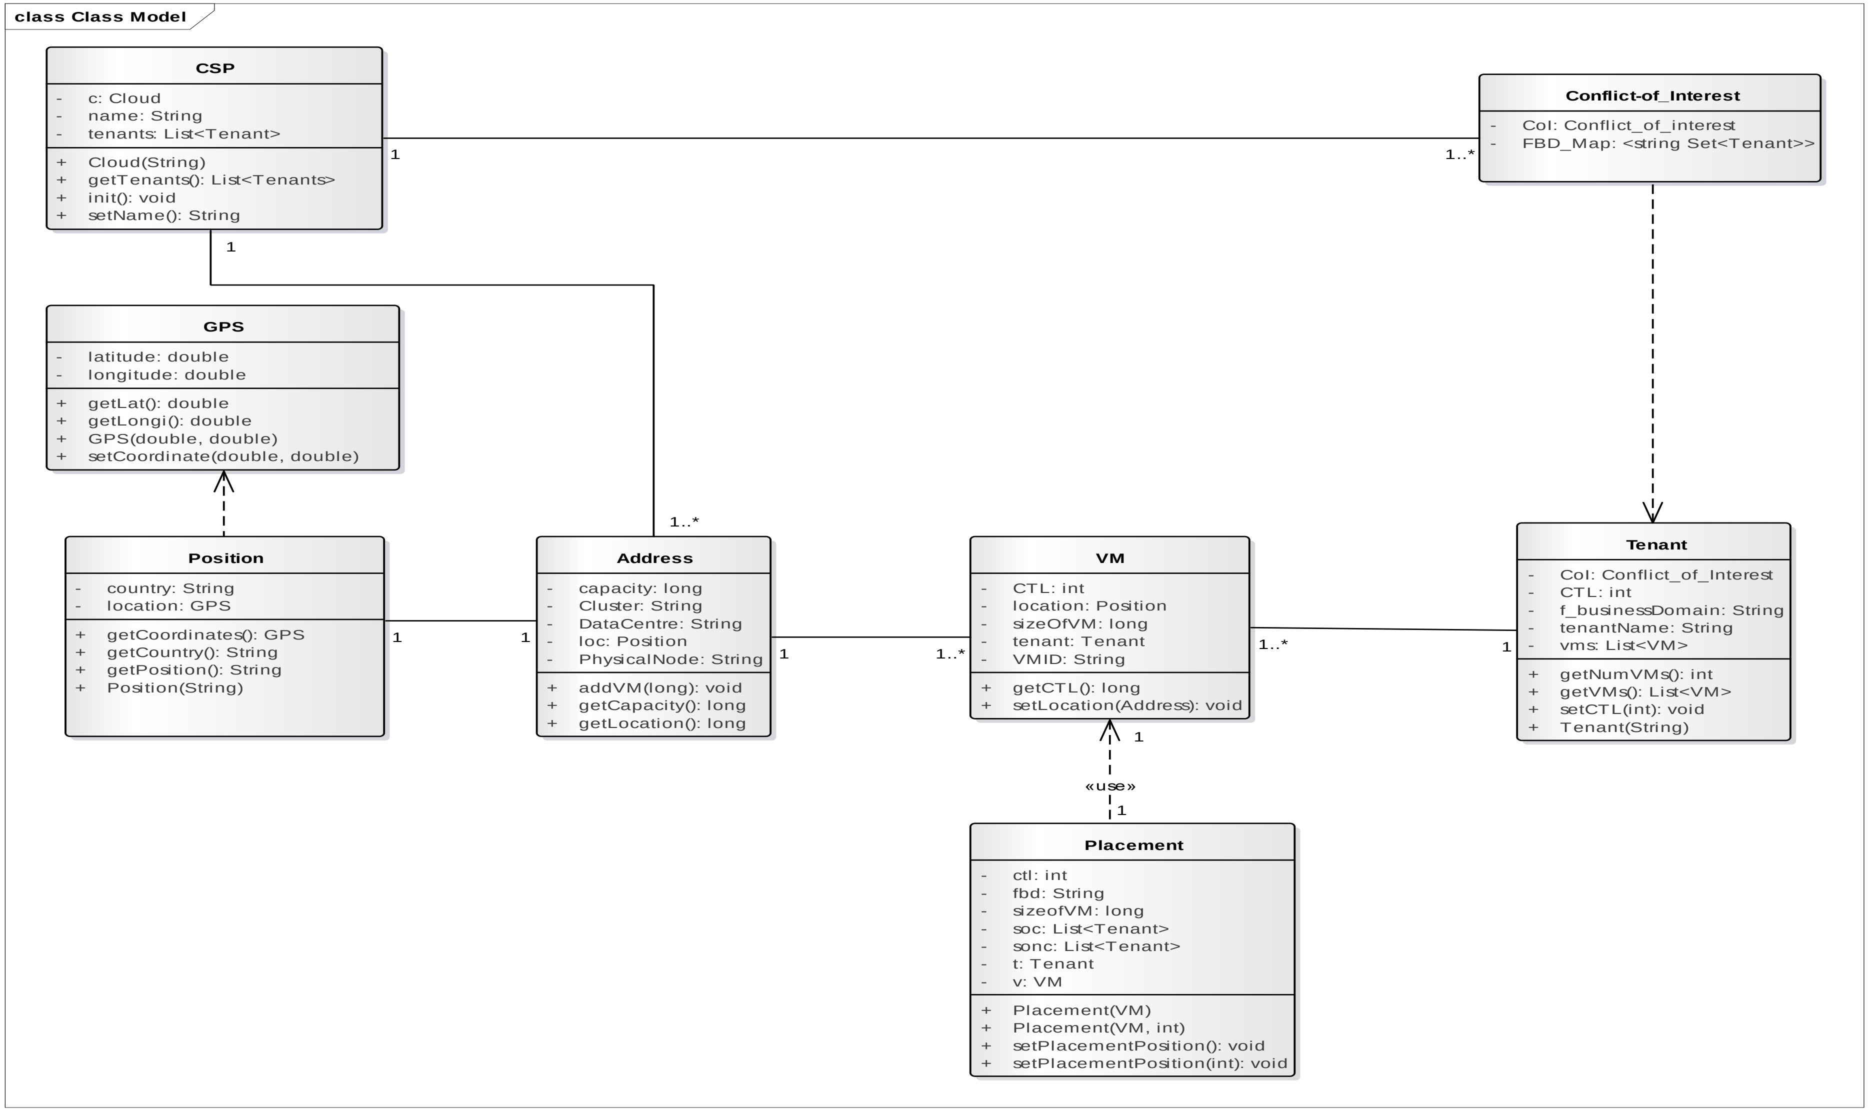

Figure 7.5:- UML Class Diagram for the Implementation of CBVMP Model

However, it suffices to say that once authenticated, tenants can only see their own VMs. Only administrators from the CSP have global access and a view of all the VMs hosted on the infrastructure.

Figure 7.6 depicts a tenant’s VM (the name of the tenant is Pick ‘n Pay) that is placed in South Africa. The tenant can view the properties of its own VMs as shown in Figure 7.7. The properties include the VM’s ID, location, current CTL, its CoI and its capacity. This capability of the solution gives tenants the crucial visibility of the location of their VMs at any given point. Over the years, cloud users have been complaining about the lack of visibility of their data once it is hosted in cloud services. Moreover, such visibility is essential for regulatory compliance issues. This is to ensure that tenants’ data is always compliant with legal regulations.

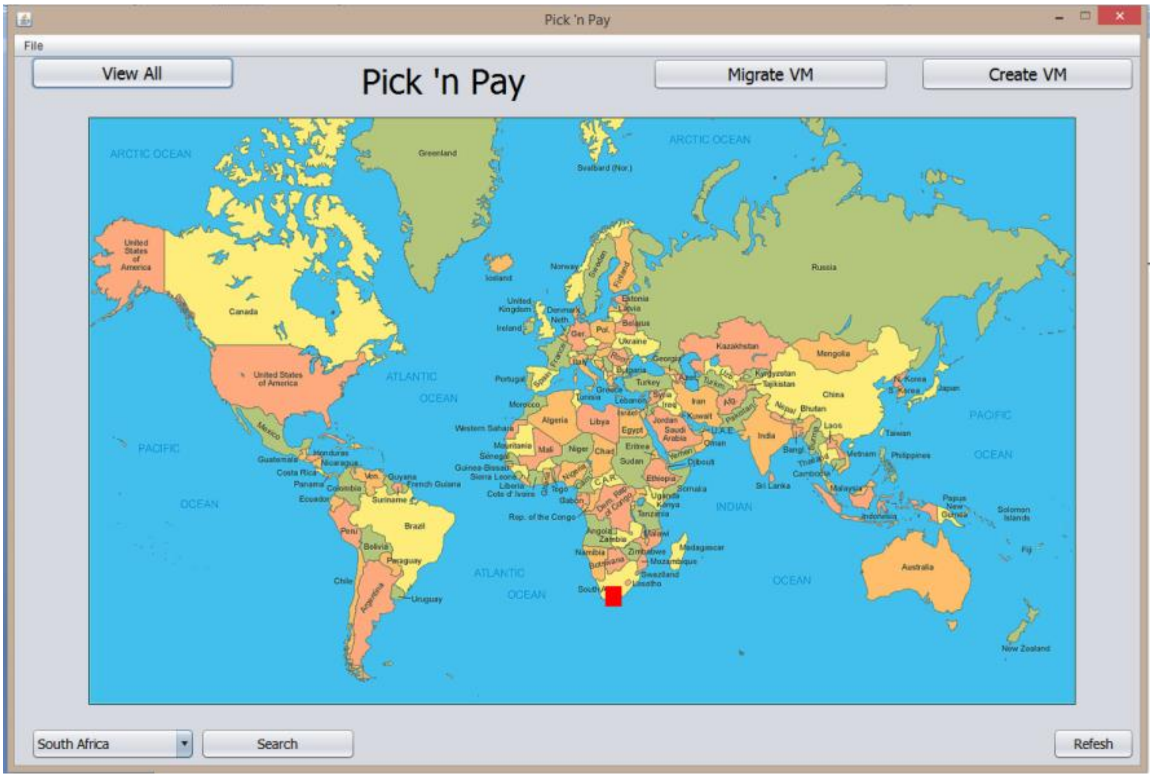

Figure 7.6:- A Screenshot Showing a Tenant’s VM

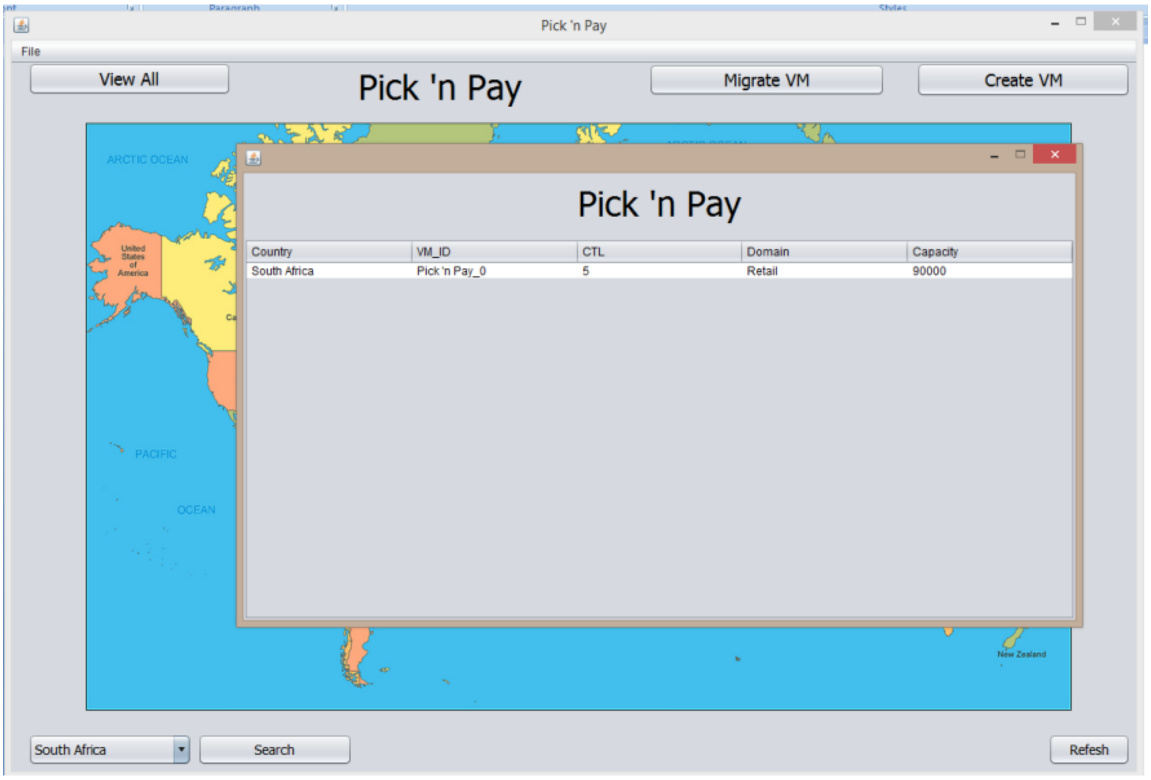

Country	VM_ID	CTL	Domain	Capacity
South Africa	Pick 'n Pay_0	5	Retail	90000

Figure 7.7:- A Screenshot Showing the Properties of VM

Figure 7.8 depicts all VM instances in the cloud infrastructure. This is an overview of the entire CSP's infrastructure across the world and is only visible to the administrators.

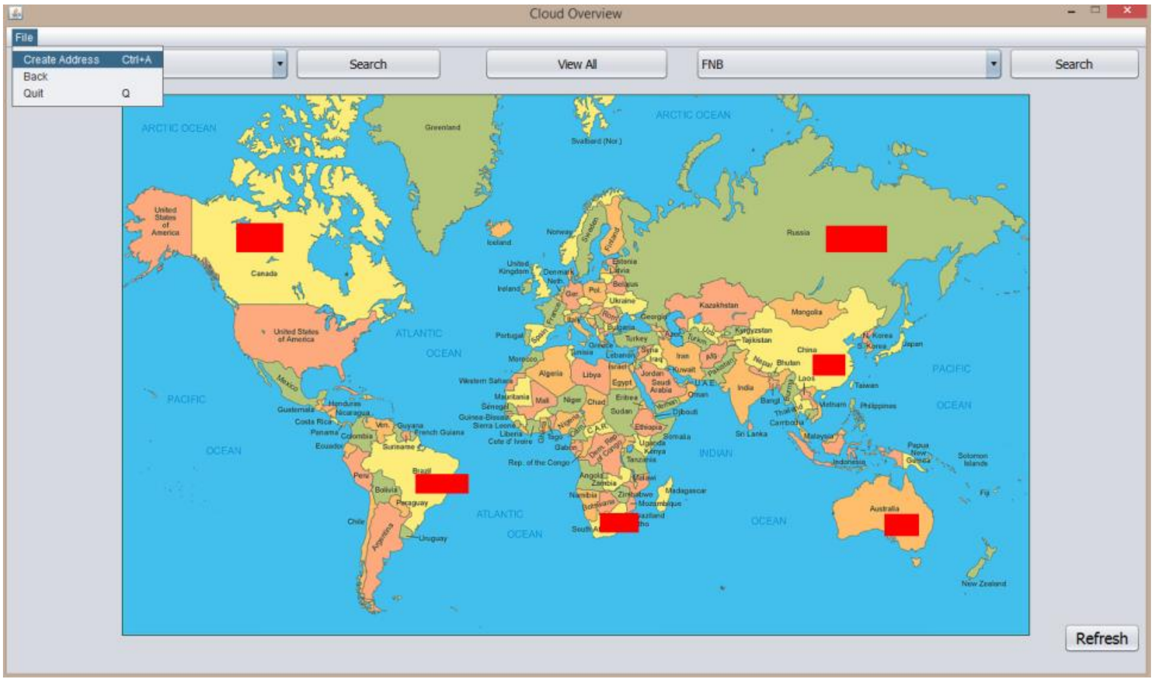

Figure 7.8:- An Administrator's Overall View of all Tenants' VMs

Figure 7.9 is a screenshot for setting up a new address on the map of the world. An administrator would use this to specify the size of the new *DC*, *Clus*, *PNs*, *Loc*.

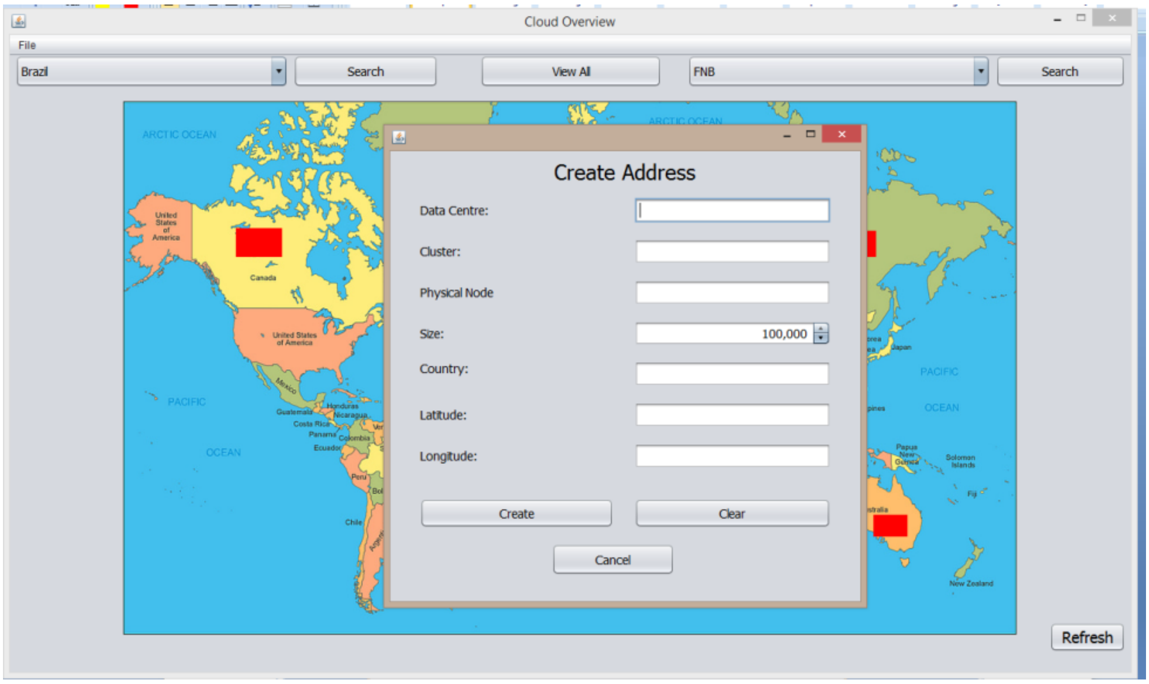

Cloud Overview

Brazil

FNB

Create Address

Data Centre:

Cluster:

Physical Node

Size: 100,000

Country:

Latitude:

Longitude:

Figure 7.9:- Creation of a New Address

On clicking ‘create’, a new *DC* is created based on the specification. This new address would be ready for prospective tenants to host their VMs. Furthermore, and at the request of a prospective tenant, an administrator can create a VM based on the tenant’s CTL and *size of VM* inputs. Figure 7.11 depicts a screenshot that is used by an administrator to create a new VM for an existing tenant. On this screenshot, the new VM ID is created automatically. The administrator can specify if the new VM must avoid conflicts. This checkbox activates the class that implements the algorithm on conflict-aware VM placement when checked. It then activates the CTL field determine the actual isolation distance of the new VM from other conflicting VMs based on the CTL. The administrator next specifies the size of the VM and the country where it would be hosted, based on the CTL.

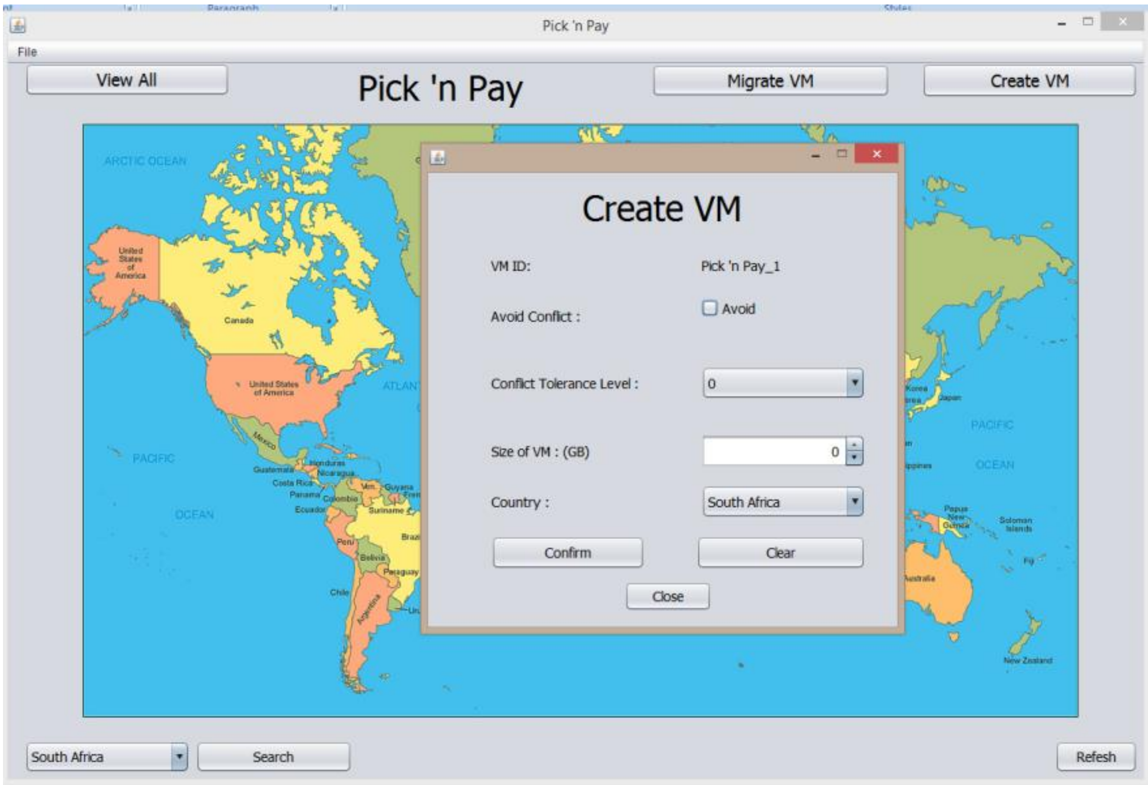

Pick 'n Pay

Create VM

VM ID: Pick 'n Pay_1

Avoid Conflict : □ Avoid

Conflict Tolerance Level : 0

Size of VM : (GB) 0

Country : South Africa

Figure 7.10:- Creation of a New VM

An administrator can also migrate an existing tenant's VM from one address to another. This could happen in case of load balancing from an administrator's perspective. However, the owner of the VM must approve all migration requests that emanate from load-balancing requests. Furthermore, it could also happen that the tenants themselves initiate a migration process. This would happen for example in cases where new legislation has been passed, which a tenant feels might compromise its data. Once again, a tenant's request for migration requires approval from the administrators to ensure that this feature is not abused to achieve co-residence with conflicting tenants. Figure 7.10 depicts a screenshot for migrating an existing VM from one address to another. Using this screen an administrator can also specify the new location and if it should avoid conflicts. Checking the 'avoid conflicts' checkbox activates the CTL dropdown, where a tenant can specify the exact CTL. On clicking 'migrate', the VM will be migrated to the new location on condition that there is sufficient space and no conflict.

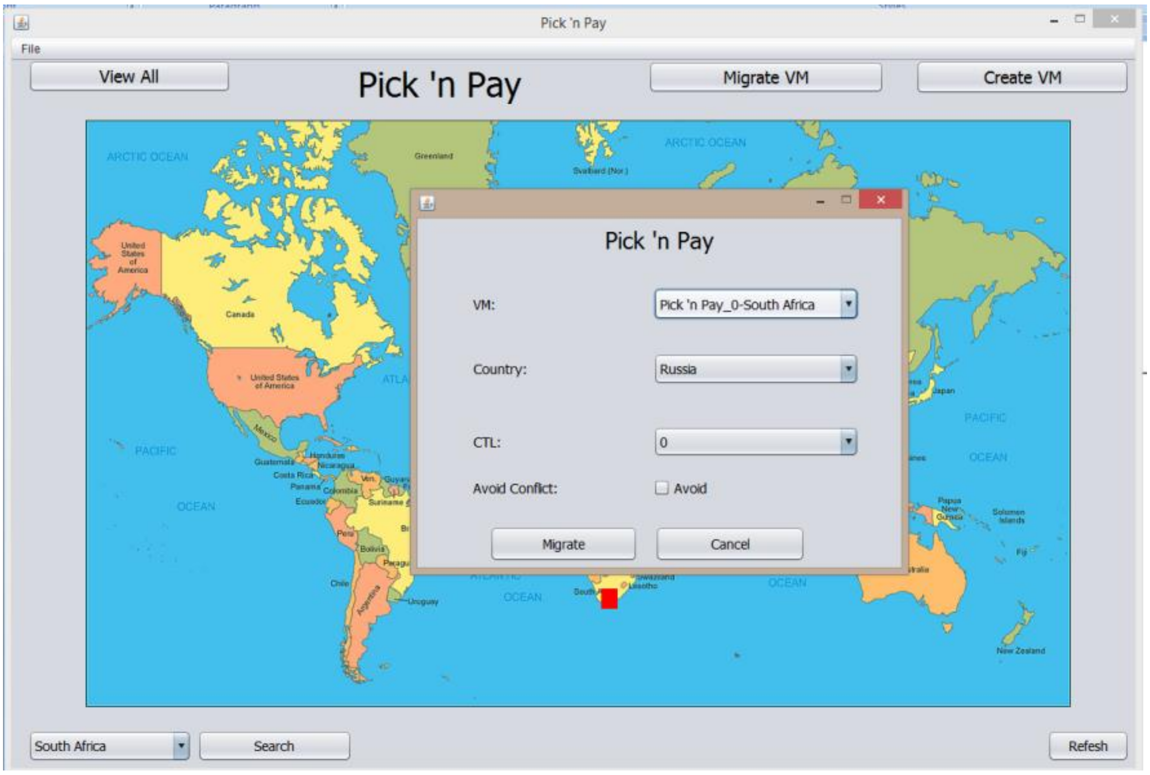

VM: Pick 'n Pay_0-South Africa

Country: Russia

CTL: 0

Avoid Conflict: □ Avoid

South Africa

Figure 7.11:- Migration of an Existing VM from One Location to Another

The above screenshots provide a glimpse of the actual processes of the proof-of-concept prototype implementation to sum up the discussion on the CBVMP model. The next section concludes Chapter 7 and points out the direction for the next chapter.

7.6 CONCLUSION

Finding the best VM placement algorithm is a challenging optimisation problem. This problem has a direct impact on costs, performance and energy consumption. Furthermore, it has a direct impact on the security of cloud-hosted data which may be leaked to unintended parties. This chapter reported on the formulation and presentation of a CBVMP model to prevent confidential data leakage threats posed by inter-VM attacks and to manage conflicts of interest between tenants in the cloud.

The model uses degrees of conflict, the constructs of a Sphere-of-Conflict and a Sphere-of-non-Conflict to provide for the physical separation of VMs belonging to conflicting tenants. Unlike most existing VM placement algorithms, the contribution made by this thesis also considers security, risk and cost involved in doing so, as well as the interdependency between these factors. The aim is to help tenants and CSPs make informed and well-calculated VM placement decisions that factor in their security profile – balanced against cost constraints and risk tolerance.

After the CSP has placed the tenants’ VMs in a manner that is conflict-aware and ensures adequate isolation, it is important that access to the VMs must be secured and constantly monitored. Hence, Chapter 8 delves into authentication, more specifically taking a risk-based multi-factor authentication approach to provide this security capability. Chapter 9 addresses the monitoring capability of the proposed solution.

CHAPTER 8 A RISK-BASED MULTIFACTOR AUTHENTICATION MODEL

8.1 INTRODUCTION

Chapter 7 zoomed into the high-level conceptual architecture presented in Chapter 6 to provide specific details on the VM placement component. It went on to provide theoretical and mathematical formulations of the CBVMP model abstracted as an optimisation problem. The model uses different degrees of CTL, as well as the constructs of SoC and SonC, S, R and C to provide VM physical separation of conflicting tenants. The CBVMP model addresses confidential data leakage posed specifically by inter-VM attacks. Chapter 7 ensures that tenants' VMs are placed in a conflict-aware manner, in the IaaS cloud. Chapter 8 now takes over to focus on how tenants' identities are verified when making attempts to access their VMs.

The authentication component of the architecture in Chapter 6 is based on the authentication requirements in Chapter 5. Chapter 8 presents a strong risk-based multifactor authentication model that authenticates users, based on their risk profile. Furthermore, this chapter demonstrates end-to-end security of user credentials and other authentication data at rest and in transit. The rest of Chapter 8 is structured as follows: Section 8.2 presents and discusses the proposed strong risk-based authentication model. Section 8.3 discusses how this research provides secure one-time passwords and out-of-band tokens (i.e. OTPs and OOB tokens). Section 8.4 briefly discusses the proof-of-concept prototype implementation of the proposed strong risk-based authentication model. Section 8.5 concludes this chapter and provides a high-level overview of Chapter 9.

8.2 A STRONG RISK-BASED MULTIFACTOR AUTHENTICATION MODEL

Figure 8.1 depicts a strong risk-based multifactor authentication model. This model attempts to address the rising issue of data leakage threats that emanate from unauthorised access to cloud resources due to inadequate authentication or compromised credentials. The model advances the current state of the art in authentication by using innovative ways to identify and verify users' identity when they access cloud resources. The model adopts and builds on the traditional username-and-password method of authentication. It adds device authentication to the user credentials. This ties users to specific devices. For example, a user might have more than one device to connect to the cloud services. Each device is associated with its own user behavioural pattern that is unique to the user. Furthermore, this chapter adopts a strong risk-based approach with the concept of a SecurityPhrase. The model also adds end-to-end security in handling credentials and authentication data. The innovation of the model is demonstrated in the manner it makes authentication decisions to grant or deny access based on a risk profile of a login attempt. It is also demonstrated in the way the proposed model transmits OTPs and OOB tokens across the vulnerable Internet using the old concept of steganography.

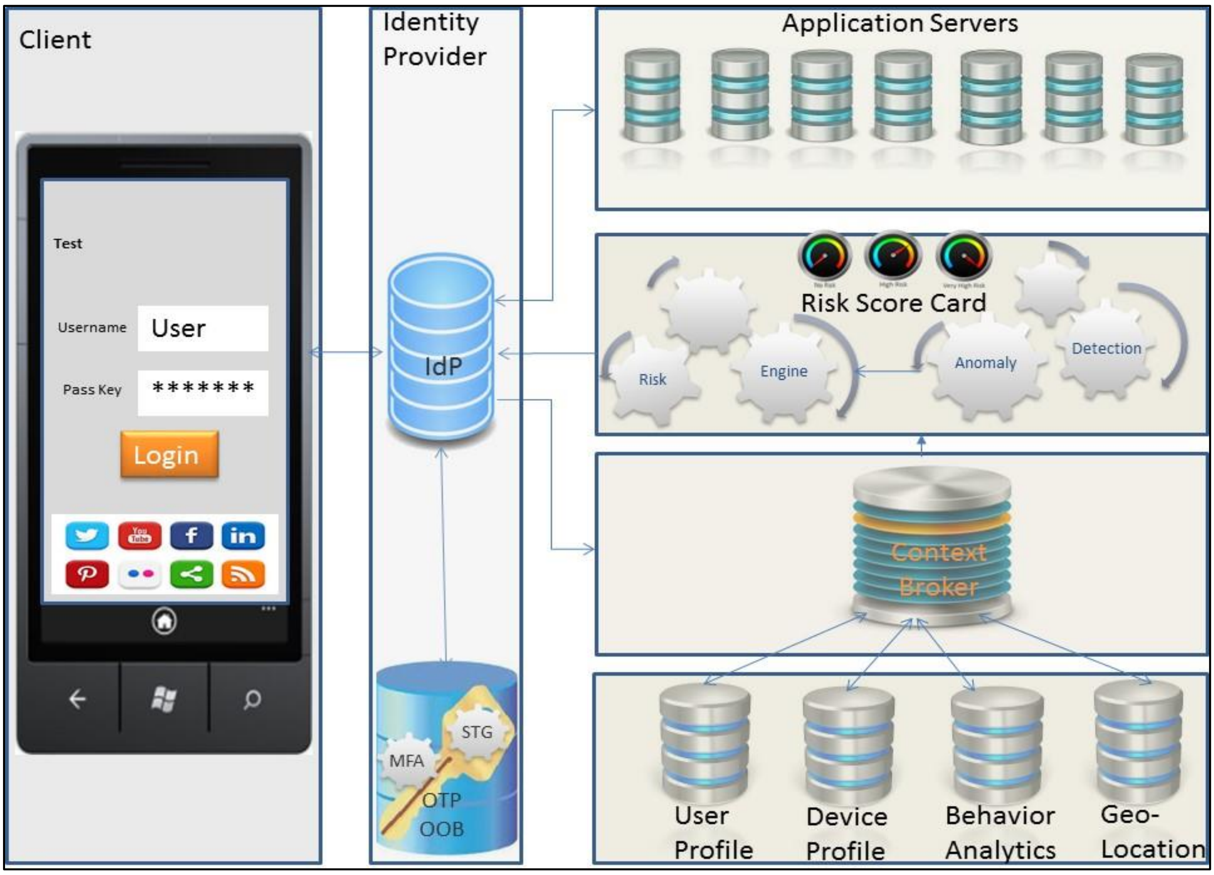

Figure 8.1:- A Strong Risk-based Multifactor Authentication Model (Dlamini et al., 2015)

Furthermore, the proposed solution makes use of geo-location data to add to the decision making of whether to grant or deny access. The other key benefit of this solution is that it uses keystroke behavioural data. Keystroke behavioural data can uniquely identify and verify users (Yeh et al., 2018; Krishnamoorthy et al., 2018). However, it requires sufficient training of the data to achieve the best results.

The model has six distinct parts: client; identity provider; risk score card; context broker; application servers; data stores. Data stores comprise of user profile, device profile, behaviour analytics and geo-locations data store instances. Each of these parts addresses one or two of the authentication requirements raised in Chapter 5. For example, the identity provider and risk score card parts assess and evaluate the potential risk of a login attempt. The risk score card then passes control to the multifactor authentication (MFA) within the Identity Provider. The MFA chooses the right authentication combination based on the risk score. Below is a discussion of each of these parts, how they each achieve their goals and their overall contribution to the entire solution.

8.2.1 Client

A user first registers on the system choosing a username and a strong password (i.e. a combination of characters made of alpha-numeric and special characters – as enforced in the password creation). These are checked for validity, character by character in real time and natively on the client device as they are being entered. A subsequent field is only enabled when the current field has been validated. For example, after correctly entering the username, the password field is activated. This is called a two-step authentication. A two-step authentication is beneficial to the model because it makes it difficult for attackers to carry out brute-force attacks. For example, it prevents attackers from even trying a password if they fail to enter a correct username. However, a two-step authentication could be used as a weapon by clever attackers to verify usernames before they start with a brute-force or dictionary attack on the password. After verifying the usernames, the only task would be to brute-force the passwords. On a positive note though, manual entry of credentials is one way to address brute-force and dictionary attacks. Therefore, this solution only considers manual entry of user credentials without using auto-complete.

Over and above user credentials, the model requires users to select a security phrase (SecurityPhrase) The idea of a SecurityPhrase is derived from the traditional concept of security questions that the information security community normally uses to authenticate users on top of the username and password combination. With the advent of social media, most of the answers to security questions are now freely available and can be accessed by attackers. Against this background, the research in hand innovated around the concept of security question to propose the SecurityPhrase concept. The SecurityPhrase is a predefined eight-character phrase that is randomly generated by the model during user registration. A valid SecurityPhrase follows the principles of a strong password with a mix of characters. For example, #Dg$H@6^ is a valid SecurityPhrase (derived from a user entering the following phrase on registration - hash-tag (#) David goes ($) Home at (@) 6:00 sharp (^)) (Dlamini et al., 2015). The SecurityPhrase is kept encrypted with the user credentials in the User Profile. On request, the user would at times be prompted to input missing characters on the SecurityPhrase. The missing characters are randomly generated, based on the risk profile of a user. The specific details of how this is done are explained later in Section 8.2.7.

8.2.2 Keystroke Dynamics

On the client device, a user manually enters their access credentials from a keyboard without using any autonomous complete or copy-and-paste functionality. This helps to extract keystroke behavioural features for each login attempt. The features include session times, typing speed, pressure on the keys, key combinations and keystroke latency. These are captured in the background without the user being aware and are stored in the behaviour analytics. This information is used as part of a user's profile. The model extracts typing speed, timestamps of each session, left and right shift, caps lock, backspace, delete, keystroke latency and keystroke pressure. These features are collectively used to model and profile a user's behavioural patterns.

Figure 8.2 depicts a system that captures and uses behavioural keystroke dynamics (Dlamini et al., 2017). The system is first trained as the user repeatedly makes login attempts and the training data is stored in a repository. The training data includes all the captured keystroke data (i.e. feature extraction and contextual data), which is then stored in a behavioural analytics' repository. The behavioural analytics' repository is further discussed later in this chapter. The training data, along with subsequent successful logins, constitutes the historical data. The solution continuously learns new behaviour patterns and updates itself for each and every successful login attempt. As the system learns with more successful login attempts, its accuracy in classifying and detecting future login attempts also improves. The historical data is drawn from the repository during a login attempt and goes into a clustering process. The clustering process is also discussed later in this chapter.

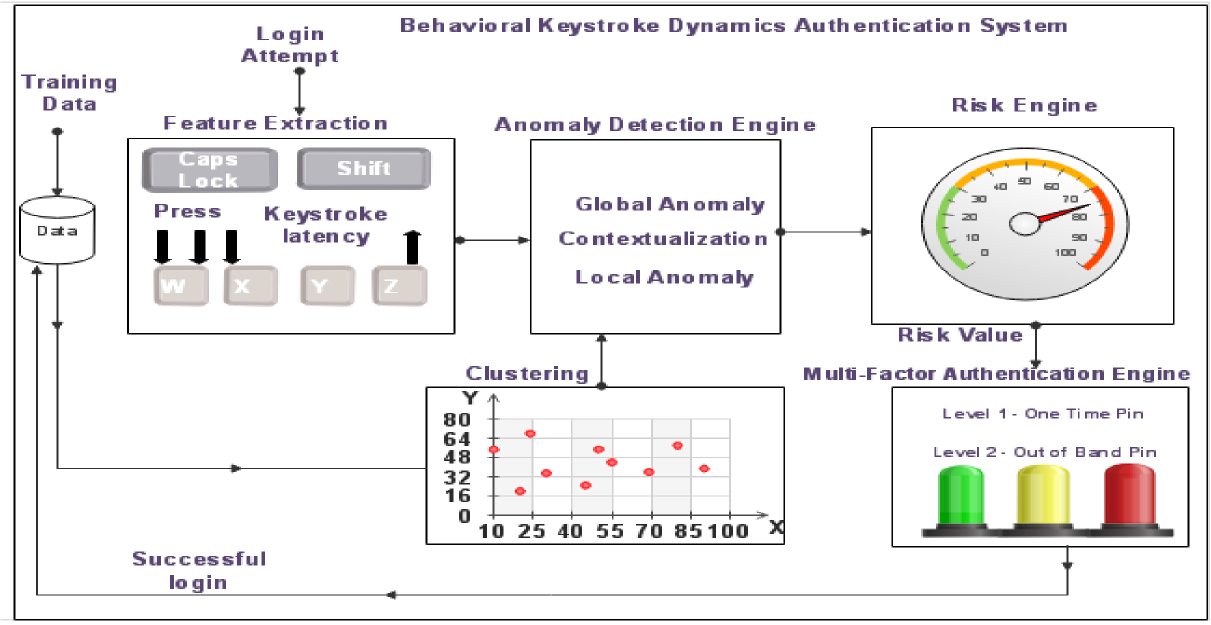

Figure 8.2:- Capturing and Using Behavioural Keystroke Dynamics (Dlamini et al., 2017)

Furthermore, the model also covertly captures device identifier, the access point's GPS (Global Positioning System) coordinates, source IP (Internet Protocol) address, device Operating System, serial numbers, SIM (Subscriber Identity Module) card number, SIM card serial number, IMEI (International Mobile Equipment Identity) number, and a unique host name, depending on the client device. For example, when a user uses a SIM-enabled mobile device, the solution collects GPS coordinates which might not necessarily be possible with desktop devices. The capturing of device information and access point locations also happens seamlessly without interrupting the user. Device-specific information is then sent to the device profile's repository. Each device will have its unique profile and the solution allows user to use different devices. Geo-location data like GPS coordinates and IP addresses are sent to the geo-location repository. This data is collected from the time the user first registers. It is stored as historical data to be used later to evaluate a user's risk level. The collected data is then passed to the Identity Provider.

8.2.3 Identity Provider

The Identity Provider acts as a proxy that takes user login and covertly acquired data from the client device and relays it to the context broker. The Identity Provider also receives the risk score card from the risk engine and makes the decision to relay it to the MFA engine, should there be a need for an OTP (Thomas and Goudar, 2018) or OOB (Sy et al., 2019) token to further authenticate a user's identity. The Identity Provider also receives the response from the MFA and routes it to the client. Once the client has completed entering either an OTP or OOB token, the request is relayed back to the MFA for verification. Should the risk score card require a SecurityPhrase (Dlamini et al. 2014), the Identity Provider relays the request through the context broker. The context broker pulls the SecurityPhrase of the user from the user profile data store and relays it straight to the risk engine, which simply passes it to the Identity Provider. The Identity Provider then gets the response from the client and confirms whether it matches the SecurityPhrase on the user profile data store. Finally, the Identity Provider relays the decision to grant or deny access back to the user. If the decision is to grant access, control is passed to the requested and authorised application servers.

8.2.4 Context Broker

The context broker takes the incoming data from the Identity Provider and performs minor processing before it pushes the data to the relevant data stores or repositories. Processing the data involves applying encryption and indexing. For example, during the registration process the context broker extracts, encrypts and indexes the data using hashes, the user credentials, device information, behaviour keystroke data and geo-location data. The encrypted and hashed datasets are then pushed to the user profile, device profile, behaviour analytics and geo-location repositories respectively. Each of these datasets has a unique identifier that relates to the specific user and device from which it was captured. The encryption part helps to secure the authentication data at rest. Indexing the encrypted datasets by using hashes helps to improve the speed of record retrieval. The context broker uses a novel cryptography system that is based on chaotic noise for encryption. The encryption system that is used in this study is explained in detail later in Section 8.3.2.

Moreover, the context broker pulls all data from the different repositories, and decrypts and relays the plain-text data to the risk score card module for detecting anomalies and determining the risk score of a new login attempt.

The next subsection discusses the process of data storage.

8.2.5 Repository

All captured data (i.e. feature extraction and contextual data) is now stored in a data store as historical data. Historical data consisting of successful login attempts is kept in a data store. The model autonomously learns new behavioural patterns and performs self-updates for each and every successful login attempt. The historical data is drawn from the database during a login attempt and taken into the risk score card to determine anomalies and their risk score.

This model distributes encrypted shares of authentication datasets in multiple repositories. The concept of shares is derived from Shamir Secret Sharing Scheme (Tieng and Nocon, 2016). Each share is assigned a unique header using a hashing algorithm. The first element of each share is the user profile identity for identification and retrieval. For example, all shares belonging to user profile *xyz* will start with *xyz-* as their unique hash identifier. The distribution of the encrypted shares is meant to increase the resistance of the model to attacks and to mitigate the increasing theft of user credentials. A compromise on one repository would not help an attacker that much. This is because other segments of the data would be held in other repositories. Hence, an attacker would have to compromise all repositories to be able to combine the different shares into a full set before trying to decrypt it using the context broker. A dataset share from one database is rendered useless and cannot be decrypted without the rest of the shares from other repositories. Furthermore, this approach stands to enhance data availability. For example, if it happens that one repository goes down or becomes corrupted, partial data could still be recovered from the rest of the databases. The context broker retrieves data from each of the repositories and then decrypts it. The plaintext data is then sent to the risk score card.

8.2.6 Risk Score Card

The risk score card process is divided into two parts: anomaly detection and risk engine. Each of these parts is responsible for a specific task in evaluating the risk score of a login attempt. The following two subsections discuss each of these processes in detail.

8.2.7 Anomaly Detection

The anomaly detection process receives plain-text data from the context broker along with the historical data from each of the repositories. It then classifies and categorises the historical data according to clusters. This means that each successful login in the historical data forms a data point within a cluster. Each of these data points is classified and grouped in clusters based on feature similarity. For example in the proposed model, this would be related to the features of each login instance that forms a data point. This means that data points (in this case login points) within a cluster must be as similar as possible. The process of clustering the login data points helps to trace different patterns of user behaviour and facilitates the anomaly detection process. The model adopts and uses a *k-means* clustering algorithm (Hartigan, 1979). The choice of using a *k-means* clustering algorithm is based on the algorithm's wide usage and simplicity. The *k-means* clustering algorithm is one of the simplest and most widely used iterative clustering algorithms that has been adapted to many problem domains. It divides *x* number of data points in *y* dimensions into *k* clusters. This is done in such a way that a within-cluster sum of squares is minimised. The author acknowledges that this might not necessarily be the best clustering option. However, it suffices to prove the concept of this thesis. Further studies may be conducted to search for the best clustering algorithm, but this is beyond the scope of this study. The *k-means* clustering algorithm (Hartigan, 1979) is as follows:

1. *Place k points into the space represented by the objects that are being clustered. (These points represent initial group centroids – centre points in each cluster.)*
2. *Assign each object to the group that has the closest centroid.*
3. *When all objects have been assigned, recalculate the positions of the k centroids.*
4. *Repeat Steps 2 and 3 until the centroids can no longer move. This produces a separation of the objects into groups from which the metric to be minimised can be calculated.*
5. *Stop when none of the cluster assignments brings about any significant change.*

The main challenge of the *k-means* clustering algorithm is to determine the optimal value of *k* (i.e. the number of clusters) with their centroids. The question about the optimal value of *k* in the *k-means* algorithm is one that still requires answers. Hence, some researchers resort to a random selection of the initial values and iterate with multiple runs until they can choose the one with the best result (i.e. minimum sum of squares within each cluster). A number of heuristics have been proposed to help determine the optimal value of *k*. Given that there is still no method that performs better than the other, this thesis adopts the elbow method for determining the value of *k* (Bholowalia and Kumar, 2014). This choice is based on the belief that the elbow method is one of the most tried and tested technique for determining the value of *k*.

The elbow method as stated in Bholowalia and Kumar (2014) starts with a minimum *k* and iterates until there is a minimum distance between each centroid and its associated data points within each cluster. Using the elbow method, one just needs to note the point where an increase in the number of *k* does not necessarily result in a significant change from previous points. This is the point where marginal changes are noted in the points on the graph, and it is called the elbow criterion. This is the point that is considered the best value of *k*. The final result of the *k-means* based on the elbow method is a number of clusters of the historical data points. Each new login attempt is assessed and classified to determine the closest cluster of points that are similar. The similarity of each new login attempt is measured using a Euclidean distance. This is done by choosing the best cluster that represents the login attempt. The context of the login attempt is also considered to ensure that the model will correctly detect outliers. For example, a login attempt that is contextualised to cluster *a*, which lies close to the centroid of cluster *b*, would still be classified as an outlier.

There are two thresholds T_1 and T_2 for determining the distance between the observed login to the centroid of the chosen cluster. This is the point where the model uses the two thresholds T_1 and T_2 to detect outliers in order to determine if the observed login is a local anomaly or global anomaly. A local anomaly is detected when the distance between the observed login point and the centroid is strictly greater than first threshold T_1 but less than the second threshold T_2. For example, a user with the correct credentials and correct SecurityPhrase, but with a non-matching geo-location would be detected as a local anomaly. A global anomaly is detected when the Euclidean distance is equal to or greater than the second threshold T_2. A global anomaly is associated with login attempts that are considered extreme anomalies. For example, a global anomaly would occur when a user who habitually uses '*caps lock*' for capital letters starts using a combination of '*shift*' and a character. Once the anomaly detection process is finished, the model switches control to the risk engine.

8.2.8 Risk Engine

The risk engine takes the results of the anomaly detection process as input to evaluate and assess the risk of a login attempt under observation. For example, a global anomaly carries a higher risk than a local anomaly. Some of the observed login attempts might reflect no anomalies at all. This would be flagged no risk. The risk levels or categories are as follows:

- No risk
- Low risk
- Moderate risk
- High risk
- Very high risk

Some login attempts might be flagged 'minimal risk' and some 'moderate risk'. Login attempts that raise a global anomaly carry a high to very high risk.

Should it happen that a login attempt under observation falls within its contextual cluster in the anomaly detection process; the risk engine will flag it as no or minimal risk. This switches control to the Identity Provider to grant access to the user for the requested resources. Local anomalies that are flagged for moderate and high-risk levels trigger the model to switch control to the MFA through the Identity Provider. The next subsection discusses what goes on in the MFA process.

8.2.9 Multifactor Authentication (MFA)

The MFA engine assesses the risk level of a user's login attempt that is received from the risk engine process. This engine has different levels of authentication measures, each of which corresponds to a specific risk level of a login attempt, i.e. no risk, low risk, moderate risk, high risk and very high risk. The MFA engine applies the corresponding level of authentication measure to the corresponding risk level. This acts as a barrier of defence. Furthermore, it allows for the model to learn from behavioural patterns (Dostalek, 2019). The model uses the following MFA mechanisms for the different risk-levels:

1. No or low risk: credentials and SecurityPhrase with minimal missing characters (25%)
2. Moderate risk: variable SecurityPhrase with an average number of missing characters (50%)
3. High risk: One Time Password (OTP) (Thomas and Goudar, 2018) and a SecurityPhrase with a minimal number of missing characters (75%)
4. Very high risk: Out-of-Band token (OOB) (Wu et al., 2018; Sy et al., 2019) and

SecurityPhrase with all missing characters (100%)

No or low risk prompts the Identity Provider to verify credentials and it requests the user to complete the SecurityPhrase by entering 25% characters. 75% of the SecurityPhrase would already be displayed with some gaps for the 25% missing characters. A moderate risk level requires verification of credentials and requests a user to key in at least 50% of their SecurityPhrase. A high-risk level prompts the MFA engine to request a user to enter at least 75% characters of the SecurityPhrase. Furthermore, it would issue an OTP to a registered device or email as an extra layer of defence. Once prompted for an OTP, the model sends an OTP to a mobile device or an email address. A very high risk prompts the MFA engine to issue the same requirements as high risk, and over and above that, to issue an OOB token. An OOB token is accompanied by a 100% SecurityPhrase. An OOB authentication is a relatively new authentication technique based on the OTP concept. Gemalto (2018) defines an OOB authentication as a strong authentication that utilises a communication channel beyond the one on which a user is being authenticated. OOB authentication redirects authentication to a remote device. A user would for example be re-directed to authenticate using an OOB token on a different device in their device profiles. Once a user is authenticated, control is passed back to the initial device where the user was authenticating and the process continues from thereon. One advantage of using OOB tokens is that they can help protect against man-in-the-browser attacks and other advanced malware that targets OTPs as they are transmitted over vulnerable networks (Pham, 2014). Hence, the next subsection discusses the approach of this thesis in dealing with advanced malware that intercepts OTPs and OOB tokens.

8.3 SECURE OTP AND OOB TOKENS

User credentials in general are sometimes transmitted in plaintext over unsecured networks where they could be eavesdropped and intercepted by malicious parties. Some researchers have illustrated how attackers use advanced malware to successfully intercept OTPs for online banking websites (Litke and Stewart, 2014). Hence, the solution in this thesis provides a secure way of transmitting OTPs, OOB tokens and SecurityPhrases.

8.3.1 Generation

Several applications that claim to produce strings of truly random numbers for OTPs or OOB tokens actually make use computer-generated random numbers that are based on deterministic Pseudo Random Number Generators (PRNGs). Random numbers based on PRNGs repeat at some point. This means that by using PRNGs, one is likely to get the same sequence of numbers at some point later on. For some of the best PRNGs, the repeating sequence might take a long time. However, this only delays the process without addressing the vulnerability. PRNGs may be suitable for modelling and simulation. However, they are not suitable for applications that require true randomness and non-deterministic properties. Hence, this thesis generates OTP and OOB tokens based on a non-deterministic True Random Number Generator (TRNG) by following a technique that is proposed in Blackledge et al. (2013), Mosola et al. (2016) and Mosola et al. (2017).

TRNGs exhibit true randomness and non-determinism properties from physical atmospheric noise which is sourced from RANDOM.ORG (Random.org, 2017). RANDOM.ORG makes use of small variations in the amplitude of atmospheric noise. The author used the atmospheric noise from Random.org and passed it to a cloud-based Eureqa system from Cornell Creative Machine Lab (Dlamini et al., 2016). The choice to use this cloud-based system was informed by its ease of use, free availability and short turn-around time from the input noise to the resultant non-linear fitness functions. The Eureqa cloud system takes the noise input and generates a number of non-linear fitness functions that resemble the input noise. For example, the noise distribution in Figure 8.3 resulted in the fitness function (i.e. Equation 8.1) (Dlamini et al., 2016).

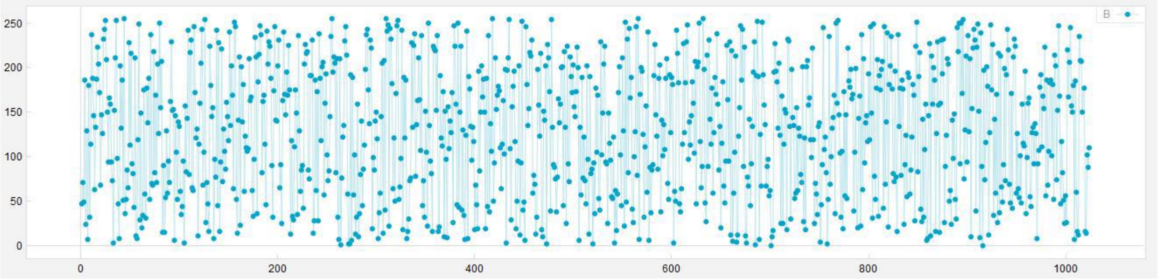

Figure 8.3:- Random Noise Distribution from RANDOM.ORG

$$\mathbf{f(x)} = \mathbf{a_o} * \cos(bx) + \mathbf{a_1} * \cos(cx^2) - \mathbf{a_2} * sin(d + ex^2 + \sin(fx) - \mathbf{a_3} * \boldsymbol{x}) \qquad (8.1)$$

where, $\mathbf{a_0}$, $\mathbf{a_1}$, $\mathbf{a_2}$, $\mathbf{a_3}$ are coefficient weights that come as a thirteen digit floating point number from the Eureqa system (Nutonian, 2018). *b*, *c*, *d*, *e* and *f* are coefficients of *x*. A decision was made to consider the original coefficient weights as they were from the Eureqa system (Nutonian, 2018) (see Equation 8.1). These were implemented in their original form in order to avoid losing any of the randomness and chaotic properties of the resultant fitness function derived from the input noise. The author subsequently took and implemented the best-fitted function to generate truly random OTPs. The same process applied for the generation of OOB tokens and encryption keys for the cryptography system that is used in this thesis. Even though this finding is not the main contribution of this thesis, it does set the current research apart from other (existing) studies. For more detail on this part, the reader is directed to Dlamini et al. (2016). The next section proceeds from this section to discuss how the encryption of the OTPs and OOB tokens is handled.

8.3.2 Encryption

The proposed solution uses a symmetric encryption scheme. This means the same key is used to encrypt and decrypt data. (A discussion on other encryption schemes falls outside the scope of this research.) The encryption process begins by converting the input OTP or OOB token sequence of four random numbers into a binary stream. The encryption algorithm then generates random floating-point integers that are normalised between zero and one. The end results of the normalised floating-point integers are a set of 64-bit binary stream encryption keys. One encryption key is randomly selected from each set and the process is repeated until the key size is equivalent to the plain-text bit stream. Each encryption key bit stream is generated from a new array of floating-point integers to further increase the randomness properties of the key.

For example, in order to encrypt a plain text of 240-bit stream, the process of key generation must be repeated at least four times with each 64-bit stream key generated from a totally new set of keys. The four keys of 64 bit each are then combined to be equal to the 240-bit stream of plain text. The extra bits are discarded, and the rest are combined using an exclusive or (i.e. XOR) operation with the plain text to form the ciphertext. The author acknowledges that an XOR operation has its own vulnerabilities when it comes to decryption. Hence, to make up for the XOR vulnerability, this study adds noise to the ciphertext by blending the key into it (Dlamini et al., 2016; Mosola et al., 2017). This is just to confuse any reverse engineering of the XOR operation in an attempt to uncover the key from the ciphertext. An ideal approach to solve this problem would be to use blocks and rounds as 3DES and AES cryptography systems. However, the issue of rounds and blocks is left as future work and is not part of this thesis.

The ciphertext bit stream and the key bit stream are blended in such a manner that the first four bits come from the first four bits of the ciphertext; the next four bits are the first four bits of the key stream, followed by the second four bits of the cipher, and so and so on until the cycles are finished. This way of adding noise to the ciphertext also helps to avoid the cumbersome key management challenges faced by most crypto systems. Security experts would shout ‘security by obscurity’ here, however, the approach does help to prevent attacks that target encryption key stores. This approach does not need any key store, which can be acknowledged as a vulnerability that, if discovered, could result in a catastrophe. Thus, research is ongoing on how best we can tackle the prevalent encryption key management in crypto systems (Mosola et al., 2017).

The idea of repeating the process for each key as stated above is meant to further increase the randomness of the generated keys and to confuse brute-force attacks. This is one way in which this thesis makes a plausible attempt to strengthen the encryption of OTPs and OOB tokens. It must be mentioned, though, that even the user credentials, SecurityPhrase, behavioural analytics data, device data and geo-location data that are also used for authentication are encrypted in the same manner.

The decryption process reverses all the steps of the encryption. It takes the blended ciphertext and splits it into sets of four bits and takes all odd sets and concatenates them to form the ciphertext. It takes all even number sets and concatenates them to form the key stream. Then, a reverse XOR operation of the key and issued the ciphertext is performed. The result is a binary stream of plain text. The plain-text binary stream is converted back using an ASCI table to get the actual plain-text OTP or OOB token. After encrypting the OTP or OOB token, it is important to ensure that they are transmitted in a secure manner from source to destination. Hence, the next section discusses how the thesis deals with transmission of OTP and OOB tokens across networks that may at times be vulnerable.

8.3.3 Transmission of OTPs and OOB Tokens

The proposed model transmits OTPs and OOB tokens using emails or SMS messages. Security is crucial for example when such messages or emails are transmitted over vulnerable and unsecured networks. The model uses steganography to embed and hide OTPs and OOB tokens in low-fidelity images. The choice to use low-fidelity images is so that they are not too big in size and can be transmitted easily, even on a slow connection. The system is designed in such a manner that the embedded OTPs and OOB tokens are only valid for a specified duration of time and they expire after a certain period. This is the time between the process of issuance and verification at the end-user device and the Identity Provider. The idea is to ensure that compromised OTPs or OOB tokens are never used on the system. Furthermore, reissuance of a new OTP or OOB token invalidates the old one. This makes the Identity Provider to anticipate a new OTP or OOB token.

The process of embedding the encrypted data within a low-fidelity image requires a least significant bit (LSB) watermarking technique. This process involves converting data into a binary format and then hiding the binary stream within the first two bits of each pixel of the image. Embedding the data in the first two bits ensures that there are no noticeable distortions on the image (Sharma and Rajni, 2012). Otherwise, the highest significant bits come with high distortion, which could reveal the hidden messages to unauthorised users. This study uses two bits to ensure that there is sufficient space to hold the OTPs and OOB tokens. An advantage of this approach is that the data would already be in a binary format from the encryption process (as explained the Section 8.3.2), and there is no need for conversion of data. Below follows a brief discussion of the process of embedding OTPs and OOB tokens into the low-fidelity image and how to retrieve it.

8.3.4 Embed OTP and OOB Data in the Low-fidelity Image

The system is designed to randomly generate a low-fidelity and small image of 640 by 320 pixels. The small size of the image is customised according to the login window where it is to be displayed. The idea is not to lose the quality of the image. Furthermore, the small size ensures that the image can be easily transmitted even in low bandwidth networks. The images are unique per user. The uniqueness is derived from a random selection of colours for each pixel of the image. Each pixel has 24 bits, i.e. eight bits for red, green and blue colours respectively. Once the unique image has been created, a request is sent for an OTP or OOB token. The time of issue is recorded and concatenated to the OTP or OOB token string. The string containing the OTP or OOB token and time of issue is encrypted (see the previous subsection 8.3.3). Then the system goes over each pixel of the generated unique image. The first two bits of each pixel of the image are replaced by a set of two bits from the binary ciphertext. This process is repeated until the end of the ciphertext is reached, iterating bit by bit over the pixels of the image. At the end of the ciphertext, the remaining first two bits in each pixel are set to empty binary values, in this case 00. The image with the embedded OTP or OOB token and timestamp is then transmitted over a network to the user.

8.4.5 Extract the OTP and OOB Token Data from the Image

On receiving the embedded image, the user may save it for temporary storage and then upload it along with their user profile ID for authentication. In order to undo the watermarking, each pixel of the image is extracted to determine its binary value. The first two bits of each pixel are extracted and concatenated. This forms the encrypted ciphertext containing the user’s OTP or OOB token and the corresponding timestamp. The resultant ciphertext is then decrypted as explained in Section 8.3.2. The plaintext is also trimmed to remove the extra empty spaces and at this point it is ready for authentication. On correct entry of the plaintext OTP or OOB, a user can then be authenticated. However, if the user fails to enter the right OTP or OOB, the system captures such failed attempts and stores them on the user’s historical data. Each login attempt is logged and stored. Depending on the request, in the case of an OTP the user is allowed three tries before the system could lock them out. However, for an OOB, the user is allowed only two attempts, after which they will be locked out.

This part of the system is designed to prevent an attacker from using a man-in-the-middle attack to try and eavesdrop on user credentials, OTPs or OOB tokens that are transmitted over unsecured networks. This is to ensure that even if an attacker were to be able to intercept the communications, he or she would not just land on plaintext OTP or OOB token. This guarantees the confidentiality and integrity of the OTPs and OOB tokens that get exchanged between the system and the user. It is a layered security approach in that the OTP or OOB tokens are first encrypted before they undergo a steganography process. In Chapter 11, which serves as an evaluation of the study, the author discusses in more detail the impact of the security mechanisms that the system uses on its overall performance and usability. The next section discusses the proof-of-concept prototype implementation of the model.

8.4 PROOF-OF-CONCEPT PROTOTYPE IMPLEMENTATION

This section presents and discusses the proof-of-concept prototype implementation of the strong risk-based multifactor authentication model. The model is implemented based on a number of packages that contain several classes and interfaces. The packages include the following: *Client; RandomNoise; Crypto;, Steganography; Multi-Factors; IdentityProvider; ContextBroker; RiskScoreCard; Repository; Frameworks.* There is also one other component namely the *Eureqa System,* which falls outside the scope of this thesis. Seven of the above-mentioned packages form the core of the risk-based multifactor authentication solution, i.e. *Repository; ContextBroker; RiskScoreCard; Multi-Factors; IdentityProvider; Steganography,* and *Crypto*. The rest of the packages (i.e. *Client*, *RandomNoise* and *Frameworks*) are only meant for support. These supporting packages fall outside the scope of this thesis and hence will only be mentioned and not discussed in detail. Due to the complexity of the risk-based multifactor authentication component, the author decided to not further unpack the UML class diagram. This means that individual classes are discussed under each package.

The *IdentityProvider* package consists of three interfaces, *Interface2ContextBroker*, *Interface2RiskEngine* and *InterfaceMFA*, and one class that is responsible for the routing of requests. At user login, the *IdentityProvider* package receives an authentication request (i.e. user credentials, device information, geo-location and behavioural analytics) from a user and routes the request through a secure socket layer using the *Interface2ContextBroker* to the *ContextBroker*. The *ContextBroker* receives the request and initiates a connection to the repositories using the *JDBC:ODBC* database interface. Once a connection has been established, the *ContextBroker* prompts the different repositories to provide existing details of the specific user (i.e. if the user is already registered) using the *RepoQuery* and *SQLRepoController* classes. The existing user login details are retrieved in an encrypted format from all the relevant repositories. The *ContextBroker* uses the *Decrypt* class in the *Crypto* package to decrypt the data and then passes it to the *Anomaly Detection* class within the *Risk Score Card* package.

The *Anomaly Detection* class checks for outliers between the target login and those in existing records. The *Anomaly Detection* class contains the implementation of the *k-mean* and *elbow* algorithms as mentioned in Section 8.2.6.1. This class is implemented in such a manner that if a login under observation is normal, the *Anomaly Detection* class would make a recommendation to the *Risk Engine* class for it to classify it as low risk and pass the result to the *IdentityProvider* through the *Interface2RiskEngine* interface. The *IdentityProvider* would then authenticate the user and grant them access to the resource that they are requesting. Thereafter, the target login instance with all its supporting details is rerouted to the *ContextBroker* for encryption, using the *Encrypt* class from the *Crypto* package, and relaying different pieces of ciphertext to the relevant repositories for storage, using the *RepoUpdate* class. This guarantees secure storage of the authentication data at rest.

For example, should it so happen that an attacker could breach the system and somehow manage to get hold of the data in the repository, they would need to decrypt it first before they can sell it in the dark web. However, the process of decrypting the data is not something that is widely used like AES or 3DES. The attacker must understand the encryption algorithm that has been used here. Much against the traditional Kerckhoff principle of crypto openness (Knoll, 2018), this study does not publicise the crypto algorithms and resultant keys. This strategy is 'security through obscurity', which refers to concealing crypto algorithms and hoping attackers would not be able to find them (Moshirnia, 2018). 'Security through obscurity' is considered bad practice in information security (Knoll, 2018). However, concealing these algorithms would presumably increase the amount of time and would potentially make it not worthwhile for the attacker. Furthermore, the attacker must be able to breach all four database instances to fully compromise this solution. On top of that, the attacker must be able to understand all the processes of the *ContextBroker,* from the way it splits the data, to how it encrypts and decrypts it. Hence, it can be confidently argued that the distribution of the different pieces of authentication data across the different instances of repositories helps to improve the resilience of the proposed solution. The distribution of the different datasets across the various repositories is handled by the *RepoSplitter* class in the *Repository* package, with specific calls to the *SQLQuaRepo, SQLRepoController* and *RepoManager* classes in the *ContextBroker* package.

Suppose that an attacker successfully compromises the *UserProfile* class and repository and is somehow able to decrypt the usernames and passwords. The attacker would take the plain-text usernames and passwords and attempt a login. For such an attacker to login successfully, he or she will need to ensure that their keystrokes match those of the compromised user account. A failure to do so would start to raise red flags and prompt the attacker to enter a *SecurityPhrase* of variable characters, depending on the risk category in which the system would have classified the user login attempt. The *SecurityPhrase* class, along with the OTP and OOB classes, is found in the *Multi-Factors* package.

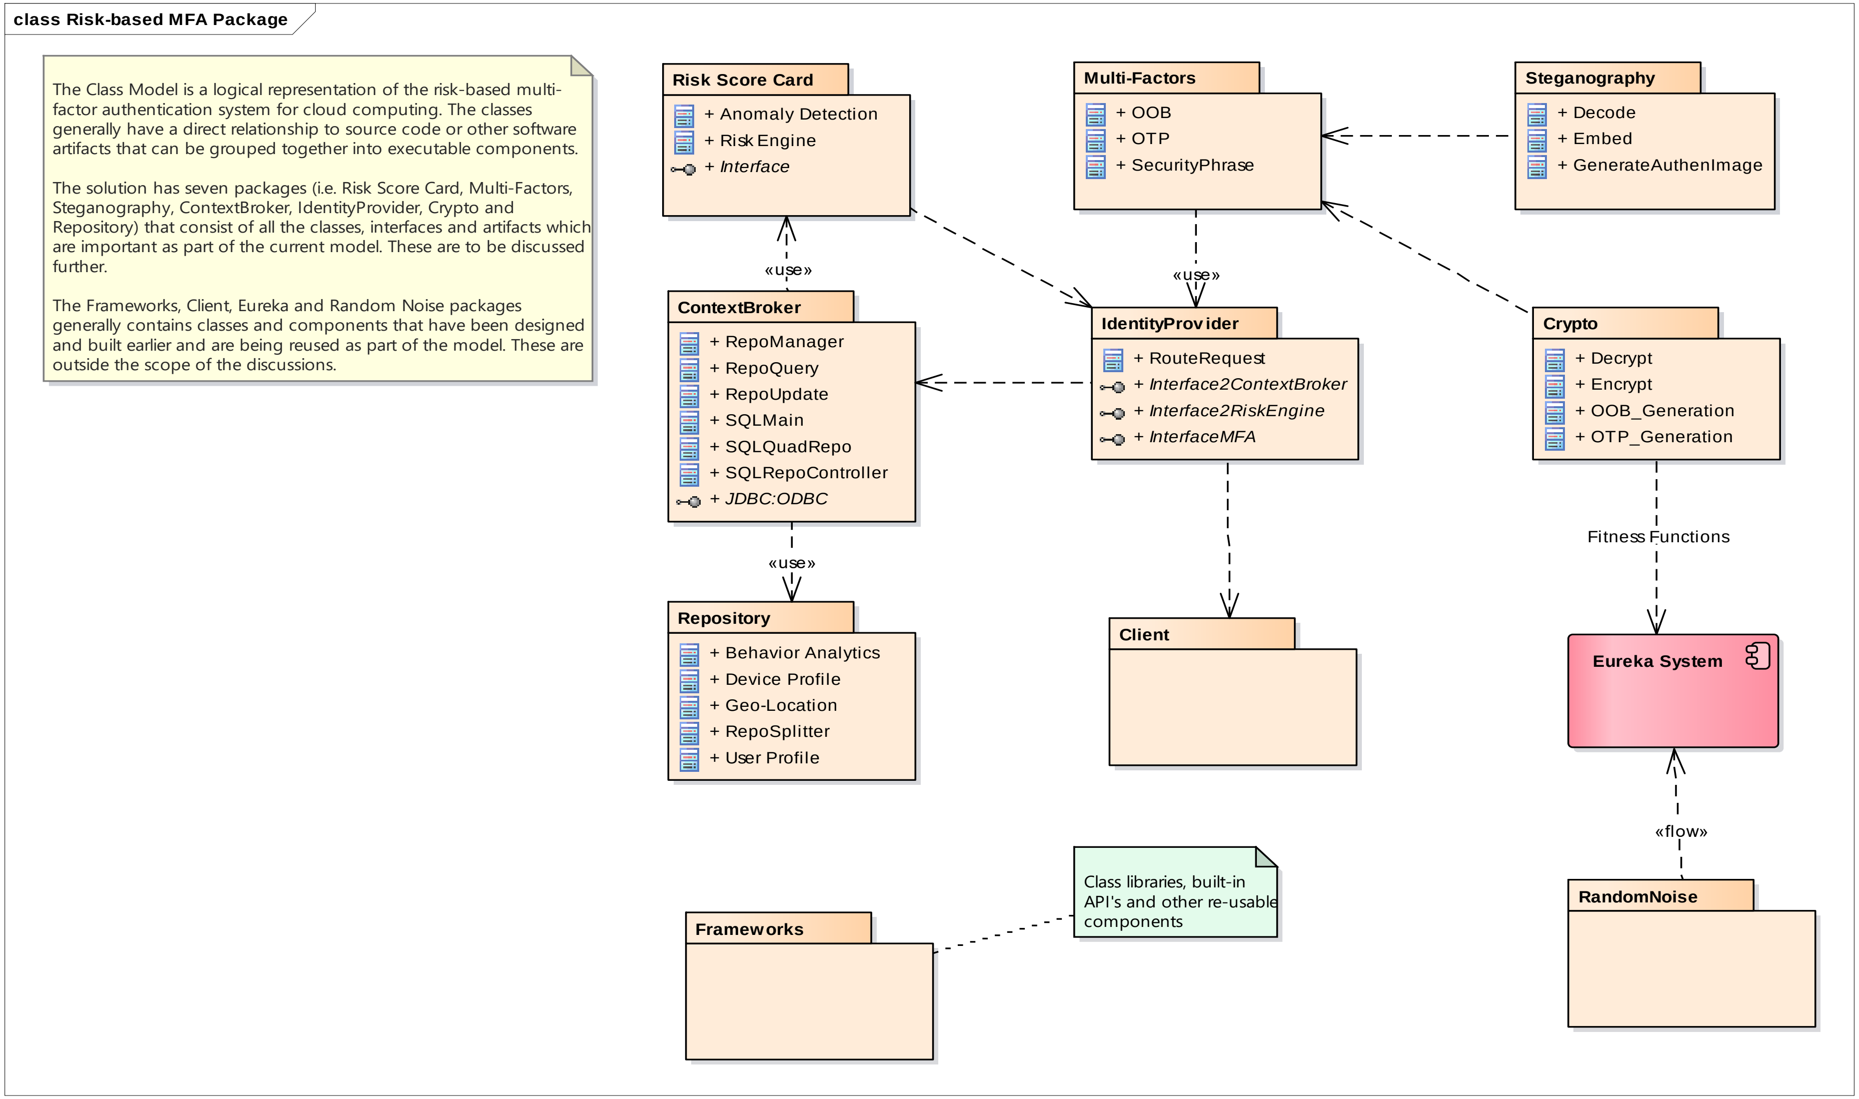

The Class Model is a logical representation of the risk-based multi-factor authentication system for cloud computing. The classes generally have a direct relationship to source code or other software artifacts that can be grouped together into executable components.

The solution has seven packages (i.e. Risk Score Card, Multi-Factors, Steganography, ContextBroker, IdentityProvider, Crypto and Repository) that consist of all the classes, interfaces and artifacts which are important as part of the current model. These are to be discussed further.

The Frameworks, Client, Eureka and Random Noise packages generally contains classes and components that have been designed and built earlier and are being reused as part of the model. These are outside the scope of the discussions.

Class libraries, built-in API's and other re-usable components

Figure 8.4:- UML Package Class Diagram for Implementing a Risk-based Multi-factor Authentication Model for Cloud Computing

The *Multi-Factors* package contains the OOB, OTP and SecurityPhrase classes. These classes are for the different factors that the solution uses for authenticating login attempts of moderate risk to very high risk. This package has two dependencies from the *Steganography* and *Crypto* packages and it uses the *IdentityProvider*, which is dependent on the *RiskScoreCard* package.

The *OOB_Generation* class inherits some of the properties of the *Encrypt* class. Once an OOB token has been generated, it is passed on to the *Encrypt* class for encryption. The encrypted OOB token is next embedded using an *Embed* class to the image that is generated by the *GenerateAuthenImage* class within the Steganography package. Once, the encrypted OOB token has been embedded in the image, it is sent through a secure channel to the *IdentityProvider* using the *InterfaceMFA*. From then, an embedded OOB token is passed to the user.

On receiving the embedded OOB token, the user will temporarily store it for upload on to the authentication user interface. The authentication requires the user to also input their profile ID. The system uses the *Decode* class to strip the OOB token from the image. It then takes the encrypted OOB token and passes it to the *Decrypt* class for decryption. After the process of decryption, the user will finally get the plain-text OOB token for authentication. Since OOB tokens are used for very high-risk login attempts, the user is only afforded two chances. Should the user fail to enter the correct OOB token on both chances, the system will automatically lock them out and request a system administrator's intervention. The same process follows for an OTP. However, the system uses the *OTP_Generation* in the *Crypto* package and *OTP* class in the *Multi-Factors* package. The rest of the classes continue as discussed in the case of an OOB token. The only difference is that with OTPs a user is given three tries instead of two, as is the case with OOB tokens.

The *SecurityPhrase* is handled a little differently from an OTP and OOB token. The difference comes mainly from the way it is generated and how it is used. The *SecurityPhrase* is generated at registration along with the username and password. This is manually done by the user and kept by the system. When a scenario arises that warrants its use, the *IdentityProvider* passes the risk level request to the *SecurityPhrase* class to randomly select the number of characters that the user has to enter to complete it. These are directly proportional to the risk posed by a login attempt. (See the example in Section 8.2.7.) The

SecurityPhrase is secured using the *Encrypt* and *Embed* classes from the *Crypto* and *Steganography* packages respectively. This is done in a similar manner to that of securing OOB tokens and OTPs. Similar to the OTP, a user is given only three tries to correctly enter the missing characters of a *SecurityPhrase* before the system could lock them out.

This marks the end of the discussion of the UML diagram of the proof-of-concept prototype implementation. The rest of this section provides examples of the screenshots of the prototype with a brief discussion.

The screenshots come from Dlamini et al. (2016) and Dlamini et al. (2017), both of which are outputs of the research in hand. Figure 8.5 depicts a screenshot for capturing a SecurityPhrase. The risk indicator at the top right corner is used by the system to determine the actual number of characters that a user would have to put in to complete the SecurityPhrase. In this case, the user has to enter 50% of the characters. This is directly proportional to the risk posed. Since, the keystroke behavioural analytics are not captured at this point, the user can make use of a soft keyboard as displayed in the screenshot.

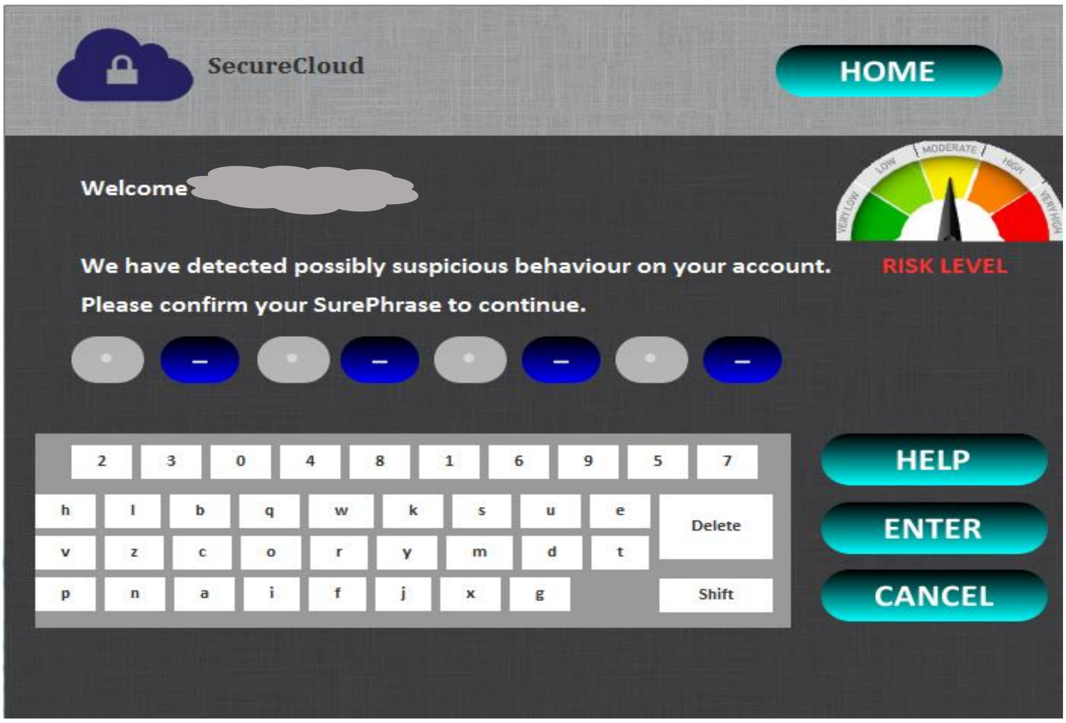

Welcome [redacted]

We have detected possibly suspicious behaviour on your account.

Please confirm your SurePhrase to continue.

Figure 8.5:- A Screenshot for Entering a SecurityPhrase

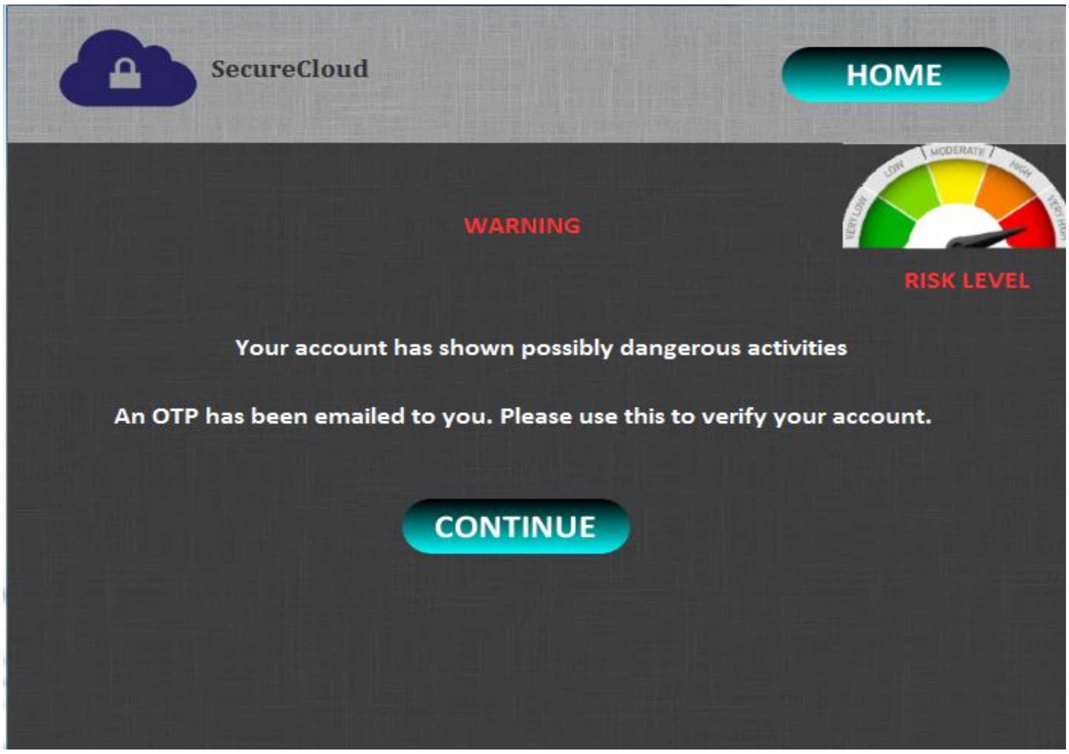

WARNING

Your account has shown possibly dangerous activities

An OTP has been emailed to you. Please use this to verify your account.

Figure 8.6:- A Screenshot for a High Risk

Figure 8.6 depicts a screenshot that is displayed when a user login attempt has been classified as very high risk. In this case, the user would only click *continue* to initiate the process of generating and transmitting an OTP. On clicking *continue*, the user will be prompted to open the embedded image as shown in Figure 8.7. This must be accompanied by the user profile ID. Should the user enter a wrong OTP on all three attempts, he or she will receive a message as reflected in Figure 8.8. The message written in red at the bottom of the screenshot helps the user to keep count of the times they have entered the OTP.

Moreover, and as most authentication systems must do, this system does not categorically state if the problem is with the OTP or the image. This is implemented the same way as with usernames and passwords. Should one of these be wrong, the system should not state exactly which one is wrong, but rather that it could be either of the two.

The next screenshots illustrate how the model handles anomaly detection using the clustering algorithm with ten clusters. The number of clusters were determined using the elbow method. For example, Figures 8.9 – 8.11 below depict an *x*-*y* plot of four login attempts as they are being classified by the solution with ten clusters.

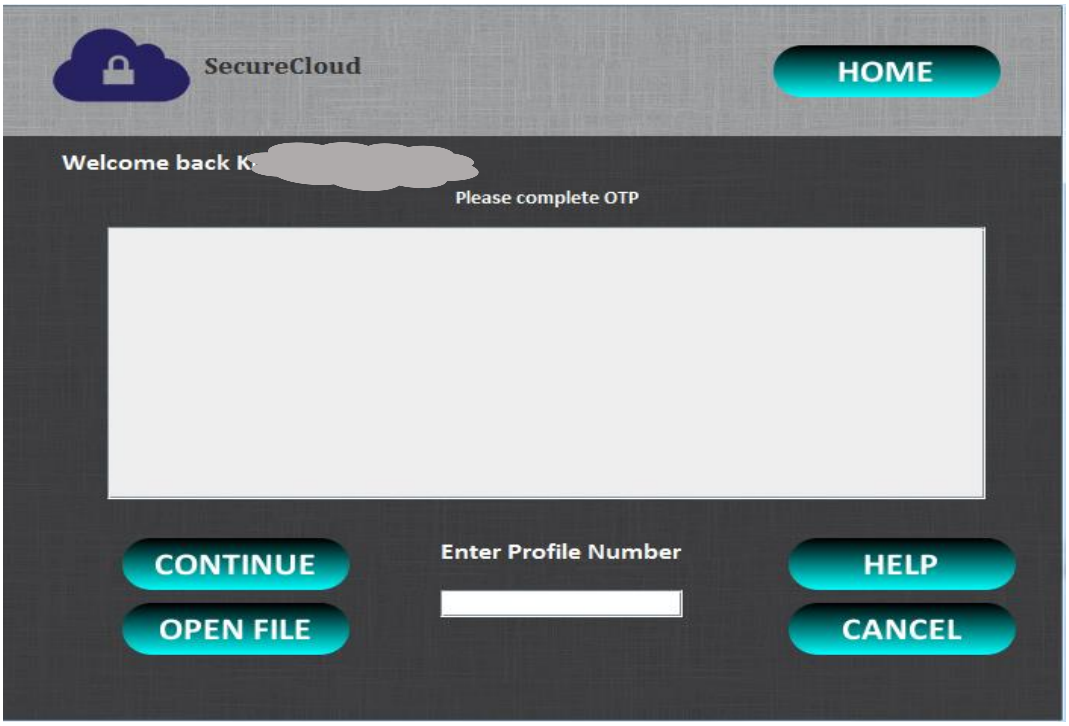

Welcome back K[redacted]

Please complete OTP

Enter Profile Number

Figure 8.7:- A Screenshot for the Embedded Image

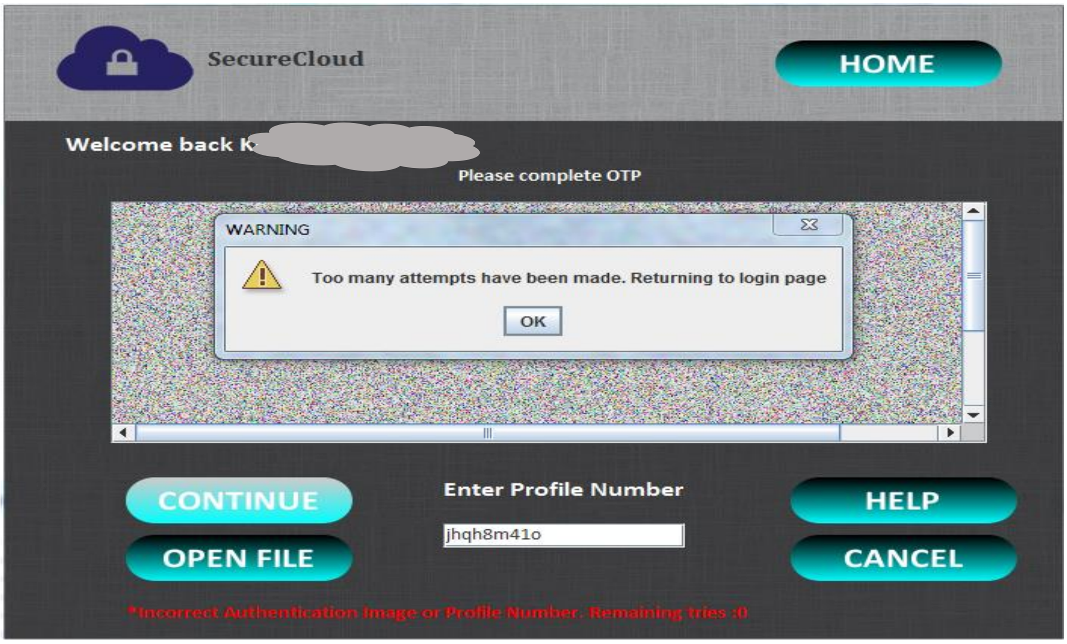

Welcome back K[redacted]

Please complete OTP

WARNING

Too many attempts have been made. Returning to login page

Enter Profile Number

jhqh8m41o

*Incorrect Authentication Image or Profile Number. Remaining tries :0

Figure 8.8:- A Screenshot for more than Three Attempts at Entering an OTP

Figure 8.9 illustrates a login attempt that fails the local anomaly check. The login attempt fails because it is too far from the cluster of context, which in this case is cluster four. However, this is not considered failing a global anomaly, because the login attempt is close to the other clusters. Some existing models would classify such an anomaly as normal pattern and raise no flags as this would cause unnecessary false positives. Hence, the rigid approach taken in this model is to ensure that such cases are correctly classified and flagged as outliers. Figure 8.10 illustrates another example of a login attempt that fails a local anomaly check. Cluster one is the cluster of context for this login attempt. Surely the login attempt in Figure 8.10 seems very close to other clusters and some models would accept it as normal behaviour, yet the proposed solution correctly classifies it as an anomaly. Figure 8.11 depicts a successful login attempt. This successful login attempt's cluster of context is cluster ten, which means it falls within the radius of cluster four's centroid.

In summary, the proposed model uses different clusters to mimic the different login scenarios of a single user. The key point to note is that a single user is associated with multiple keystroke dynamics, which are then represented as different clusters. This depends on a number of factors such as fatigue, loss of concentration and others. For example, when a user is highly alert, the solution classifies him or her towards the bottom right corner of the *x*-*y* plot. As the user starts to get tired or lose concentration, their classification starts to move further up and towards the *y*-axis of the *x*-*y* plot. Cluster one is a good example of fatigue creeping in on the user. The proposed solution caters for all these different scenarios.

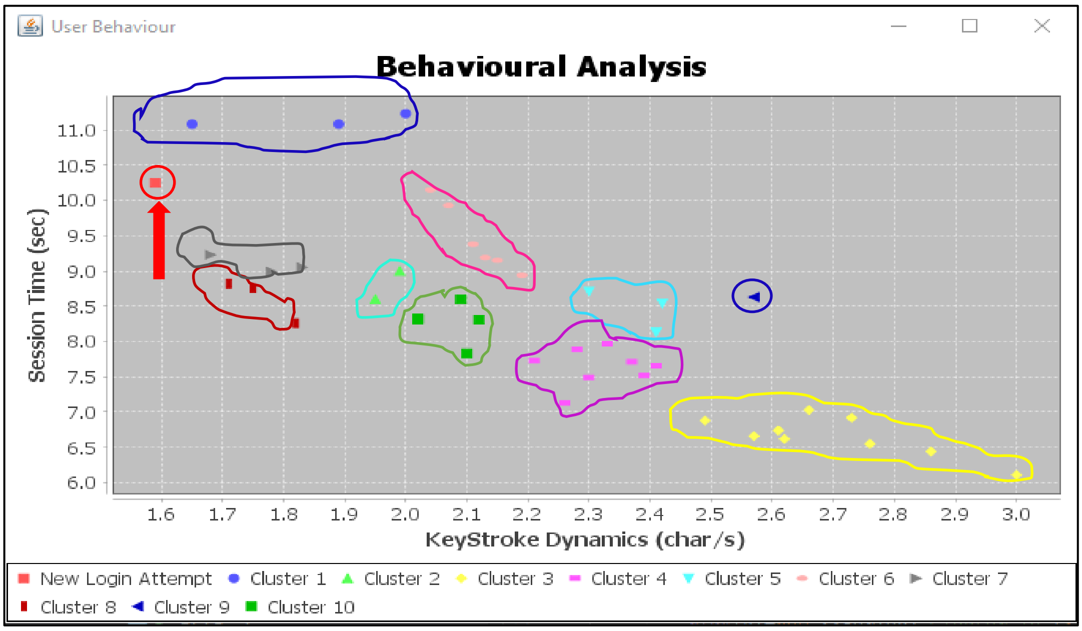

Figure 8.9:- A Failed Local Anomaly Check

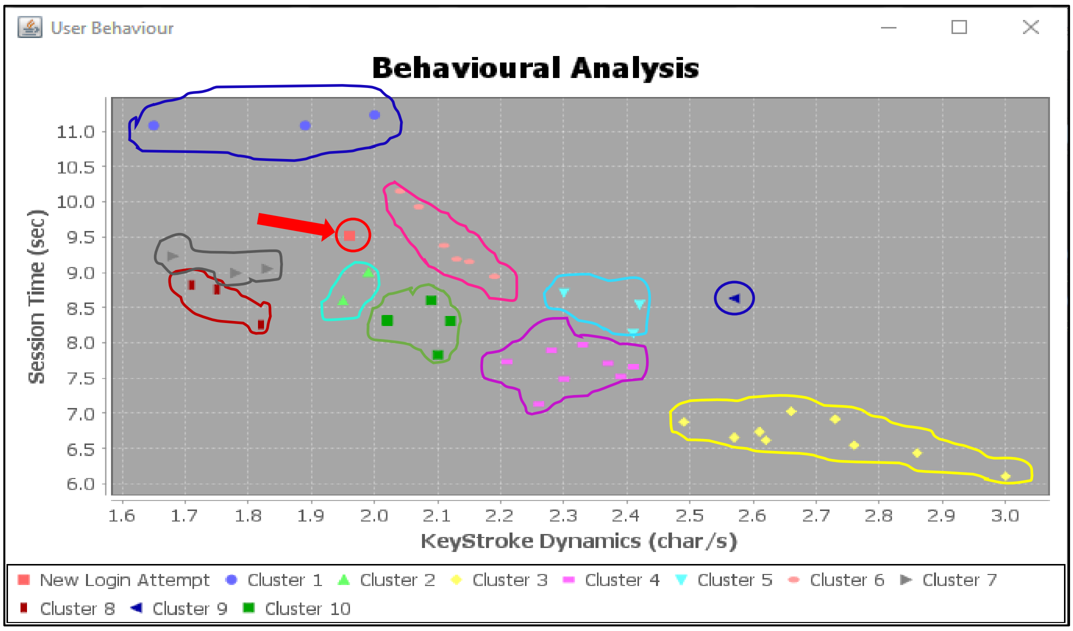

Figure 8.10:- A Failed Local Anomaly Check

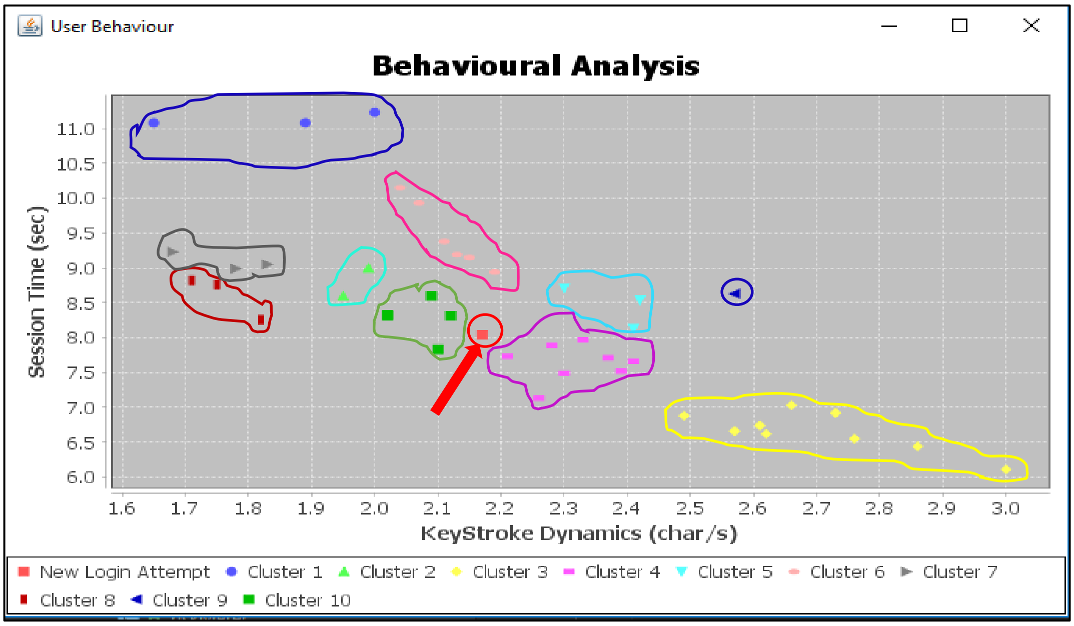

Figure 8.11:- A Successful Login Attempt

In closing the discussion on the proof-of-concept prototype implementation, it can be deduced that the proposed model provides a plausible solution to solve the challenges posed by attackers using compromised user credentials to access systems they are not authorised to access specifically in an IaaS cloud. Even though the tests were done on a small sample, it was interesting to note that the model achieved its authentication goals. It is expected that with time, as users continue to use the model, its accuracy on classification and anomaly detection will improve. The expectation is that as one uses the model over and over again, the data points would tend to converge towards the centroid in such a manner that any slight deviation from the norm would be detected as an outlier. Such an exercise would require a very large data set to increase the sensitivity, accuracy and efficiency of the model. Fortunately, even with a small data set, the model was able to successfully detect both global and local anomalies, as well as normal logins.

The next section concludes Chapter 8 and introduces Chapter 9.

8.5 CONCLUSION

The proposed solution provides strong and risk-based multi-factor authentication that scales up and down, based on the threat levels or risk posed on the system. It provides end-to-end security of user credentials and other authentication data – at rest and in transit. Chapter 8 employed an innovative encryption algorithm that takes chaotic random noise as input to generate encryption iterators that help encrypt and secure authentication data which is then stored in multiple storage locations to increase its resiliency. It also introduced the use of steganography to transport encrypted OTPs, OOB tokens and SecurityPhrases to the user. Furthermore, the solution that was proposed makes use of a number of techniques to effectively deny unauthorised users with “supposedly” stolen user credentials from accessing systems that they are not authorised to use. The techniques include user credentials such as usernames and passwords, behavioural keystroke patterns, device information and geo-location data. The proposed solution has proven to be effective and resilient in barring users with compromised credentials from accessing systems they are not authorised to use.

It is important to ensure strong and risk-based authentication and even more critical to secure the authentication data from criminals. However, it is also vital to monitor the system for malicious behaviour. Such monitoring is key for regulatory compliance purposes and for anticipating a legal or disciplinary hearing. Therefore, this thesis also monitors and records all accesses in a digital forensically sound manner. This component is handled in Chapter 9.

Chapter 9 discusses how this study tackles the digital forensic readiness aspect of cloud computing. It is primarily concerned with proactive collection and preservation of digital evidence in the cloud in anticipation of a legal lawsuit or policy breach investigation. The proactive collection and preservation of digital evidence must be done in a forensically sound manner, and Chapter 9 demonstrates how this is to be achieved. The main goal is to reduce the time it takes to conduct a digital forensic investigation in the cloud.

The evaluation chapter (i.e. Chapter 10) later adds more detail on how the solution affects performance. For example, it will show how the innovative encryption employed in securing authentication data compares to other encryption algorithms like Triple-DES, AES 256, etc.

CHAPTER 9 A 3-TIER DIGITAL FORENSIC READINESS MODEL

9.1 INTRODUCTION

Chapter 8 focused on the strong risk-based multifactor authentication component of the high-level conceptual architecture. Its focus was on dealing with the challenge of unauthorised access to cloud resources due to inadequate or compromised credentials. The model proposed in Chapter 8 took a risk-based approach to authenticate users using multiple factors (i.e. *SecurityPhrase*, *OTP* and *OOB*) and based on specific risk indicators such as keystroke dynamics, device information, and geo-location. The model ensures that users' login attempts are classified based on different risk indicators and it would scale up or down authentication factors based on the risk posed by each login attempt. A combination of security tools was used to secure the authentication factors both in storage and transmission. The end result was a strong risk-based multifactor authentication model.

Chapter 9 now focuses on the digital forensic readiness component of the high-level architecture presented in Chapter 6. Chapter 9 tackles the digital forensic readiness aspect of cloud computing with respect to proactive collection and preservation of digital evidence in anticipation of a legal lawsuit or policy breach investigation. The main goal of this chapter is to reduce the time it takes to conduct a digital forensic investigation in the cloud. Furthermore, this part of the study aims to ensure that digital forensic investigations may be carried out without any or with minimal disruptions to the cloud service provider's infrastructure or business processes. Inherent to the main goal of this chapter, this part of the study aims to notify digital investigators of security threats as they happen or before they occur. The proposed model acts as a system that monitors and captures all user activity in a forensically sound manner. The key contribution of Chapter 9 is a 3-tier digital forensic readiness model that consists of a planning, implementation and assessment tier. The rest of this chapter is structured as follows: Section 9.2 outlines the model. Section 9.3 outlines the proof-of-concept prototype implementation of the 3-tier digital forensic readiness model. Section 9.4 concludes the chapter and provides an overview of Chapter 10.

9.2 A 3-tier Digital Forensic Readiness Model for Cloud Computing

The proposed 3-tier digital forensic readiness model consists of a planning, implementation and assessment tier as shown in Figure 9.1. Each of these tiers is discussed below.

9.2.1 PLANNING TIER

The planning tier consists of processes that are necessary for preparing the cloud for digital forensic readiness. These planning processes include organisations articulating clear business goals and objectives for digital evidence; defining a clear cloud scenario; scrutinising all the applicable laws, regulations, ethics and judiciary admissibility; analysing pre-incident collection, preservation, storage and manipulation of potential admissible digital evidence; planning incident detection; making a pre-incident analysis; and defining a clear cloud system architecture that takes care of all the above. Taking all these into consideration ensures a comprehensive digital forensic readiness system that meets all the requirements stipulated in Chapter 5.

9.2.1.1 Articulating Business Goals and Objectives

For digital forensic readiness to remain effective, it should be driven by the overall business goals and objectives. This will ensure that it gets the necessary buy-in and support from top management. Hence, the planning tier should clearly define the business goals and/or objectives for the collection, preservation, storage and manipulation of potential digital evidence. This should clearly state why a digital forensic readiness capability is required, exactly what digital evidence is to be captured and for what purposes (e.g. in anticipating a litigation or corporate dispute; to demonstrate compliance to regulatory mandates).

9.2.1.2 Defining a Clear Cloud Scenario

This process requires an investigator to carefully study the cloud that a particular organisation wants to adopt (e.g. public or private or hybrid) and consider the type of services sought (e.g. SaaS or IaaS or PaaS or any combination of these). Based on this, an investigator should carefully examine and scrutinise all the scenarios that might pan out and require digital evidence. This process should be able to identify all potential threats and vulnerabilities that might result in a corporate dispute or litigation and may require concrete digital evidence to either support or defend an anticipated case. Within the cloud scenario definition process, it is important for an investigator to carefully study the currently implemented information security measures to identify and plan for the gaps that exist. This is more like performing a risk assessment in order to get a clear understanding of how to proactively retrieve digital evidence in preparing for a fast and effective investigation in the cloud. All these must be done in an ethical manner and comply with the relevant laws and regulations to ensure the judiciary admissibility of digital evidence.

9.2.1.3 Scrutinising Laws, Regulations, Ethics and Judiciary Admissibility of Digital Evidence

Within a cloud computing environment, there is normally a client, network service and cloud service providers – each of which is governed by its own policies and governance, legal, ethical and regulatory frameworks. It becomes important for investigators and law enforcement agencies to consider and familiarise themselves with the laws with which each of these (client, network and cloud service providers) must comply. They must also make an effort to ensure that these laws and regulations are at least aligned to avoid conflicts later on. For this process, organisations should seek the help of a legal expert who can advise them on where their data or applications could be safe. Legal experts should also be called in to advise cloud service providers on the type of clients and the type of data or applications that they are allowed to host, in accordance with the applicable laws in certain jurisdictions of both the clients and service providers.

9.2.1.4 Identifying Potential Sources of Admissible Digital Evidence

The client, network and cloud service providers could all be the potential sources of digital evidence within a cloud environment. The client’s network and access devices of a cloud service provider’s IT systems could hold potentially relevant digital evidence (Chung et al., 2012). Furthermore, the web browser, identity and access control management systems (wherever they sit) could also present sources of potential digital evidence. Hence, it is important to look at all the components that make a cloud infrastructure a reality. In addition to identifying the sources of digital evidence, it is essential to identify the potential admissible digital evidence.

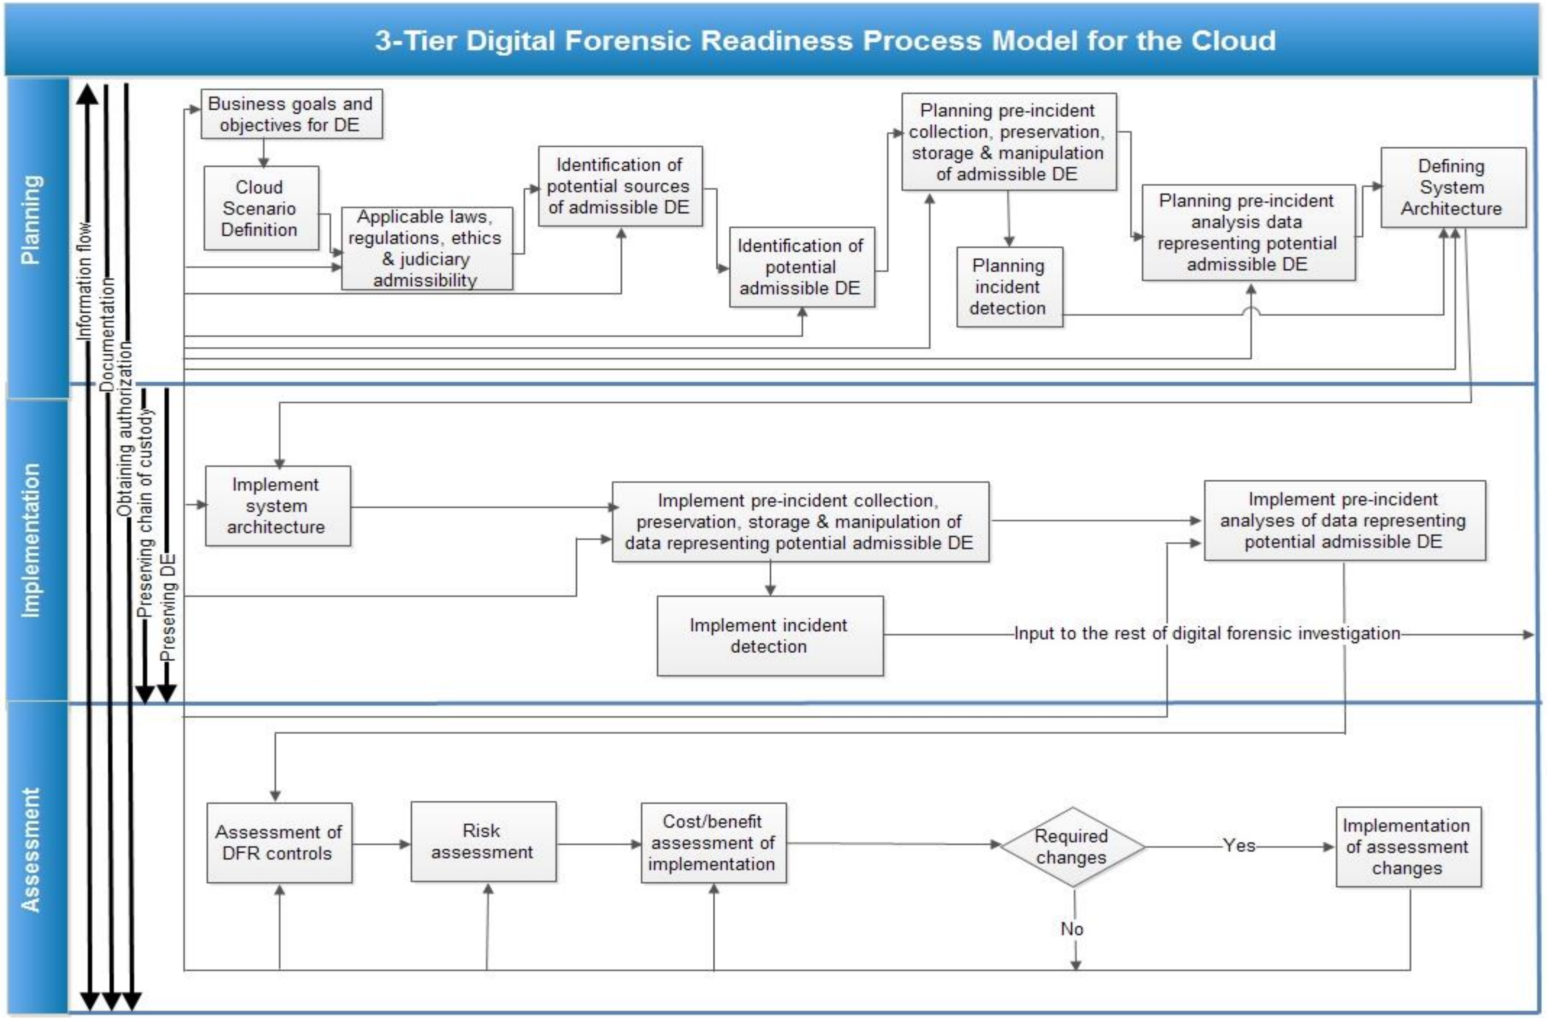
3-Tier Digital Forensic Readiness Process Model for the Cloud
Planning
Implementation
Assessment
Information flow
Documentation
Obtaining authorization
Preserving chain of custody
Preserving DE
Business goals and objectives for DE
Cloud Scenario Definition
Applicable laws, regulations, ethics & judiciary admissibility
Identification of potential sources of admissible DE
Identification of potential admissible DE
Planning pre-incident collection, preservation, storage & manipulation of admissible DE
Planning incident detection
Planning pre-incident analysis data representing potential admissible DE
Defining System Architecture
Implement system architecture
Implement pre-incident collection, preservation, storage & manipulation of data representing potential admissible DE
Implement pre-incident analyses of data representing potential admissible DE
Implement incident detection
Input to the rest of digital forensic investigation
Assessment of DFR controls
Risk assessment
Cost/benefit assessment of implementation
Required changes
Yes
No
Implementation of assessment changes

Figure 9.1:- A 3-Tier Digital Forensic Readiness Model for the Cloud

9.2.1.5 Identifying Potential Admissible Digital Evidence

Once the sources of potential evidence are known, the next step is to determine what sort of digital evidence should be captured that would constitute concrete, credible, authentic and admissible digital evidence in a court of law. Clients should seek the help of a legal expert even on this process to ensure that they capture only concrete, credible, authentic and admissible digital evidence. Potential admissible digital evidence such as activity and audit logs could be obtained from the client's access devices, network provider systems and the service provider's identity and access management systems. Given the magnitude of the anticipated litigation or corporate dispute, cloud clients, investigators and legal experts should ensure that they look at the cost and benefit of acquiring digital evidence to ensure that they observe the proportionality doctrine as well as privacy and human rights of the cloud clients. Once potential admissible digital evidence is identified, it is necessary to start planning its collection, preservation storage and manipulation.

9.2.1.6 Planning Pre-incident Collection, Preservation, Storage and Manipulation of Digital Evidence

This process deals with planning on how to handle potential digital evidence from the point of collection, preservation, storage and manipulation. Involving a legal expert in it could help ensuring that all employees and third parties of the cloud clients are well aware of how to handle potential digital evidence; how to preserve the digital evidence's chain of custody; how not to contaminate the digital evidence at collection, preservation, storage, transportation and manipulation; and how long to preserve digital evidence according to legal mandates and internal retention policies. This plan should include incident detection triggers to raise alarms when malicious incidents are detected.

9.2.1.7 Planning Incident Detection

This process is a plan to monitor on a real-time basis all the activities and/or behaviour of cloud clients and service providers. It outlines the actions to be performed when an incident is reasonably anticipated or actually detected. Within the incident detection plan, the activity monitoring service defines different rules to raise a green flag for all legal activities, a yellow flag for a likely anomalous activity, and a red flag for an anomaly. All activities should be remotely logged in a secure repository to avoid tampering. Moreover, activities that raise a yellow flag should send an alert to an information security officer, while those that raise a red flag should be sent to both an information security officer and an investigator for potential investigation.

9.2.1.8 Planning the Pre-incident Analysis of Data Representing Potential Admissible Digital Evidence

The pre-incident analysis plan pre-collates all the events that might lead to an incident, more like intrusion detection systems do, to outline or pre-conceive some of the pre-conditions of an attack. This pre-incident analysis plan analyses a series of activities in order to detect suspicious events or activities that might be deemed malicious. It then raises yellow or red flags to the responsible personnel and logs all such activities in a forensically sound manner. The plan pre-collates all events that have been flagged as malicious. It does this in a chronological order to reflect on what actually happened, how it started, from whom and when, in order to get to the root cause of the incident and reflect on what took place leading to an incident.

9.2.1.9 Defining a System Architecture

This process takes into consideration all the other processes within the planning tier. It defines and outlines a comprehensive system architecture that would achieve the objectives of digital forensic readiness and would prepare the cloud for a fast and effective investigation.

9.2.2 IMPLEMENTATION TIER

The implementation tier takes over from the planning tier to actually execute and implement the digital forensic readiness plan for a specific cloud scenario that could either be a public, private or hybrid cloud deployment, offering either IaaS or PaaS or SaaS, or a mixture of the services. This tier comprises the actual system architecture; pre-incident collection, preservation, storage and manipulation of potential admissible digital evidence; incident detections, pre-incident analysis; and escalation of other incidents for investigation. Each process of the implementation tier is discussed in detail below.

9.2.2.1 Implement System Architecture

The system architecture that was defined in the planning tier gets implemented as part of this process and covers all the aspects of the planning phase (as comprehensive as detailed therein). The system architecture is implemented according to all the defined plans within the planning tier.

9.2.2.2 Implement Pre-incident Collection, Preservation, Storage and Manipulation of Data that Represents Potential Admissible Digital Evidence

All the digital forensic readiness measures that were outlined in the planning tier to help collect, preserve, store and manipulate the potential admissible digital evidence are put in place to do just that. These include the implementation of an activity monitoring system along with all the defined rules of detecting suspicious activities that may result in adverse consequences. This process must be carried out as outlined in the system architecture. This is the actual implementation of all the above pre-incident plans into the actual operational system.

9.2.2.3 Implement Incident Detection

The activities from the incident detection plan are implemented in this process. All activities that raise red flags are sent to an information security officer and investigator for further investigation. All yellow flags are sent to an information security officer for further analysis before the activities concerned could escalate to potential incidents. Those that seem to be critical, get escalated to an incident and are then sent to an investigator as red flags by the information security officer. All green flags for other activities that are viewed as appropriate, get logged for regulatory and compliance purposes but not necessarily for digital forensic readiness purposes.

9.2.2.4 Implement Pre-incident Analyses of Data that Represents Potential Admissible Evidence

As part of this process, an investigator or information security officer must collate all events that might lead to a yellow or red flag being raised. If done properly, this could be used to prevent an incident from happening in the first place. An investigator will have to monitor a subset of the events and if they play out, he or she would be called to react fast to stop an incident before it occurs. However, if the incident has already occurred, this process could help investigators to re-construct the events in a chronological order to be able to get to the bottom of the incident. Feeding this back to the pre-defined rules and pre-conditions of system monitoring could help prevent a recurrence of incidents.

9.2.3 ASSESSMENT TIER

The assessment tier periodically reviews and assesses all the planned and implemented digital forensic readiness processes to determine if they are all effectively gathering credible digital evidence that is admissible in court and from all identified sources. For example, this tier should constantly check to determine if clients have increased the number of service providers, check if they have switched providers, and find out if they are using other types of access devices that might be different from the ones in the initial plan. This tier also reviews all the necessary processes in order to determine if they all still meet the defined business goals and objectives of proactively collecting digital evidence in the cloud. Furthermore, the assessment tier also checks if the collected evidence is actually gathered in an ethical and legal manner that respects the users' privacy and human rights; and if it is collected according to all applicable laws and complies with regulatory mandates. Should this not be the case, the whole process has to be repeated from the planning tier through to the assessment tier. The assessment tier consists of the assessment of DFR controls; risk assessment; cost/benefit analysis; and then it goes back to the implementation of the assessment changes discussed below.

9.2.3.1 Assessment of Digital Forensic Readiness Controls

This process is meant to review all the controls that have been planned and implemented to proactively collect, store, preserve and manipulate relevant potential digital evidence. These controls have to be assessed to determine if they are collecting and preserving digital evidence as outlined in the initial plan to help achieve business goals and objectives; according to the applicable laws; and are still responsive to the green, yellow and red flags raised by the incident detection controls. Furthermore, the controls are assessed to check if they are not picking false positives, if log collection measures are not compromised, whether evidential data is stored in a secure manner that preserves its integrity, and whether a clear chronological chain of custody is kept that is inaccessible to unauthorised persons. The DFR controls are also assessed to check if they still adhere to the stipulated rules of digital evidence admissibility as stated in the different laws from different jurisdictions.

9.2.3.2 Risk Assessment

This process is based on the fact that the security landscape is constantly changing, with online criminals constantly being a step ahead of the security officers. This race requires that the proposed model be equipped with the capability to constantly review the security threat landscape to identify changes that might require some modifications in the implemented DFR controls. It reviews already existing and new technologies to see if they pose any threats and vulnerabilities. This process also assesses the client devices, their web browsers and connections in terms of threats and vulnerabilities. It goes further to identify trends that criminals use to ensure that the DFR controls do not miss any key incriminating evidence.

9.2.3.3 Cost/Benefit Analysis

This process addresses the ‘proportionality doctrine’ requirement. It assesses the recommended changes coming from the assessment of DFR controls and risk assessment to check if the cost of implementing some of the recommended changes on the DFR capability would not out-weight the benefits of doing so. It is during this process that recommended changes get either rejected or implemented. Sometimes, this process could lead to a decision that requires going back to the starting point.

9.2.3.4 Implementation of Selected Assessment Changes

During this process, only assessment changes that do not violate the laws of cost/benefit analysis get to be implemented into the implementation tier. This means that only those changes where the benefits far outweigh the cost, get implemented. Otherwise, the whole process has to be repeated from the first step. The next section discusses the implementation of the three-tier DFR model.

9.2 IMPLEMENTATION OF PROOF-OF-CONCEPT PROTOTYPE

This section discusses the three-tier digital forensic readiness model's proof-of-concept prototype implementation. This part of the model is implemented as a proof-of-concept in notepad++ php using the following core packages – *Collect Logs*, *Encrypt*, *Index*, *Store* and *Framework* – as depicted in Figure 9.2. The *Framework* package contains all other supporting classes that lie beyond the discussions in this thesis. This package, as in the case in Chapter 8, consists of other user-defined classes that do not form the core of the solution, but provide support. Hence, the author's decision not to include its details in the discussions.

The *Collect Logs* package consists of classes that create the necessary connections that are required to collect user activities. This package has *hooks*, *event listeners* and *triggers* for accessing and collecting the necessary logs as digital evidence. The *Hooks* class is responsible for accessing and gathering user activities. This class, which creates the *Event Listeners* and *Triggers*, aims to prevent malicious activities from causing incidents in the first place. The *Log* class receives all data from the *Hooks* class. *Event Listeners* and *Triggers* classes pick anomalies from the *Log* class. The *Alert & Notify* class sends notifications to investigators when an *Event Listener* or a *Trigger* picks up a log with a red flag indicating malicious activities. The *Capture* class is responsible for capturing the digital evidence and having it ready for encryption and hashing.

The *Encrypt* package consists of AES, Chaos-Based, DES, Triple-DES and Key Management classes that aim to preserve the confidentiality of digital evidence. These classes are collectively responsible for securing the captured digital evidence and providing confidentiality guarantees. An investigator can select their preferred encryption algorithm. The choice is based on the sensitivity of the case at hand and the importance of the digital evidence. The choice of an encryption algorithm is also based on admissibility requirements of individual jurisdictions. For example, investigators in different jurisdictions may be mandated to use a specific encryption algorithm for digital evidence admissibility. The *Key Management* class is meant to secure the encryption keys after they have been used to encrypt the digital evidence. This also depends on the encryption algorithm that is being used.

The *Index* package consists of the one-way *Hash* and *Integrity Checker* classes. These two classes are responsible for ensuring that the captured digital evidence is not altered to compromise its integrity. They ensure that digital evidence is not tampered with right through the entire life cycle of an investigation process. One-way hashes provide one of the fundamental elements for digital evidence admissibility with respect to its reliability and integrity (Stone, 2015; Schmitt and Jordaan, 2013). Schmitt and Jordaan argue that the quickest way for digital evidence to fail the admissibility test in court would be to collect it in an illegal manner or to modify it after it was captured. Hence, the *Hash* class ensures that digital evidence is admissible in court.

This is done using CRC 32b cryptographic hashing algorithm. The choice of CRC 32b is based on the evaluation results of the different hashing algorithms that are covered in Chapter 11. At it this point, it can be highlighted that CRC 32b uses eight-character hashes, which require less storage space compared to other hashing algorithms like MD5. Furthermore, the choice is motivated by the fact that CRC 32b is good at detecting hash collision. For example, a single character difference in digital evidence would produce a completely different hash entry. The *Integrity Checker* class is used to compare and demonstrate that hashes of an original and a copy of digital evidence are the same. The idea is basically to check and verify that the integrity of the digital evidence is not compromised. It illustrates that digital evidence has not been altered or modified in any way, thus demonstrating that its integrity has been maintained and preserved from its original state (Schmitt and Jordaan, 2013).

The *Store* package consists of *Handshake*, *Establish & Verify Secure Channel*, *Verify Repo* and *Verify Source* classes. These classes collectively initiate communication with the remote repository, transmit and store the digital evidence in the remote repository. The *Handshake* class is responsible for establishing a dedicated communication channel between the source of digital evidence and the remote repository. The *Handshake* class calls the *Verify Repo* class to ensure the authenticity of the remote repository. Furthermore, the remote repository calls the *Verify Source* class to verify the authenticity of the source initiating the communication. This is to authenticate the source (capturing device) and destination (remote repository) of digital evidence. The *Establish & Verify Secure Channel* class ensures that the source and destination communicate over a secure channel in order to reduce the effects of a man-in-the-middle attack from eavesdropping while digital evidence is being transmitted.

The rest of this section provides some of the screenshots of the proof-of-concept prototype with a brief discussion. The proof-of-concept was implemented on a cloud platform called OwnCloud, which will be discussed further in the evaluation chapter.

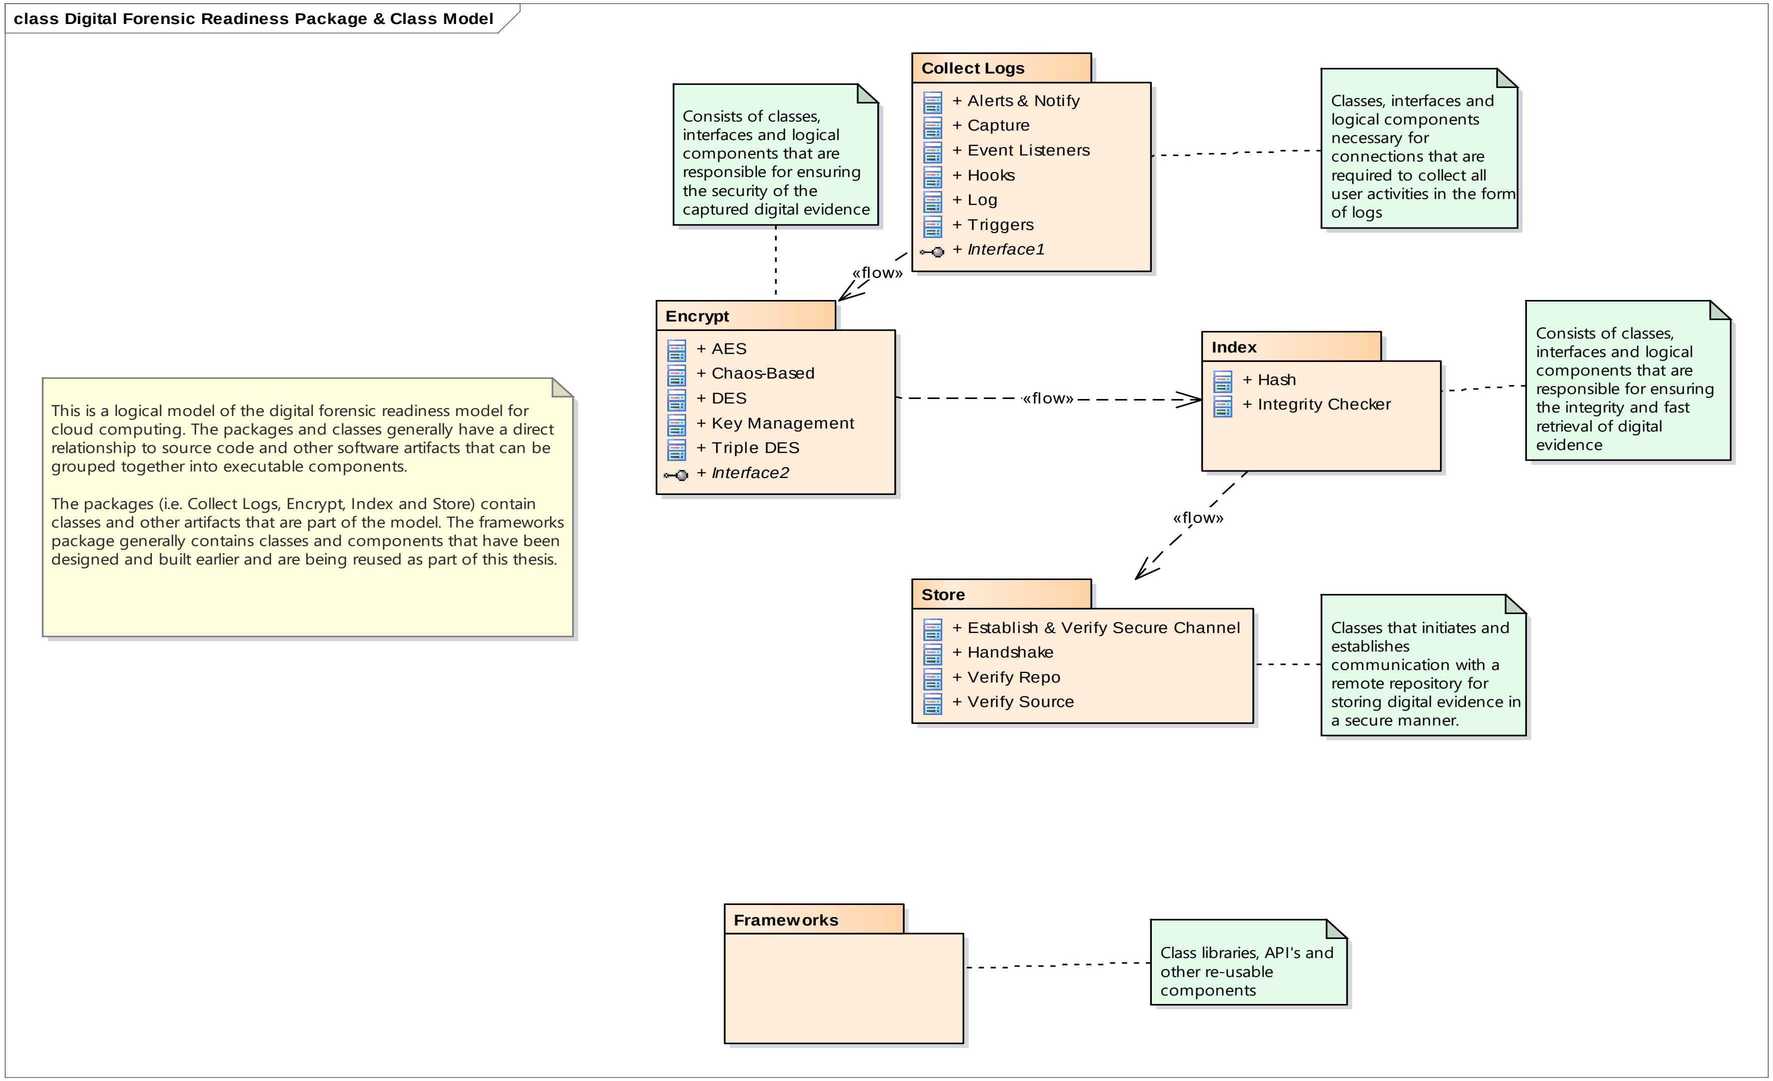
class Digital Forensic Readiness Package & Class Model
Collect Logs
+ Alerts & Notify
+ Capture
+ Event Listeners
+ Hooks
+ Log
+ Triggers
+ Interface1
Classes, interfaces and logical components necessary for connections that are required to collect all user activities in the form of logs
Consists of classes, interfaces and logical components that are responsible for ensuring the security of the captured digital evidence
«flow»
Encrypt
+ AES
+ Chaos-Based
+ DES
+ Key Management
+ Triple DES
+ Interface2
«flow»
Index
+ Hash
+ Integrity Checker
Consists of classes, interfaces and logical components that are responsible for ensuring the integrity and fast retrieval of digital evidence
«flow»
Store
+ Establish & Verify Secure Channel
+ Handshake
+ Verify Repo
+ Verify Source
Classes that initiates and establishes communication with a remote repository for storing digital evidence in a secure manner.
Frameworks
Class libraries, API's and other re-usable components
This is a logical model of the digital forensic readiness model for cloud computing. The packages and classes generally have a direct relationship to source code and other software artifacts that can be grouped together into executable components.
The packages (i.e. Collect Logs, Encrypt, Index and Store) contain classes and other artifacts that are part of the model. The frameworks package generally contains classes and components that have been designed and built earlier and are being reused as part of this thesis.

Figure 9.2:- UML Package Class Diagram for Implementing a 3-Tier Digital Forensic Readiness Model for Cloud Computing

However, the author reflects on some of the screenshots taken from previous work (Maistry, 2015) as part of this thesis. The discussions exclude the screenshots for login to the system. This is because digital forensic investigators, law enforcers and compliance officers login as illustrated in Chapter 8. All their activities are also tracked for accountability to ensure a proper chain of custody and improve the admissibility of potential digital evidence in court for legal lawsuits. Once authenticated, investigators can start searching for digital evidence belonging to a specific user under investigation. The search results could be segmented based on specific dates as depicted in the screenshot in Figure 9.3

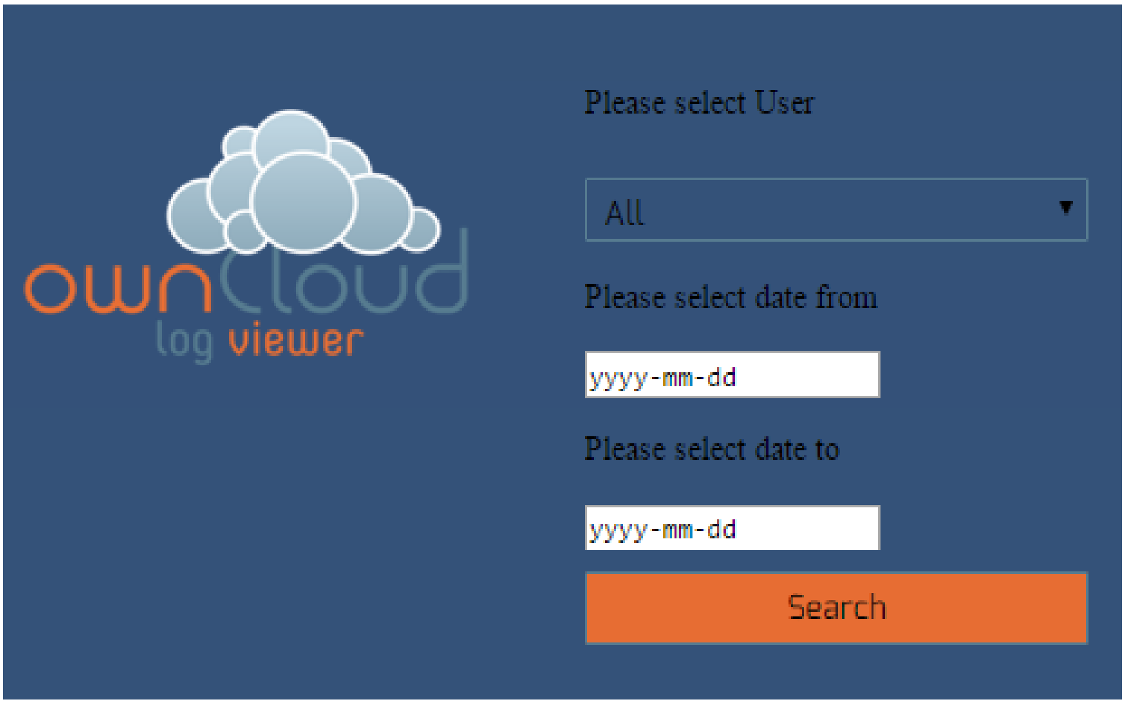

Please select User

All ▾

Please select date from

yyyy-mm-dd

Please select date to

yyyy-mm-dd

Figure 9.3:- A Screenshot Showing User Search Criteria

An investigator may select to see all digital evidence for all users or may choose to select it for a specific user as shown in Figure 9.4. A compliance officer might be interested in knowing that all user activities are monitored.

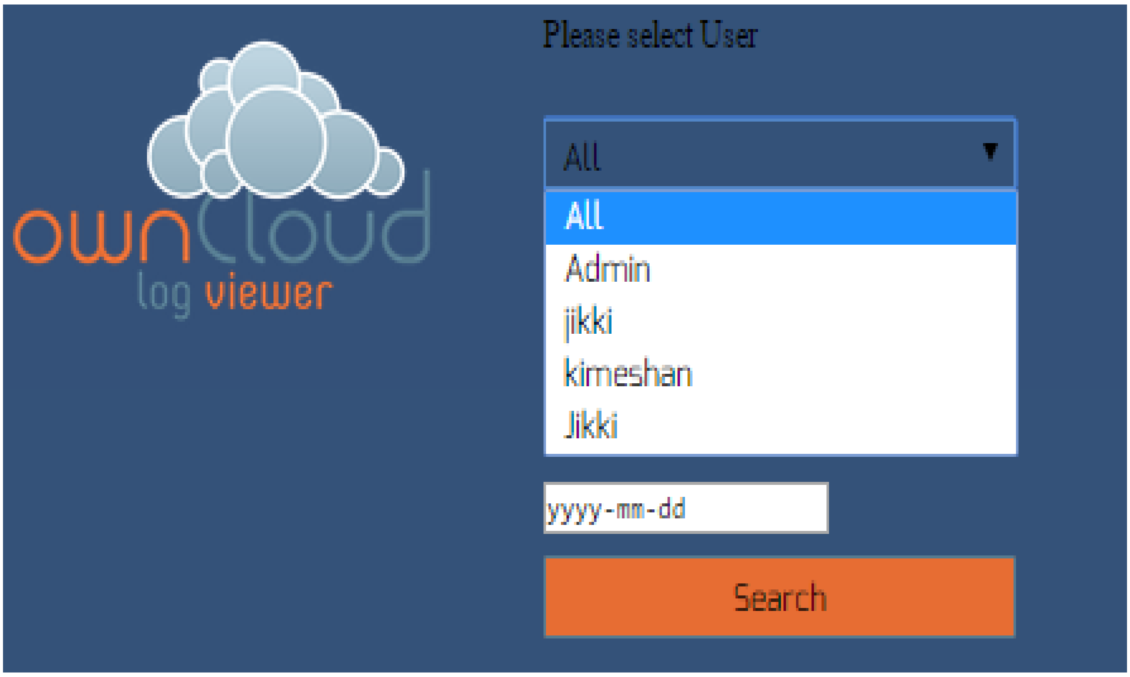

Figure 9.4:- A Screenshot for User Selection

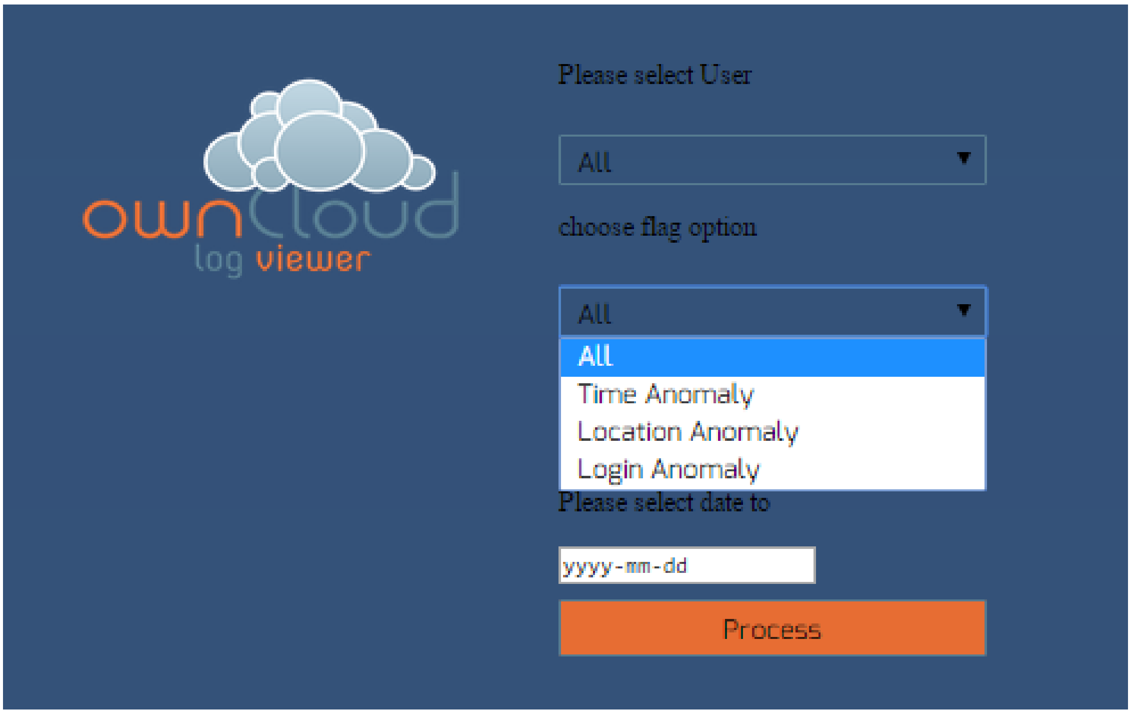

Figure 9.5:- Screenshot for Searching Digital Evidence based on Anomalies

The screenshot in Figure 9.5 can be used as supporting evidence for compliance regulatory purposes to prove that indeed the cloud service provider and all its clients have put measures in place to proactively collect digital evidence in anticipation of digital forensic investigations. The proposed solution has an option for digital forensic investigators to search evidence based on specific anomalies (e.g. login; location; time anomaly; all types of anomalies). A time anomaly would for example occur when a user whose normal sessions are between 08:00 – 17:00 all of a suddenly starts a session at 01:00. The system would put a red flag on this session and add it as an anomaly because it deviates significantly from normal behaviour.

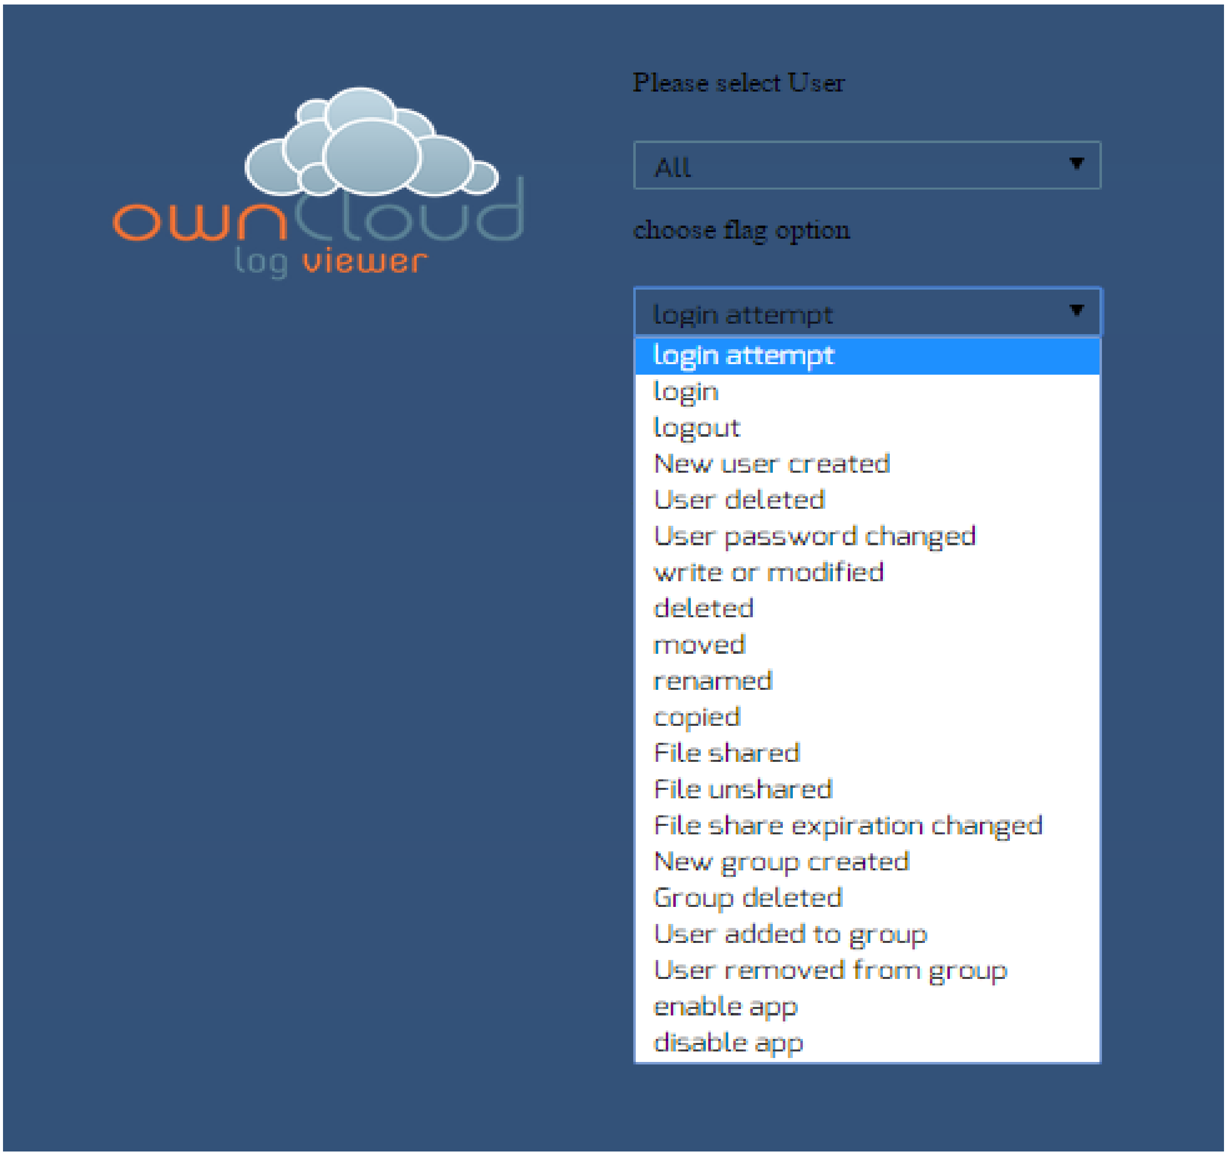

Please select User

All

choose flag option

login attempt

Figure 9.6:- A Screenshot for Filtering Digital Evidence based on Specific User Action

An investigator may also decide to search only based on specific user actions. For example, a search for digital evidence related to logins, user administration, file processing and sharing. This search can also be restricted to specific dates. (See Figure 9.6.)

Figure 9.7 shows a colour-coded result of a search based on login attempts. This screenshot illustrates unsuccessful and successful login attempts, and logouts. The colour codes are not in any way related to the criticality or severity of an anomaly. They are merely used to group a combination of anomalies together. For example, at ID 54 and 55 a user's first attempt to login at 19:15:32 is unsuccessful and at 19:16:20 is successful. This grouping of anomalies can help to construct events and to illustrate the chronological order of events.

Figure 9.8 provides a snapshot overview of the statistics on all activities of individual users. This screenshot can also be used by investigators to quickly identify problematic user accounts without going through the entire process of analysing the captured digital evidence. For example, user *ikki* has sixteen unsuccessful login attempts with just three successful ones. Furthermore, this user seems to be very active with a *write*, *move*, *rename* and has shared a file. Given the number of unsuccessful login attempts compared to the successful ones and the suspicious activity, this user might need tight monitoring.

Figure 9.9 is a screenshot that illustrates the implementation of CRC 32b cryptographic hash algorithm to verify the integrity of potential digital evidence. This screenshot shows unreliable entries in a darker colour, e.g. the entry at ID 7. This indicates either that the file *system diagram.png* has been tampered with by another investigator, or that it has been updated at the source without the update reflecting on the digital evidence side. If the evidence is tampered with by an investigator, this must be verified on the accountability records of the system. However, should it be that the versions are not synchronised, the investigator must do a refresh that updates both versions.

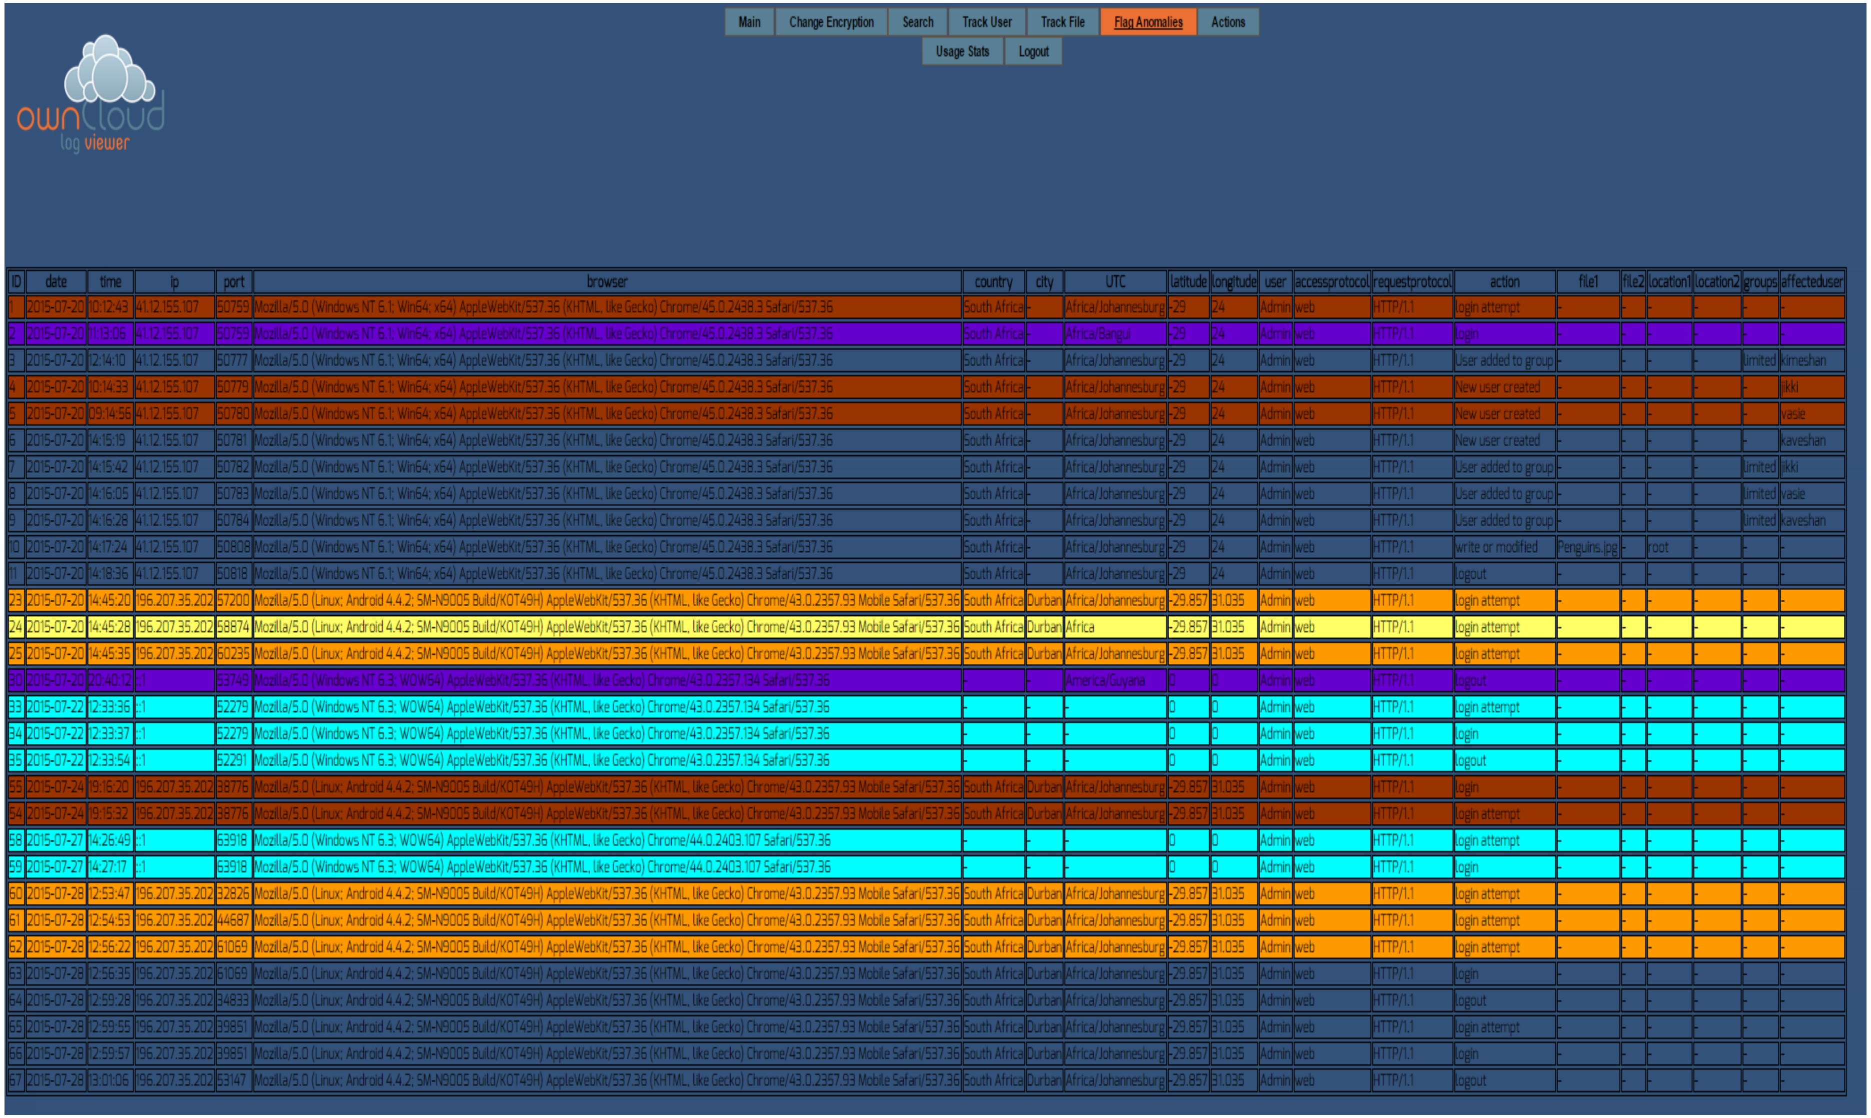

ID	date	time	ip	port	browser	country	city	UTC	latitude	longitude	user	accessprotocol	requestprotocol	action	file1	file2	location1	location2	groups	affecteduser
1	2015-07-20	10:12:43	41.12.155.107	50759	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Johannesburg	-29	24	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
2	2015-07-20	11:13:06	41.12.155.107	50759	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Bangui	-29	24	Admin	web	HTTP/1.1	login	-	-	-	-	-	-
3	2015-07-20	12:14:10	41.12.155.107	50777	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Johannesburg	-29	24	Admin	web	HTTP/1.1	User added to group	-	-	-	-	limited	kimeshan
4	2015-07-20	10:14:33	41.12.155.107	50779	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Johannesburg	-29	24	Admin	web	HTTP/1.1	New user created	-	-	-	-	-	jikki
5	2015-07-20	09:14:56	41.12.155.107	50780	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Johannesburg	-29	24	Admin	web	HTTP/1.1	New user created	-	-	-	-	-	vasie
6	2015-07-20	14:15:19	41.12.155.107	50781	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Johannesburg	-29	24	Admin	web	HTTP/1.1	New user created	-	-	-	-	-	kaveshan
7	2015-07-20	14:15:42	41.12.155.107	50782	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Johannesburg	-29	24	Admin	web	HTTP/1.1	User added to group	-	-	-	-	limited	jikki
8	2015-07-20	14:16:05	41.12.155.107	50783	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Johannesburg	-29	24	Admin	web	HTTP/1.1	User added to group	-	-	-	-	limited	vasie
9	2015-07-20	14:16:28	41.12.155.107	50784	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Johannesburg	-29	24	Admin	web	HTTP/1.1	User added to group	-	-	-	-	limited	kaveshan
10	2015-07-20	14:17:24	41.12.155.107	50808	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Johannesburg	-29	24	Admin	web	HTTP/1.1	write or modified	Penguins.jpg	-	root	-	-	-
11	2015-07-20	14:18:36	41.12.155.107	50818	Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2438.3 Safari/537.36	South Africa	-	Africa/Johannesburg	-29	24	Admin	web	HTTP/1.1	logout	-	-	-	-	-	-
23	2015-07-20	14:45:20	196.207.35.202	57200	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
24	2015-07-20	14:45:28	196.207.35.202	58874	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa	-29.857	31.035	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
25	2015-07-20	14:45:35	196.207.35.202	60235	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
30	2015-07-20	20:40:12	::1	53749	Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.134 Safari/537.36	-	-	America/Guyana	0	0	Admin	web	HTTP/1.1	logout	-	-	-	-	-	-
33	2015-07-22	12:33:36	::1	52279	Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.134 Safari/537.36	-	-	-	0	0	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
34	2015-07-22	12:33:37	::1	52279	Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.134 Safari/537.36	-	-	-	0	0	Admin	web	HTTP/1.1	login	-	-	-	-	-	-
35	2015-07-22	12:33:54	::1	52291	Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.134 Safari/537.36	-	-	-	0	0	Admin	web	HTTP/1.1	logout	-	-	-	-	-	-
55	2015-07-24	19:16:20	196.207.35.202	38776	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	login	-	-	-	-	-	-
54	2015-07-24	19:15:32	196.207.35.202	38776	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
58	2015-07-27	14:26:49	::1	63918	Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.107 Safari/537.36	-	-	-	0	0	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
59	2015-07-27	14:27:17	::1	63918	Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.107 Safari/537.36	-	-	-	0	0	Admin	web	HTTP/1.1	login	-	-	-	-	-	-
60	2015-07-28	12:53:47	196.207.35.202	32826	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
61	2015-07-28	12:54:53	196.207.35.202	44687	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
62	2015-07-28	12:56:22	196.207.35.202	61069	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
63	2015-07-28	12:56:35	196.207.35.202	61069	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	login	-	-	-	-	-	-
64	2015-07-28	12:59:28	196.207.35.202	34833	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	logout	-	-	-	-	-	-
65	2015-07-28	12:59:55	196.207.35.202	39851	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
66	2015-07-28	12:59:57	196.207.35.202	39851	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	login	-	-	-	-	-	-
67	2015-07-28	13:01:06	196.207.35.202	53147	Mozilla/5.0 (Linux; Android 4.4.2; SM-N9005 Build/KOT49H) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/43.0.2357.93 Mobile Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.857	31.035	Admin	web	HTTP/1.1	logout	-	-	-	-	-	-

Figure 9.7:- A Screenshot for Colour-Coded Digital Evidence Results

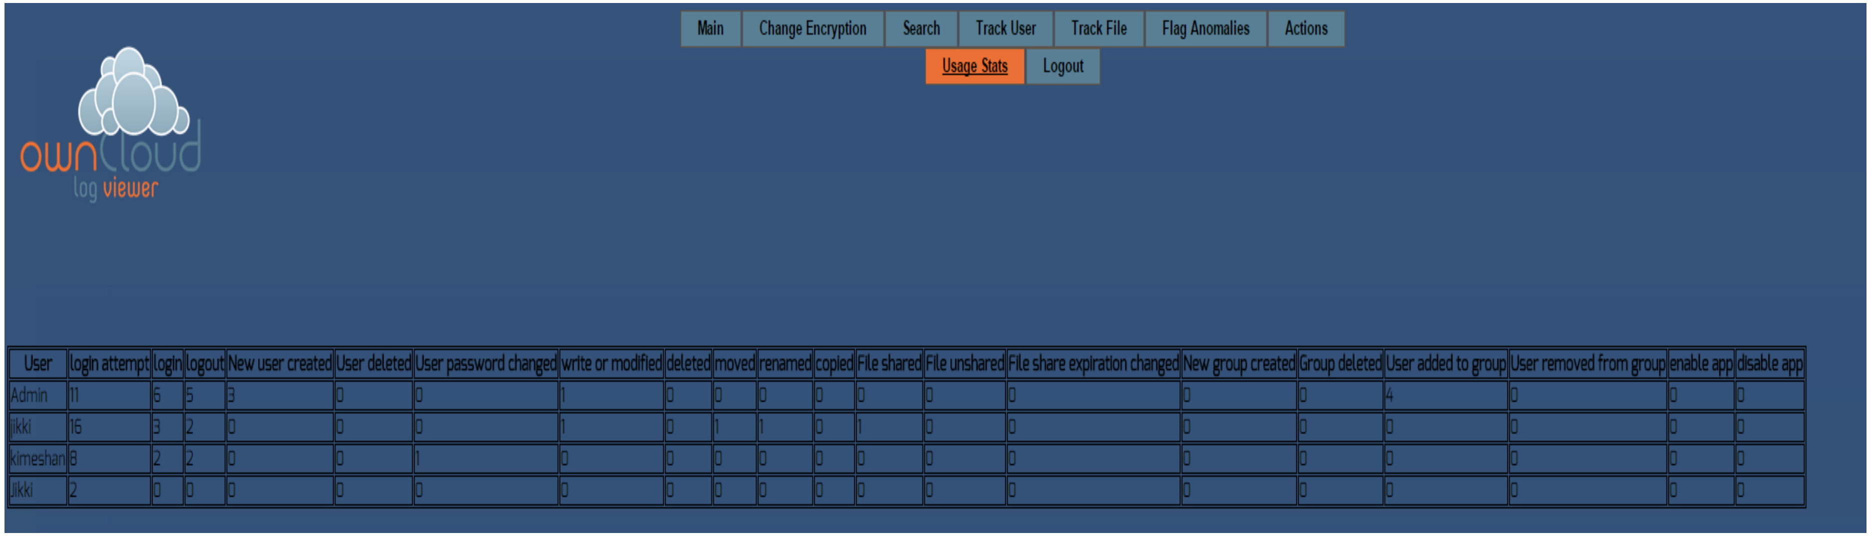

User	login attempt	login	logout	New user created	User deleted	User password changed	write or modified	deleted	moved	renamed	copied	File shared	File unshared	File share expiration changed	New group created	Group deleted	User added to group	User removed from group	enable app	disable app
Admin	11	6	5	3	0	0	1	0	0	0	0	0	0	0	0	0	4	0	0	0
jikki	16	3	2	0	0	0	1	0	1	1	0	1	0	0	0	0	0	0	0	0
kimeshan	8	2	2	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Jikki	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 9.8:- Screenshot for a Snapshot Overview Statistics of Usage

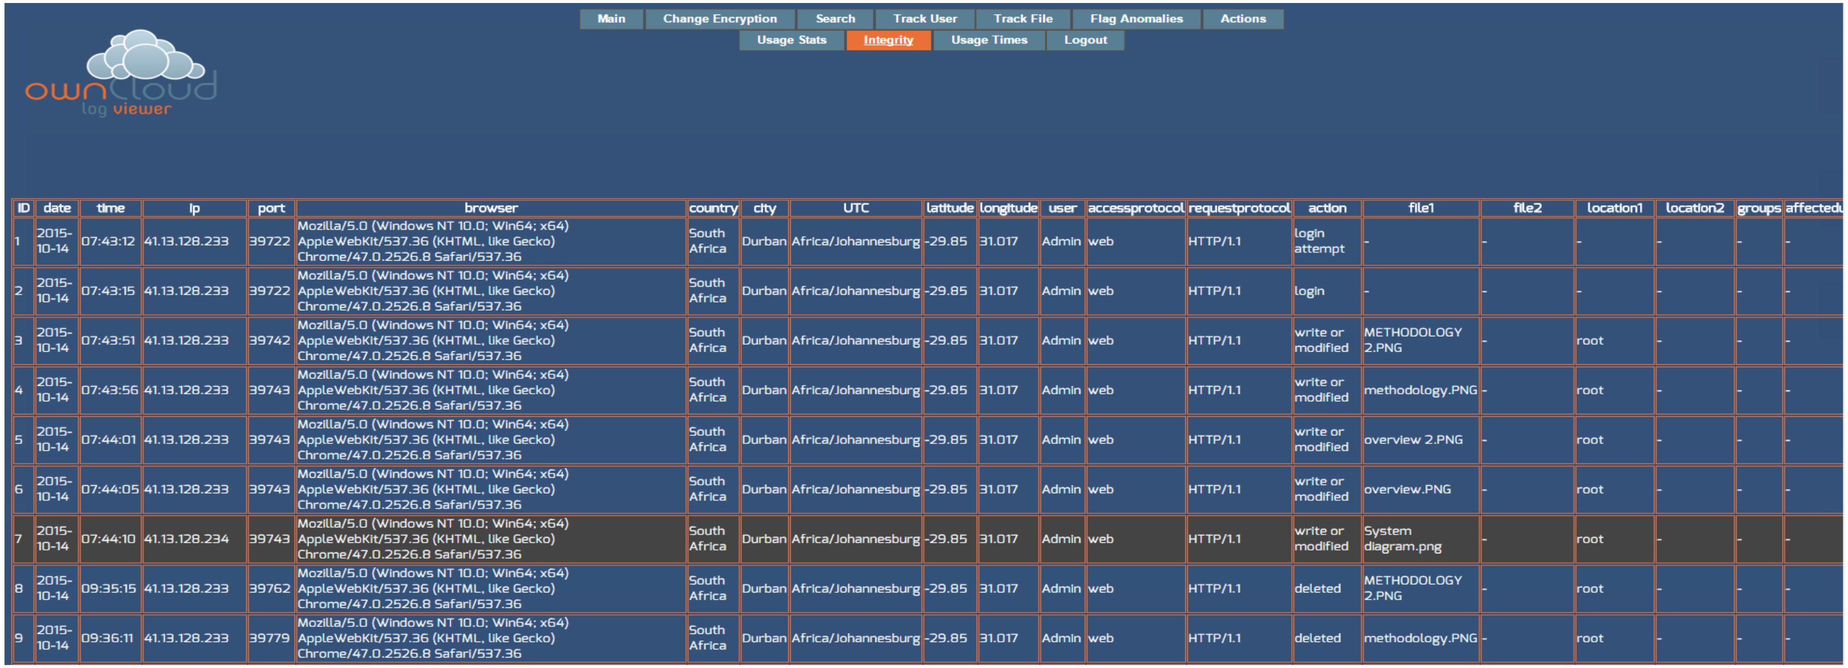

ID	date	time	ip	port	browser	country	city	UTC	latitude	longitude	user	accessprotocol	requestprotocol	action	file1	file2	location1	location2	groups	affectedu
1	2015-10-14	07:43:12	41.13.128.233	39722	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.8 Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.85	31.017	Admin	web	HTTP/1.1	login attempt	-	-	-	-	-	-
2	2015-10-14	07:43:15	41.13.128.233	39722	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.8 Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.85	31.017	Admin	web	HTTP/1.1	login	-	-	-	-	-	-
3	2015-10-14	07:43:51	41.13.128.233	39742	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.8 Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.85	31.017	Admin	web	HTTP/1.1	write or modified	METHODOLOGY 2.PNG	-	root	-	-	-
4	2015-10-14	07:43:56	41.13.128.233	39743	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.8 Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.85	31.017	Admin	web	HTTP/1.1	write or modified	methodology.PNG	-	root	-	-	-
5	2015-10-14	07:44:01	41.13.128.233	39743	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.8 Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.85	31.017	Admin	web	HTTP/1.1	write or modified	overview 2.PNG	-	root	-	-	-
6	2015-10-14	07:44:05	41.13.128.233	39743	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.8 Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.85	31.017	Admin	web	HTTP/1.1	write or modified	overview.PNG	-	root	-	-	-
7	2015-10-14	07:44:10	41.13.128.234	39743	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.8 Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.85	31.017	Admin	web	HTTP/1.1	write or modified	System diagram.png	-	root	-	-	-
8	2015-10-14	09:35:15	41.13.128.233	39762	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.8 Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.85	31.017	Admin	web	HTTP/1.1	deleted	METHODOLOGY 2.PNG	-	root	-	-	-
9	2015-10-14	09:36:11	41.13.128.233	39779	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.8 Safari/537.36	South Africa	Durban	Africa/Johannesburg	-29.85	31.017	Admin	web	HTTP/1.1	deleted	methodology.PNG	-	root	-	-	-

Figure 9.9:- A Screenshot to Determine the Integrity of Digital Evidence

9.3 CONCLUSION

This chapter tackled the digital forensic readiness aspect of cloud computing from the perspective of a proactive collection and preservation of digital evidence in anticipation of a legal lawsuit or policy breach investigation. The main contribution of Chapter 9 is the presentation and discussion of the three-tier digital forensic readiness model. The presented model aims to reduce the time it takes to conduct a digital forensic investigation in the cloud. Furthermore, Chapter 9 ensures that digital forensic investigations are carried out with minimal disruptions to the infrastructure or business processes of the cloud service provider. The model also notifies digital investigators of security threats before they occur in order to prevent them from happening in the first place. This introduces the proactive part of digital forensics and is a move beyond its reactive nature (i.e. acting on security threats after they have already occurred).

The proposed model monitors and proactively captures all user activity in a forensically sound manner. It helps to prepare the cloud environment to become ready for digital forensic investigations prior to an investigation. The proposed model has been designed and implemented to meet all the requirements (i.e. D_REQ 03.1 – 03.6) as listed in Chapter 5 and covers all the parts (i.e. regulatory requirements, proactive incident detection, monitoring and alerts, secure digital evidence handling and accountability) of the DFR component in the overall solution architecture as presented in Chapter 6.

The next chapter (i.e. Chapter 10) evaluates each of the individual models discussed in chapters 7, 8 and 9 respectively.

CHAPTER 10 EVALUATION AND DISCUSSION OF RESULTS

10.1 INTRODUCTION

Chapters 7, 8 and 9 provided a discussion of each component and its implementation of the overall solution. In particular, Chapter 7 focused on addressing the virtual machine placement problem by proposing a conflict-aware VM placement model. Chapter 8 focused on addressing the challenge of inadequate or weak authentication by proposing a strong risk-based multifactor authentication. Finally, Chapter 9 focused on addressing delays in digital investigations by proposing the 3-tier digital forensic readiness component for the monitoring and proactive collection of digital evidence from the cloud.

Chapter 10 now provides an evaluation and discussion of the results of each component of the solution. The main aim of this chapter is to validate the proposed solution in order to determine its viability. Furthermore, it aims to discuss the implications of the results of each component. This chapter is structured as follows: Section 11.2 provides an evaluation and discussion of the results of the conflict-aware VM placement component. Section 11.3 focuses on the evaluation and discussion of the results of the risk-based MFA component. Section 11.4 evaluates and discusses the results of the digital forensic readiness component. Section 11.5 offers an overall evaluation, which sums up the results discussion before Section 11.6 concludes the chapter and gives an overview of the next chapter.

10.2 EVALUATION AND DISCUSSION OF THE CONFLICT-AWARE VM PLACEMENT COMPONENT

Unlike most studies that only evaluate their VM placement algorithms on a simulated environment; the experiments in the case of this thesis were conducted on both a simulated and live cloud infrastructure. The choice of the simulation and live cloud environment is based on how well these two matched the proposed cloud structure, i.e.

$$VM \subseteq PN \subseteq Clu \subseteq DC \subseteq Loc$$

This section is therefore divided into two subsections: the first subsection discusses the results of the CBVMP model experimentation using CloudSim. CloudSim is a framework for modelling and simulation of cloud computing and evaluation of resource provisioning algorithms (Calheiros et al., 2010; Goyal, Singh and Agrawal, 2012; Mehmi, Verma and Sangal, 2017). The choice of CloudSim is based on its wide adoption and flexibility to allow customisation. The second subsection discusses the experimentation of the model on an OpenNebula cloud infrastructure. The decision to choose OpenNebula over other cloud infrastructures such as OpenStack, Eucalyptus and Nimbus is based on the fact that OpenNebula supports the IaaS structure that is proposed in Figure 10.1.

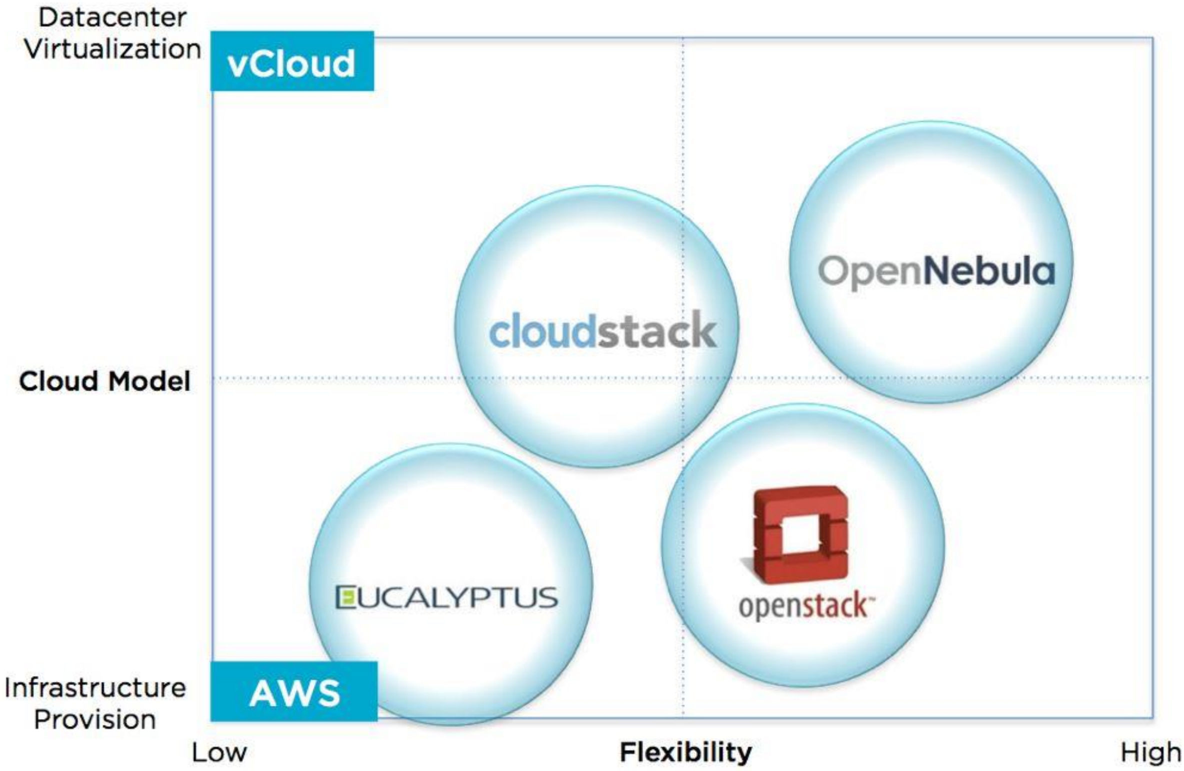

Figure 10.1:- A Comparison of Open-source Cloud Management Platforms (Llorente, 2013)

The VM placement strategy in this thesis requires a hierarchical structure to the underlying cloud infrastructure. This means that for some other cloud infrastructure, it would require a complete change to the underlying infrastructure before the solution could be applied. Thus, most of the existing cloud infrastructures were rendered unsuitable for experimentation. Furthermore, the choice of OpenNebula as cloud infrastructure is motivated by the comparison and analysis in Llorente (2013) and as depicted in Figure 10.1. Llorente (2013) takes four main open-source players in the cloud ecosystem, namely Eucalyptuss, CloudStack, OpenStack and OpenNebula.

Figure 10.1 shows that the open-source cloud platforms are clustered between the two proprietary ones, i.e. Amazon Web Services (AWS – bottom corner) and VMware (vCloud – top corner). AWS and vCloud are considered the extreme cases between which lie all four open-source models. These are analysed and compared based on their ability to adapt to data centre virtualisation that can be customised to provide differentiated cloud services. The results (Llorente, 2013) are in favour of OpenNebula, which demonstrates a high level of flexibility and high virtualisation of cloud services. Hence, OpenNebula is a good choice for this research.

10.2.1 Evaluation and Discussion of CloudSim Experimental Results

A number of experiments were conducted to test and evaluate the effectiveness and efficiency of the algorithms of the proposed conflict-aware VM placement model. The author defined a set of four VM instances, i.e. small, medium, large and extra-large, for both CloudSim and OpenNebula cloud. This is similar to the naming convention in Amazon Web Services IaaS. Table 10.1 shows the corresponding number of VM instances used in each of the three experiments on CloudSim, with a *CTL* of 0 for all VM placements. For example, Experiment 1 comprised of 10 small, 15 medium, 30 large and 50 extra-large VM instances. Experiment 2 comprised of 15 small, 30 medium, 50 large and 75 extra-large VM instances. Experiment 3 comprised of 30 small, 75 medium, 115 large and 150 extra-large VM instances.

Table 10. 1:- VM Instances in CloudSim

	VM Instance Classes			
Experiment No.	**Small**	**Medium**	**Large**	**Extra Large**
1	10	15	30	50
2	15	30	50	75
3	30	75	115	150

The specification of the different VM instances is shown in Table 10.2. The table shows that small VM instances are 2 GB, medium VM instances are 5 GB, large VM instances are 10 GB and extra-large instances are 30 GB. The RAM is 128, 512, 768 and 1024 MB for small, medium, large and extra-large VM instances respectively. The results of running these experiments are shown in Figure 10.2.

Table 10. 2:- Specification of VM Instances in CloudSim

	Specification of VM Instance Classes in CloudSim			
	Small	**Medium**	**Large**	**Extra Large**
Storage (GB)	2	5	10	30
RAM (MB)	128	512	768	1024
No. of CPU	1	1	1	1
No. of vCPU	1	1	1	1

They show the time it took to place each of the VM instances in CloudSim. For all VM instances, the CTL was kept at zero, which means they were placed wherever there was sufficient space.

Figure 10.2 reflects on the results of the three experiments. Experiment 1 shows that it took about 2.5 seconds to place 10 small VM instances, approximately 6 seconds to place 15 medium VM instances, approximately 12 seconds to place 30 large VM instances and 15 seconds to place 50 extra-large VM instances.

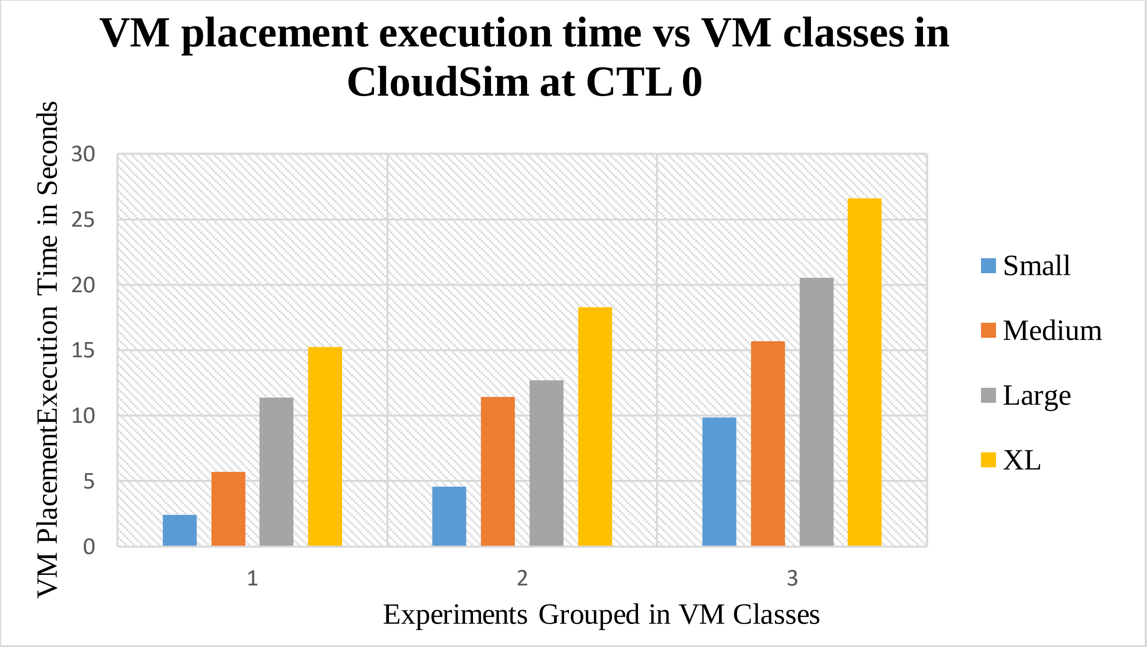

Figure 10.2:- VM Placement Performance across the Different VM Classes in CloudSim

Experiment 2 reflects approximately 5 seconds taken to place 15 small instances, approximately 12 seconds taken to place 30 medium VM instances, approximately 13 seconds to place 50 large VM instances, and approximately 19 seconds to place 75 extra-large VM instances. If we consider Experiment 3, it took approximately 10 seconds to place 30 small VM instances, approximately 16 seconds to place medium VM instances and approximately 21 seconds to place large VM instances. At the other end of the spectrum, extra-large VM instances took more time to place. For instance, it took about 27 seconds to place 150 extra-large instances.

Given the constant CTL in all three experiments, the results reflect a linear increase across the different types of VM instances. Furthermore, small VM instances appeared to take less time to be placed, compared to the other classes of instances. The author consequently deduced that the number and size of the VM instances had an effect on the VM placement performance. At this point it is important to note that the increase in placement times related to the number and size of VM instances. However, in Experiment 2, a deviation occurred from the normal linear increase. This is where it took about 13 seconds to place 50 large instances, which is almost the same time it took to place 30 medium instances in the same experiment. This is one of the rare occurrences which might be attributed to some external factor that could not be determined.

Figures 10.3, 10.4 and 10.5 reflect on the results of re-running the experiment with different *CTL*s ranging from zero to four and an increasing number of VMs.

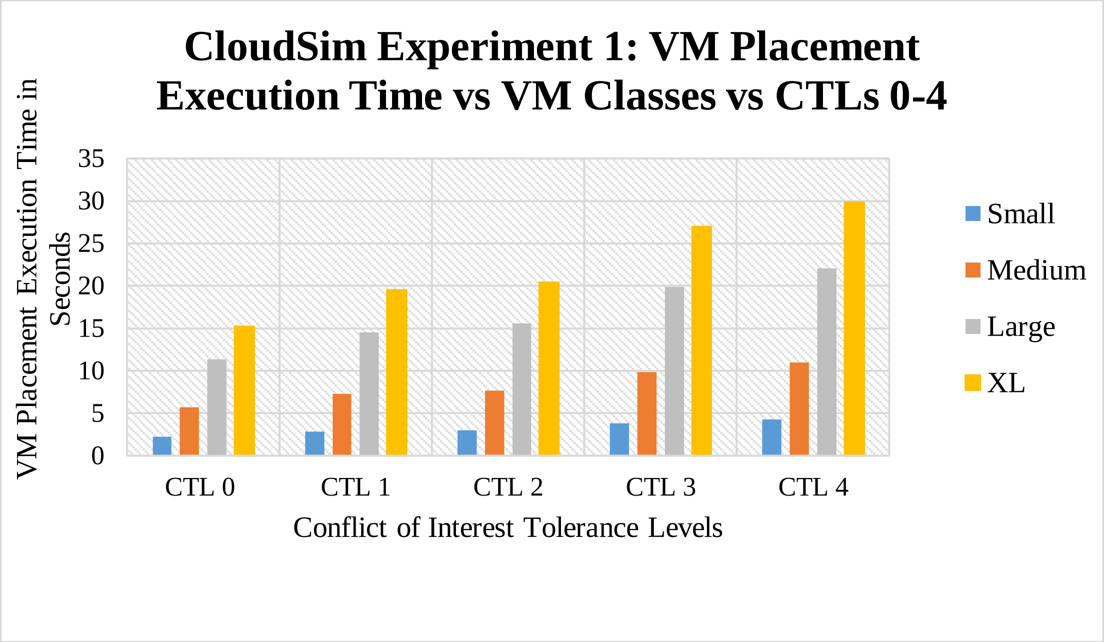

Figure 10.3:- Experiment 1: Results of VM Placement Execution Time against Different CTLs in CloudSim

From these three figures it can be deduced that the VM placement execution time increased linearly with a linear increase in *CTL*. For an example, an extra-large VM instance of CTL zero took approximately half the time to place, compared to an exact same size with a CTL of 4. This means that it took more time to place VMs of high *CTL* than those of a lower one. In general, the results illustrate that the CTL of a VM had a direct impact on the time it took to place a VM in CloudSim. The higher the CTL of the VM instance to be placed, the more time it took to do the actual placement.

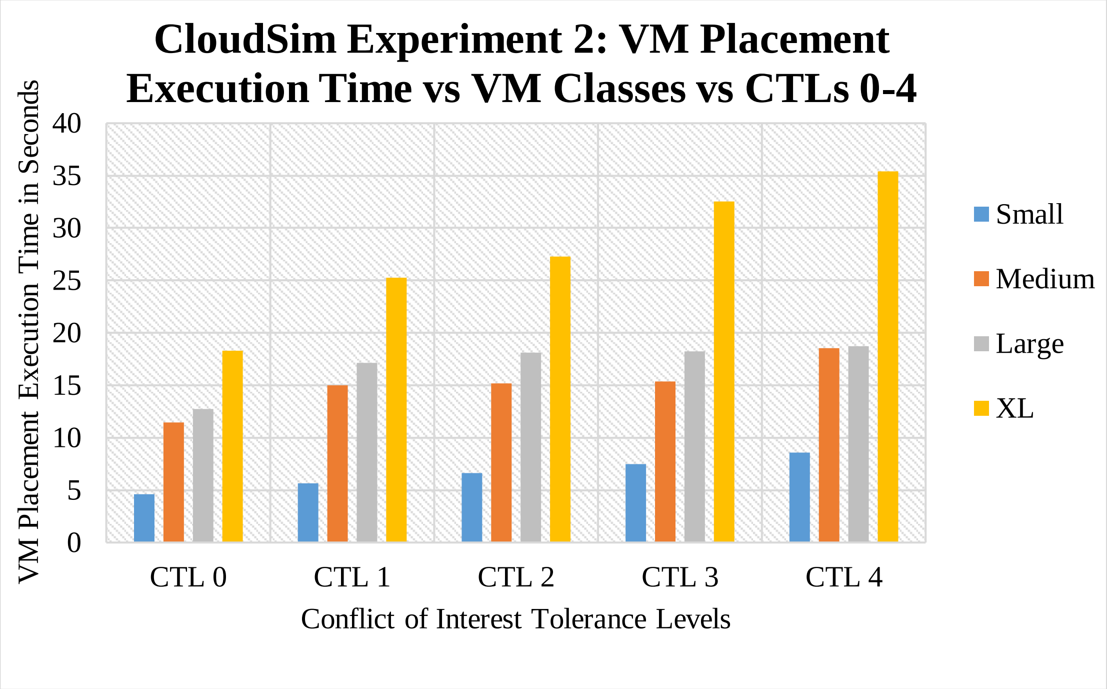

Figure 10.4:- Experiment 2: Results of VM Placement Execution Time against Different CTLs in CloudSim

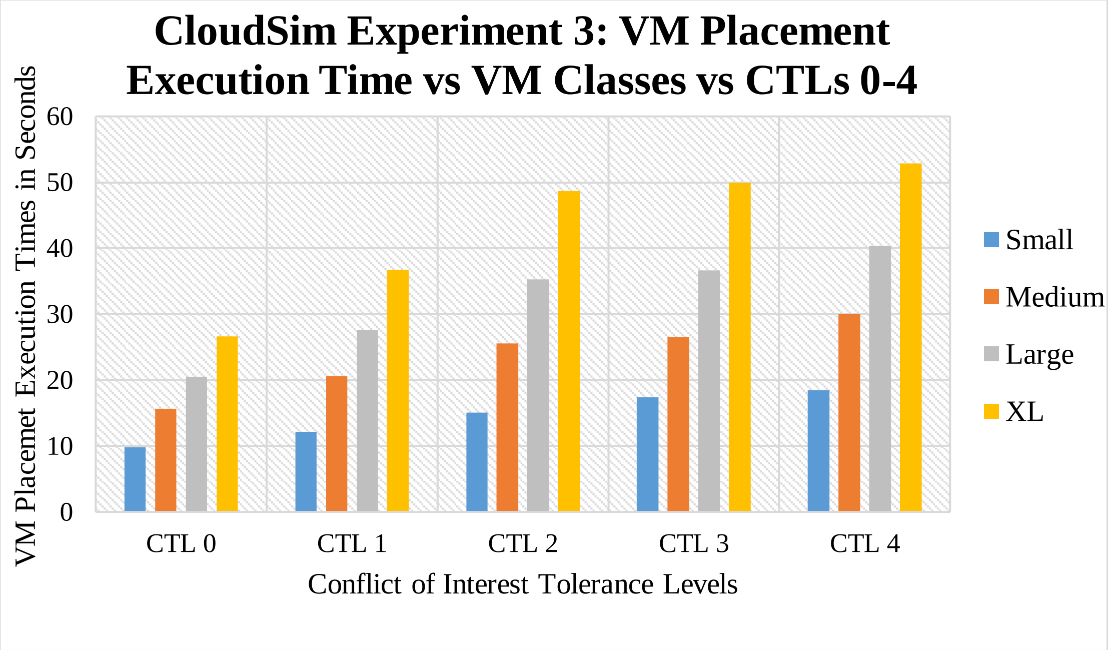

Figure 10.5:- Experiment 3: Results of VM Placement Execution Time against Different CTLs in CloudSim

Furthermore, it follows that the bigger the size of the VM instances that were to be placed, the more time it took to do the VM placement. Therefore, the author concludes that the introduction of the proposed conflict-aware VM placement (CBVMP) model introduced a time lag in the time it took to place VM instances in CloudSim. Based on the results, the author argues that it was better to allocate small VMs of zero CTL than large ones of higher CTLs. Furthermore, the author deduced that distributing tenants' data in many small VMs, even with higher CTLs, would also improve the resilience of the cloud, more especially when other hardware drives fail. The tenants would still be able to retrieve parts of their data elsewhere. However, this approach might have a negative impact on data retrieval. Retrieving data from several distributed small VMs might introduce time delays. This could be more of a challenge if the data needs to be assembled in a sequential order at the tenant's side. An inherent challenge to having many small VMs distributed across the cloud infrastructure is that with more tenants coming on board, it would become very difficult to find non-conflicting slots. However, this problem can be addressed by a conflict-aware load balancing (Zuikeviciute and Pedone, 2008; Amza, Cox and Zwaenepoel, 2003; Wang et al. 2018). Since conflict-aware load balancing falls outside the scope of this thesis, it is recommended as future work.

In concluding this section, the simulated conflict-aware VM placement provides good guidance on what to expect in a real cloud environment. However, cloud simulations cannot be relied upon for a true reflection of what could happen on the real cloud infrastructure. Therefore, this study also evaluates the proposed model on a real OpenNebula cloud infrastructure. The next section discusses the results of the conflict-aware VM placement thereof.

10.2.2 Evaluation and Discussion of OpenNebula Experimental Results

The following table shows the number of VM instances that were used in the live OpenNebula cloud infrastructure. For Experiment 1, there were 5 small, 5 medium, 10 large and 25 extra-large VM instances. Experiment 2 had 6 small, 15 medium, 20 large and 43 extra-large VM instances. Experiment 3 had 6 small, 12 medium, 34 large and 60 extra-large VM instances, and Experiment 4 had 8 small, 18 medium, 40 large and 65 extra-large VM instances. These numbers are summarised in Table 10.3.

Table 10. 3:- VM Instances in OpenNebula Cloud Platform

	VM Instance Classes in OpenNebula			
Experiment No.	**Small**	**Medium**	**Large**	**Extra Large**
1	5	5	10	25
2	6	15	20	43
3	6	12	34	60
4	8	18	40	65

In this setup, the number of VM instances are quite small compared to the experiments conducted in CloudSim. This is because in a live cloud environment, the number of instances is constrained by the size of the available physical infrastructure of the cloud. The host machines that were used for this experiment were limited to approximately 980 GB in total. Hence, the VM instances for OpenNebula are fewer than those used in CloudSim. However, this difference in VM instances is not expected to have so much of an effect on the generalisation of the results and the conclusions drawn.

Table 10. 4:- Specification of VM Instances in OpenNebula

	Specification of VM Instance Classes in CloudSim			
	Small	**Medium**	**Large**	**Extra Large**
Storage (GB)	1	2	4	8
RAM (MB)	128	512	768	1024
No. of CPU	1	1	1	1
No. of vCPU	1	1	1	1

Table 10.4 shows the specification of the VM instances that were used for the OpenNebula cloud platform. This table shows that a small VM instance is 1 GB, medium VM instance is 2 GB, a large VM instance is 4 GB and an extra-large VM instance is 8 GB. The RAM is 128, 512, 768 and 1024 MB for small, medium, large and extra-large VM instances respectively. Regardless of its size, each instance is allocated one CPU and one vCPU.

Figure 10.6 shows the results of placing the actual VMs at *CTL* zero on an OpenNebula cloud for the different sizes. Four runs were performed at CTL of zero, whereby VMs can be placed anywhere as long as there is sufficient storage capacity. In these four experiments, it took between 18 and 34 seconds to place the small VM instances. The medium VM instances took a turnaround time of between 32 and 57 seconds. Large instances took between 42 and 70 seconds and extra-large VMs were placed within 48 to 82 seconds. The results further show that the placement time increased with the size of the VMs. The results of Experiment 2 seemed to defy the odds, in that the time it took to place 10 large VM instances was almost the same as that of placing 25 extra-large VMs. This is despite the fact that in all four experiments, all VM instances were at CTL of zero, which does not cater for any conflicts. The only variable that changed was the number of VM instances.

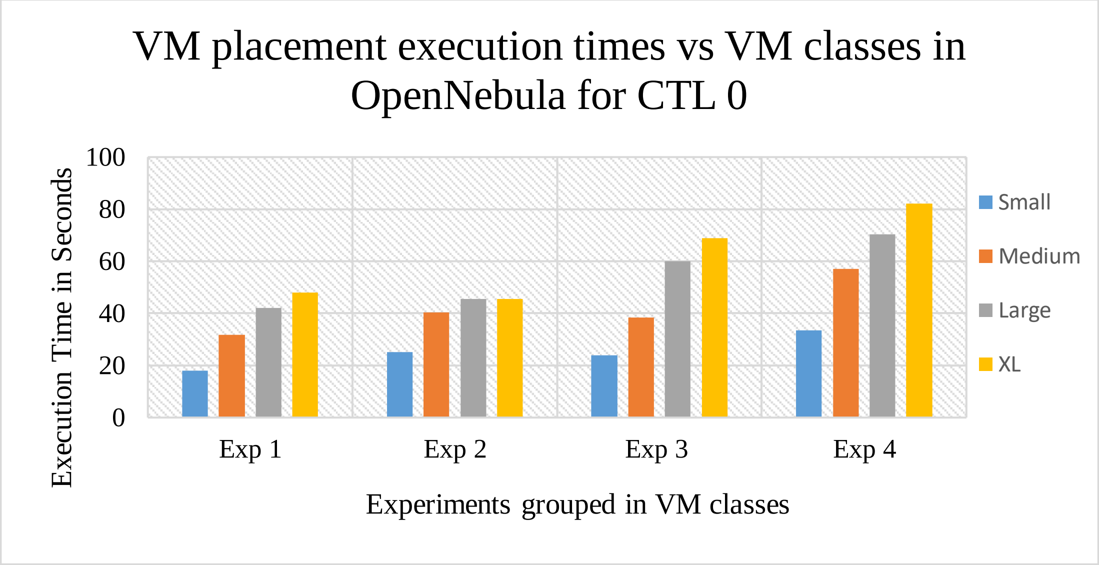

Figure 10.6:- VM Placement Performance across the Different VM Classes in OpenNebula at CTL 0

However, it can be postulated that this deviation might be as a result of large VM instances being placed with different physical hosts and the fact that most of the extra-large VM instances were placed on the same physical host. However, this is strange, and the deviation becomes the odd one out, even when one looks at the next results in Figure 10.7. Furthermore, the results in this experiment (Figure 10.6) show that decreasing the number of VM instances also reduces the time it takes to do the actual placement. This is demonstrated in Experiments 2 and 3 where the number of VM medium instances are reduced from 15 to 12 and the VM placement turnaround time also gets reduced from approximately 40 to 38 seconds. Moreover, there is a peculiar result between Experiments 2 and 3 for the small VM instances. The number of VMs that are being placed is the same, namely 6, but the results in placement execution time differ. This is a result that requires further research to explain the noted deviation.

Figure 10.7 shows the results of placing VMs at CTL zero to four. The small VM instances were placed between 18 and 35 seconds, while the medium VM instances required a turnaround time of between 31 and 45 seconds. Large instances took between 42 and 66 seconds and extra-large VMs were placed within 48 to 68 seconds. As demonstrated in Figure 10.7, the results in Figure 10.8 also show that the placement time increased with a linear increase in the CTL and an exponential increase in the sizes of the VMs.

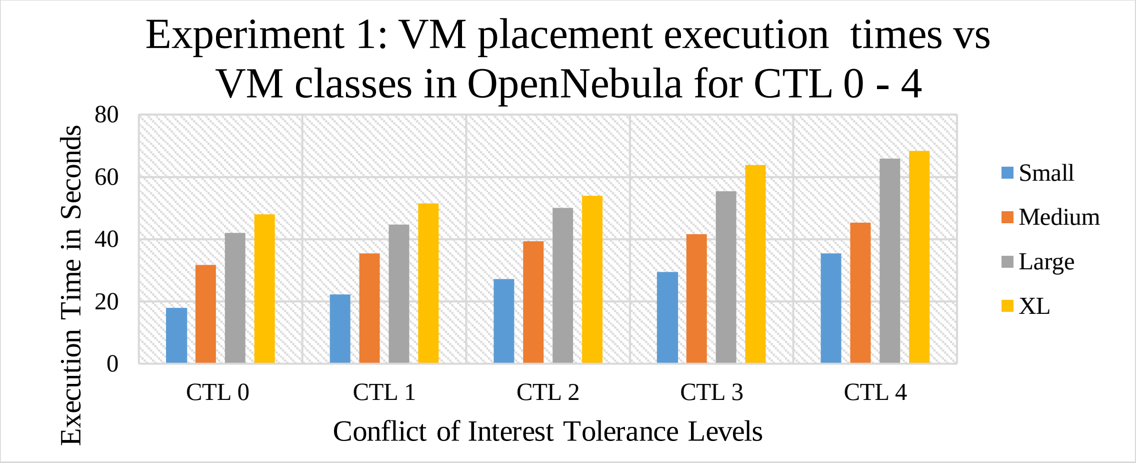

Figure 10.7:- Experiment 1: VM Placement Performance across the Different VM Classes in OpenNebula at CTL 0 – 4

Figure 10.8 also shows that small VM instances took between 25 and 55 seconds to be placed. Medium VM instances required a turnaround time of between 40 and 59 seconds. Large instances took between 47 and 66 seconds and extra-large VMs were placed within 45 to 79 seconds. These results reflect an insignificant increase between CTL zero and one for medium and large VM instances. However, for small and extra-large instances the difference was more significant. At CTL two, the difference in turnaround time between medium and large was very small, with medium instances taking more time than large instances. This is another unusual result that requires further analysis to correctly determine its cause.

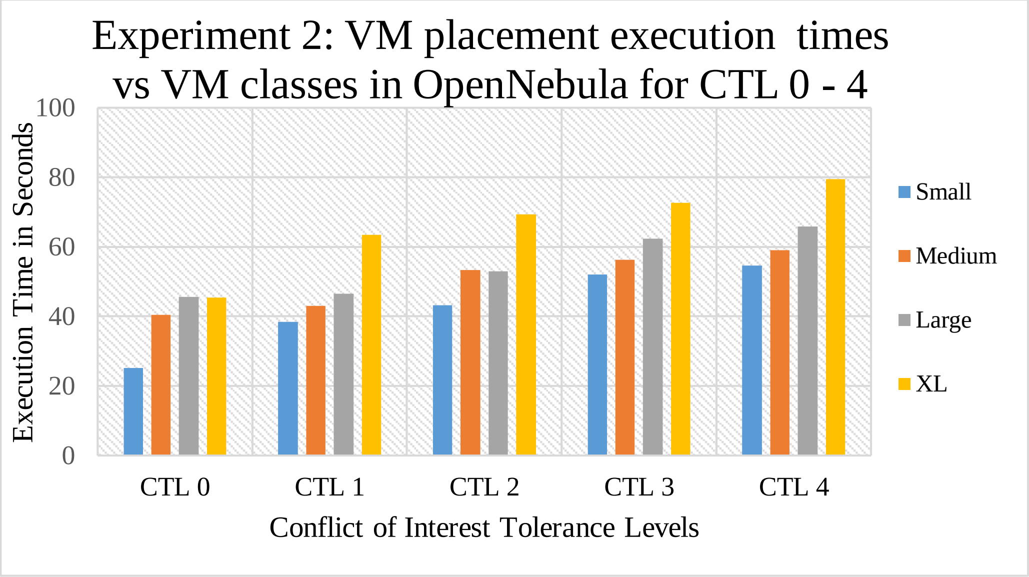

Figure 10.8:- Experiment 2: VM Placement Performance across the Different VM Classes in OpenNebula at CTL 0 – 4

Figure 10.9 shows that small VM instances were placed within a timeframe of 24 to 43 seconds. Medium VM instances were placed from a minimum of 24 to a maximum of 62 seconds. Large VM instances took 60 seconds minimum and 73 seconds maximum. Lastly, placing extra-large VM instances took 69 to 87 seconds. The results of Figure 10.9 reflect a relatively linear increase right across all VM instance classes (i.e. small, medium, large and extra-large) and the different CTLs.

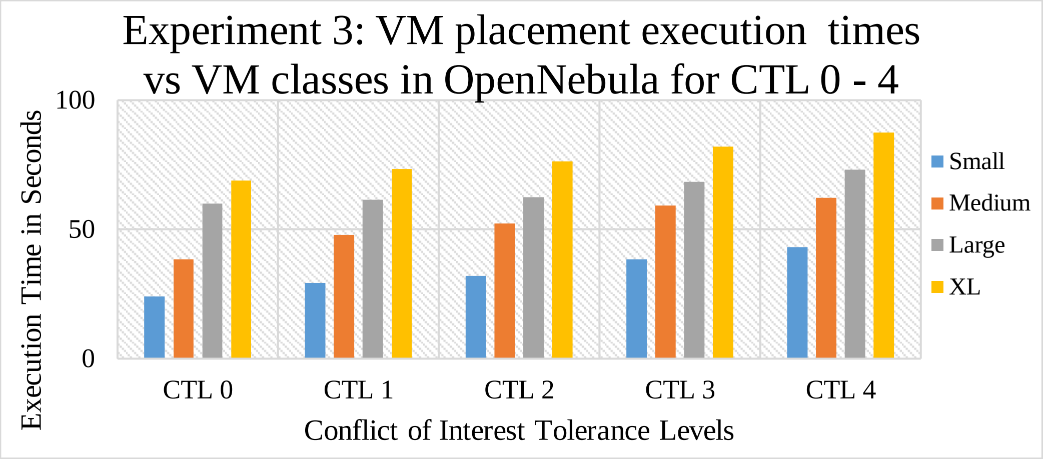

Figure 10.9:- Experiment 3: VM Placement Performance across the Different VM Classes in OpenNebula at CTL 0 – 4

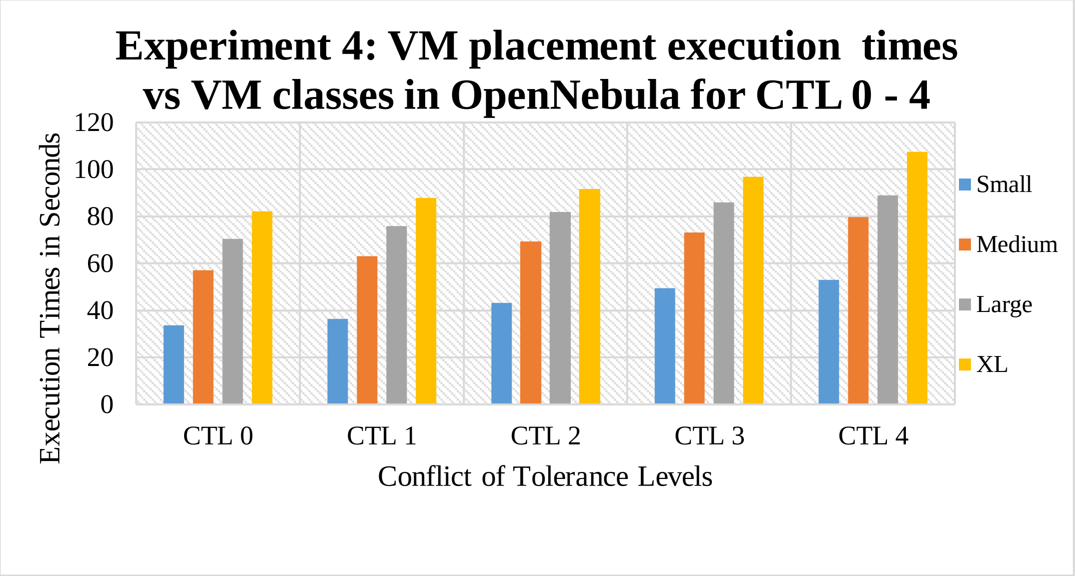

Figure 10.10:- Experiment 4: VM Placement Performance across the Different VM Classes in OpenNebula at CTL 0 - 4

Figure 10.10 shows that small VM instances were placed within a timeframe of 34 to 53 seconds. Medium VM instances were placed from a minimum of 57 to a maximum of 80 seconds. Large VM instances took 70 seconds minimum and 89 seconds maximum. Lastly, the placement of extra-large VM instances required 82 to 108 seconds. The results of Figure

10.10 are similar to those of Figure 10.9. Both figures reflect a relatively linear increase right across all VM instance classes and the different CTLs.

In summary and without any loss of generality, one can deduce from the results that the higher the CTL of VM instance, the longer it took to do the actual placement. Moreover, one can also deduce that the size of the VM instances had a direct effect on the time it took to do the actual placement. Furthermore, when doing the actual placements of VMs, it was noted that the VM instance templates were created on the front-node and had to be transferred via the network for deployment to the respective physical nodes. This adds some latency to the response time. Hence, it would be better if the VM instances were created locally on the nodes where they were to be placed. In comparison, comparing the results of an actual VM placement on an OpenNebula cloud environment to those of CloudSim (simulated environment) clearly shows that a simulated environment produces results with a significant error margin. For example, CloudSim placed 150 instances of extra-large VMs in less than 30 seconds, yet doing the actual placement on OpenNebula with less than half the number of instances took more than 80 seconds – almost three times the time it took to do a similar placement in CloudSim. Hence, it can be argued that although a simulated environment is fine to show an estimated glimpse of how much time it would take to place instances, it is nowhere close to the actual VM placement on an existing cloud. Having done the experiments in an OpenNebula cloud, it would be interesting to compare the results to other cloud computing environments such as CloudStack, Eucalyptus, and OpenStack. However, such a comparison is left as future work.

The next section discusses the results of the strong risk-based MFA component.

10.3 DISCUSSION OF THE RESULTS OF THE RISK-BASED MFA COMPONENT

This section discusses the evaluation of the strong risk-based MFA component of the solution. Emphasis is placed on evaluating the effectiveness and efficiency of security controls to secure login data in storage and processing. The login data includes usernames, passwords, geo-locations, OTPs, OOB tokens, SecurityPhrase and device identifiers. Furthermore, this section deals with the response time (i.e. performance) of the entire authentication component.

10.3.1 Evaluation and Discussion of the Results of the Cryptographic Element

Figure 10.11 illustrates eleven iterations of encrypting and decrypting login data (i.e. 10496-bit stream) and its response times.

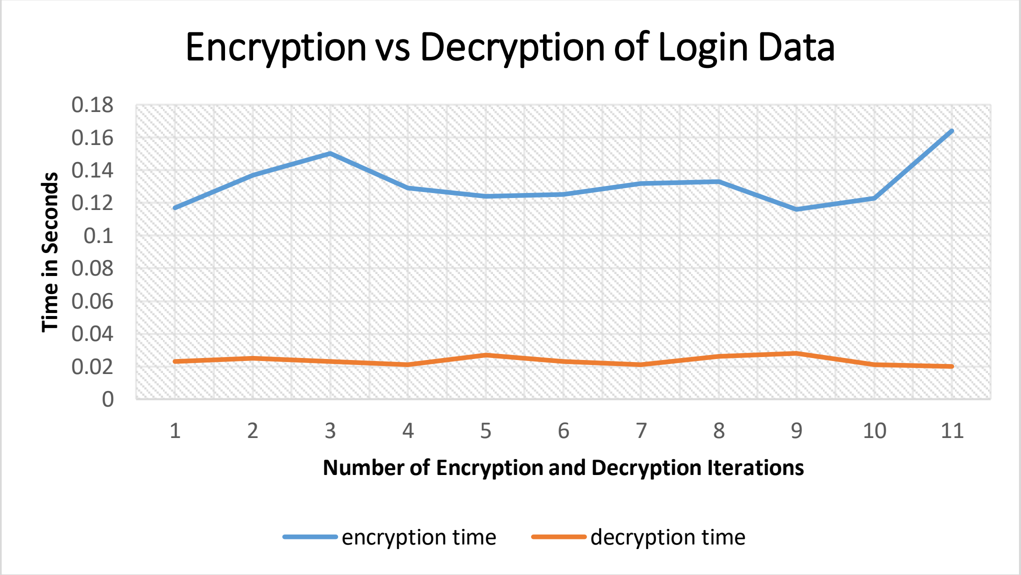

Figure 10.11:- Performance of the Cryptography Algorithm on Encrypting and Decrypting Login Data

The graph shows that both encryption and decryption of login data happen in a small fraction of a second. For an example, encryption takes approximately 0.12 to 0.16 seconds and decryption takes approximately 0.02 to 0.025 seconds. This is an indication that the cryptography algorithm employed in this study is effective, efficient and works without getting in the way of users. It also means that users would not even notice the effects of the encryption and decryption on their usual login process. This result is key, especially when one considers the on-going discussions around balancing security and usability of security mechanisms (Adams and Sasse, 1999; Shen, Yang and Zhou, 2018; Alashwali and

Rasmussen, 2018). Furthermore, the results show that encryption takes longer than decryption. Although there is a significant time difference, the effects of encryption would not have a negative impact on the login process. This is because the encryption process happens in the background after a user has already been authenticated, and therefore it has no effect on the login turnaround times. The good thing about the on-the-fly decryption of login data is that it ensures a good login turnaround time. Hence, to have it as low as it currently is, is good. Moreover, the results do not show any significant correlation between the encryption and decryption response time. However, it does appear that when the encryption's response time increases, the decryption's response time decreases at a different rate. For example, from iteration one to three, the encryption response time is increasing, yet the decryption response time is decreasing. The same can be noticed from iterations nine to eleven.

Furthermore, the ciphertext is based on the entire ASCII table with 128 characters. This is beyond the normal 26-character set and nine numeric characters that are found in most cryptographic systems. The 128 ASCII characters make brute-forcing our solution to be computationally unfeasible. This is because it would require many more resources. Besides, the cryptographic solution based on the ASCII table is to a certain extent immune to frequency analysis attacks. This is mainly because some of the characters in the ASCII table are not recognisable. Moreover, the strength of the ciphertext comes from the noise that is added after encryption, which blends the key with the ciphertext. This is 'security through obscurity', where the solution depends on the attacker not being able to tell that the encryption key is embedded and hidden within the ciphertext. This element could be a vulnerability of the solution. Since there is still a need for a good key management system, future research will explore a keyless encryption option, where one encrypts, uses neural networks to learn the key, and then discards or destroys the key. This part is also left for future work and not part of this thesis.

The next subsection evaluates and discusses the results of the steganography element of the strong risk-based MFA component.

10.3.2 Evaluation and Discussion of the Results of the Steganography

This section evaluates and discusses the effectiveness and efficiency of the steganography aspect of the strong risk-based MFA solution. Steganography is used to embed both OTPs and OOB tokens for login attempts that are classified as high risk. This is done to secure OTPs and OOB tokens as they are transmitted to user login devices across sometimes vulnerable networks. Figure 10.12 depicts the response time to embed and decode OTPs and OOB tokens from a low-fidelity steganography image.

The figure shows that both embedding and decoding of OTPs and OOB tokens take less than a fraction of a second to complete. Embedding an OTP or OOB token in the images appears to take more time than decoding it out of the image. This is almost the same sequence as the encryption and decryption showed in Figure 10.11. Furthermore, the results show that embedding an OTP or OOB token varies between 0.3 to below 0.5 seconds, with an average of 0.3772 seconds. On the other hand, the process of decoding OTPS and OOB tokens seems to be consistent, with negligible variations. Both these add insignificant time delays to the turnaround time of the overall authentication process.

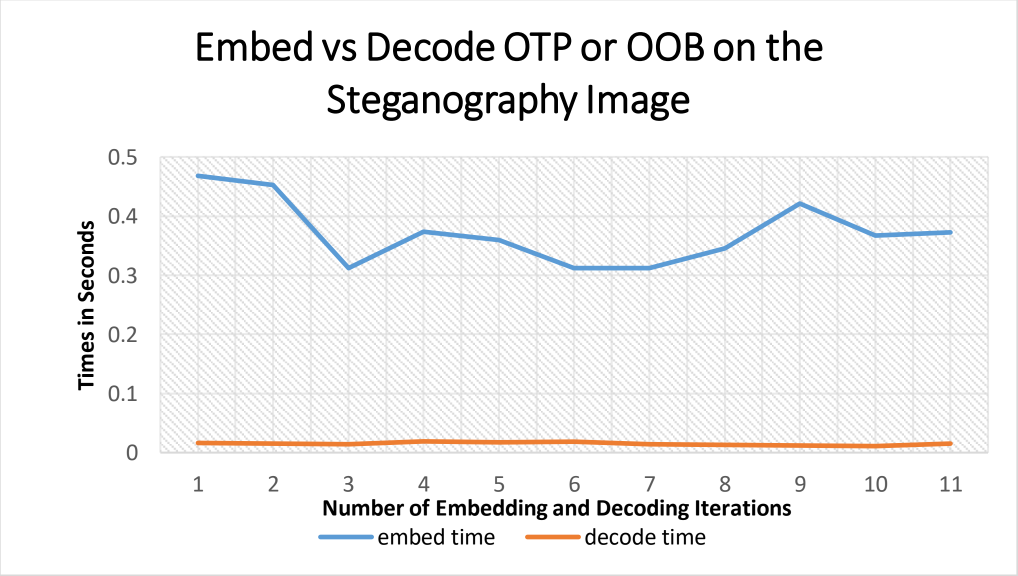

Figure 10.12 :- Performance of the Steganography Algorithm on Embedding and Decoding OTPs and OOB Tokens

It can be argued that embedding and decoding OTPs and OOB tokens happen on the fly and that it is seamlessly done in the background without disturbing the users. Since steganography was used effectively and efficiently to deliver OTPs and OOB tokens in this study, this has improved the authentication process without affecting the users. However, word processors have been used before by attackers to decode messages hidden in images (Ahvanooey et al., 2018; Kumari and Singh, 2018). In order to thwart such an attack, the author used the Least Significant Bit to change the first two bits of each and every pixel of the image. This is to ensure that word processor attacks would not be able to decode the text (i.e. OTPs or OOB tokens) from the image.

The next subsection evaluates and discusses the results of user login risk analysis turnaround times. This refers to the time it takes for the strong risk-based MFA to evaluate the risk and eventually authenticate a user. It excludes the time that it takes the user to enter their credentials, OTPs and OOB tokens from the input device. The focus is now placed on the automated processes after the user has clicked login, to the time when the system finally grants or denies them access. Figure 10.13 shows that the response time of risk analysis and login. This is also a fraction of a second. At worst, it takes approximately 0.37 seconds to evaluate the risk of a login instance. At best, it takes approximately 0.24 seconds. This confirms that the authentication of users happens on the fly.

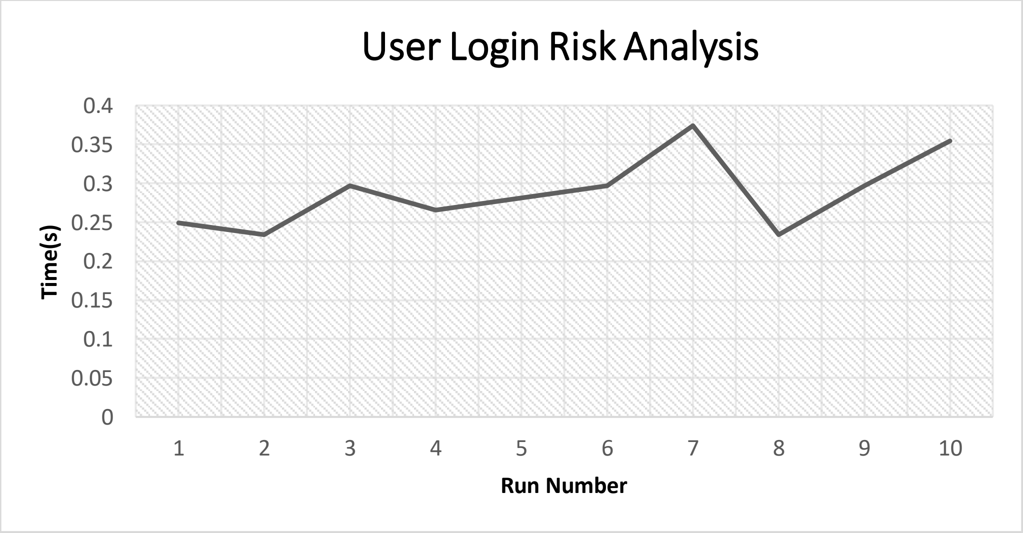

Figure 10.13:- User Login Risk Analysis

Furthermore, the keystroke behavioural analytics element of the risk-based MFA component was tested using session duration time and typing speed. This was to see if the authentication mechanism would be able to deny access to illegitimate users with authentic login credentials. This test followed after a process of training the system with login credentials from authentic and authorised users. The test was done with a sample size of ten users, which might not be significant enough to draw conclusive results. Hence, it would be necessary to get a bigger sample size in future work. However, this small sample size was sufficient to illustrate the point and demonstrate the practicability and effectiveness of the solution to deny access to unauthorised users with the correct credentials. It was interesting to note that all login attempts of the illegitimate users were flagged as either second-order or first-order outliers. This means all users that needed either an OTP or OOB token on top of the normal user credentials were correctly classified as suspicious logins. Despite the fact that the unauthorised users were making login attempts with correct credentials, they were all denied access.

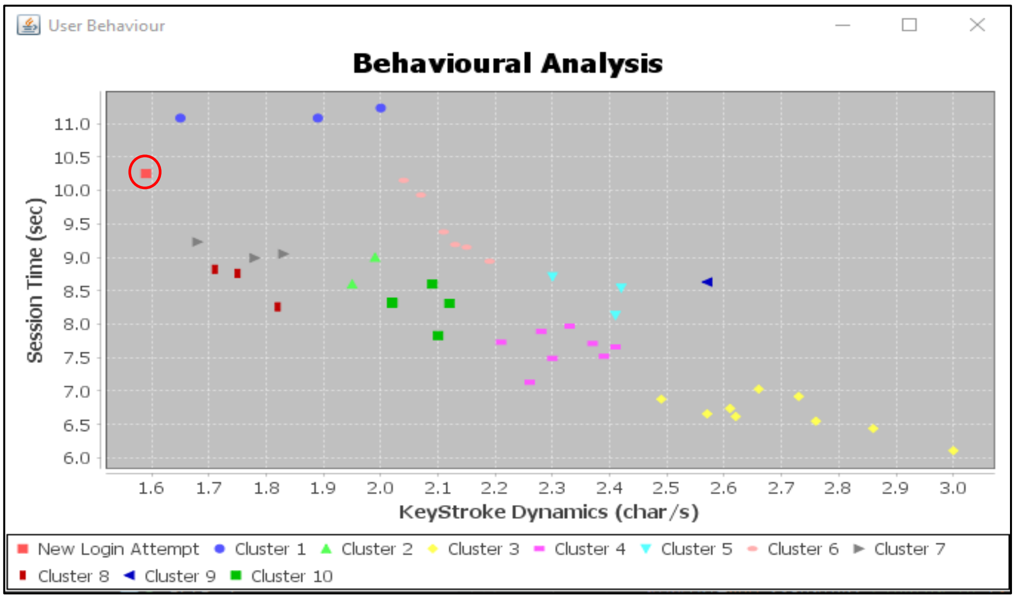

Figure 10.14 :- A Login Attempt that Fails a Local Anomaly Check and Prompts for an OTP

Figure 10.14 shows an example of a login that fails a local anomaly check. This is classified as a first-degree outlier and prompts the use of an OTP. In Figure 10.14, the cluster of context is cluster one (i.e. blue circular dots). Existing behavioural keystroke dynamic models would classify such an anomaly (which is not necessarily too far from the centroid of cluster one) as a normal pattern and grant access to the user. However, this is not the case with the proposed solution, as this is correctly classified as an outlier. Figure 10.15 shows an example of a login attempt that failed a global anomaly check and was classified as a second-degree outlier. The user was consequently prompted for an OOB token.

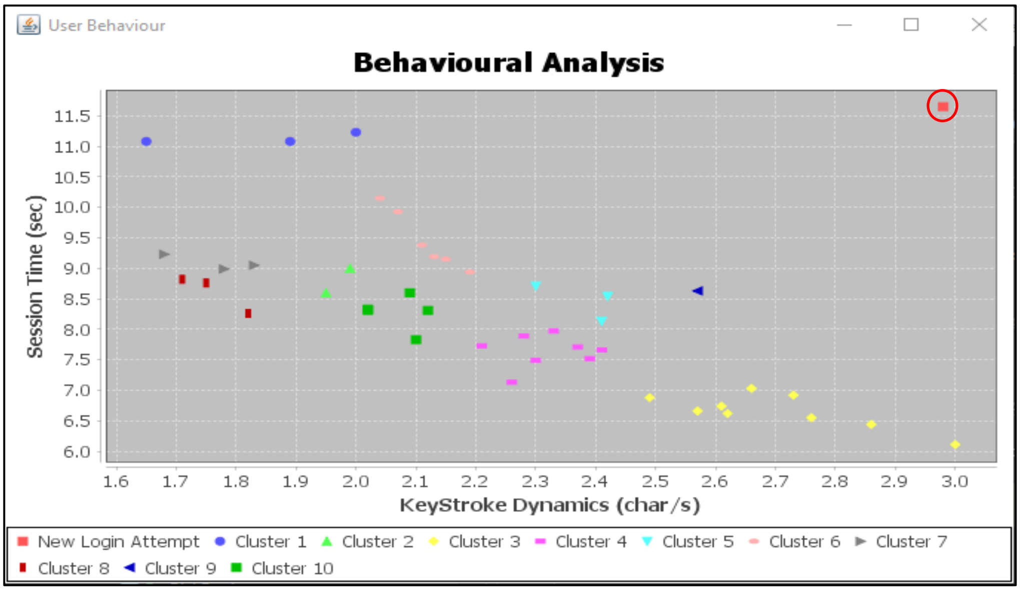

Figure 10.15:- A Login Attempt that Fails a Global Anomaly Check and Prompts for an OOB Token

The login attempt is located far from any of the ten login clusters, i.e. at the top right corner of the *x*-*y* plot. This login attempt is correctly classified as a second-degree outlier even though the user credentials are correct. This would require a user to enter an OOB token. Sixty per cent of the testing users failed the global anomaly check.

Figure 10.16 is a graph depicting the statistics of the different outliers based on the experiments.

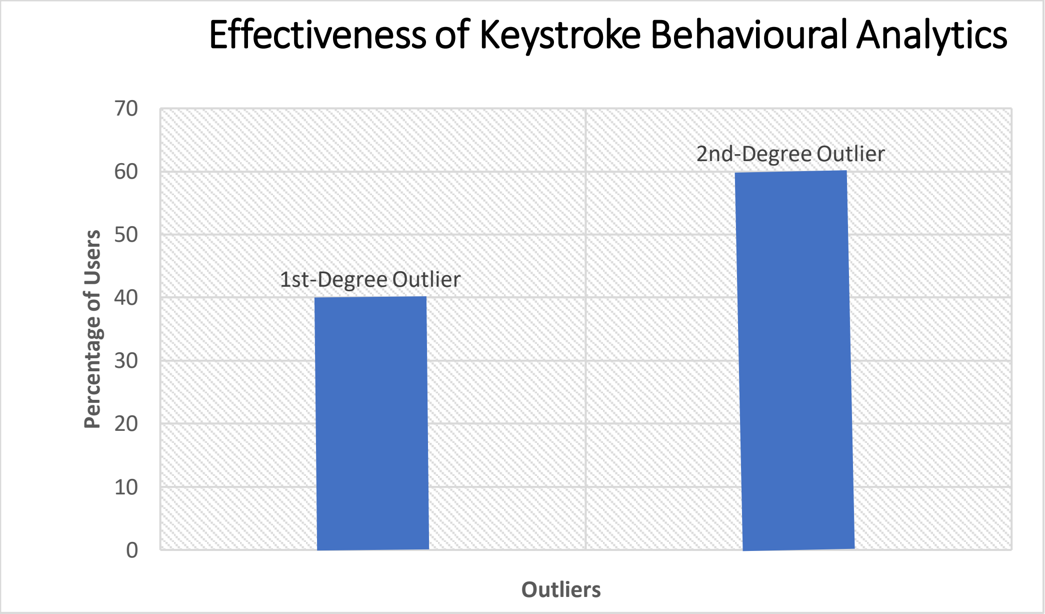

Figure 10.16 :- Percentage of Users' Login Attempts that Failed the Local vs Global Anomaly Check

Figure 10.16 shows that 40% of the users' login attempts failed the local anomaly check and they were therefore prompted for an OTP. It also shows that 60% of the login attempts failed the global anomaly check and were prompted for OOB tokens. This a good indication of the effectiveness of the solution. However, an ideal case would be to have all these test user login attempts classified as second-degree outliers.

Despite the current work's failure to reach an ideal case, it can be concluded that it yields a plausible solution to solve challenges posed by attackers using compromised user credentials. The results reflect that the strong risk-based MFA solution that is proposed in this thesis is effective and to a certain extent also efficient in preventing unauthorised users who are armed with legitimate credentials from gaining access to systems that they are not authorised to access.

The results also demonstrate the uniqueness of this approach, in that even minor deviations from the norm are classified as anomalies. However, there is room for improvement. For example, more training data with a larger sample size would cause the classification of login attempts to be very sensitive, in that the slightest deviation from the norm would result in a second-degree outlier. This could be done in future work.

The next section discusses the evaluation of the DFR component.

10.4 EVALUATION AND DISCUSSION OF RESULTS OF THE DFR COMPONENT

This section evaluates and discusses the results of the DFR component. The DFR component was tested for performance, correct flagging of incidents, preserving integrity and for ensuring the security of digital evidence on a private cloud environment called OwnCloud. The first test evaluated the effects of the DFR component on performance. This is important in the security domain because most security solutions are associated with time lags. Most users try to find ways to circumvent security solutions that are meant to protect their data and systems from attackers. Hence the DFR component was tested for its response time for capturing, transmitting, writing and deleting a file of 973 KB which was considered to be potential digital evidence. This test file was quite big when compared to normal sizes of log entries. It can be argued that this would be a worst-case scenario, but it is important to use a worst-case scenario to avoid under-estimation. Hence, this test provided a good indication of what to expect. Figure 10.17 shows the response time of some operations when the DFR component had not yet been added and when it was added to OwnCloud. It must be mentioned though, that the scenario is based on a one-to-one relationship between the remote server and the client device. A more realistic scenario would have been to have multiple clients doing the same operation and to do a stress test on how the server could scale with more traffic. However, this scenario was not covered as part of this thesis.

Figure 10.17 reflects that all the operations (capture, transmit, write and delete) on the potential digital evidence took almost double the time they would take when the DFR component was not added. This is an indication that the DFR adds a significant time lag to the normal operations. The expectation is that when the size of the digital evidence increases, the response time would also increase. Furthermore, it could be deduced that adding more clients would also increase the time it takes to do each of the operations, especially writing and transmitting the potential digital evidence to the remote server.

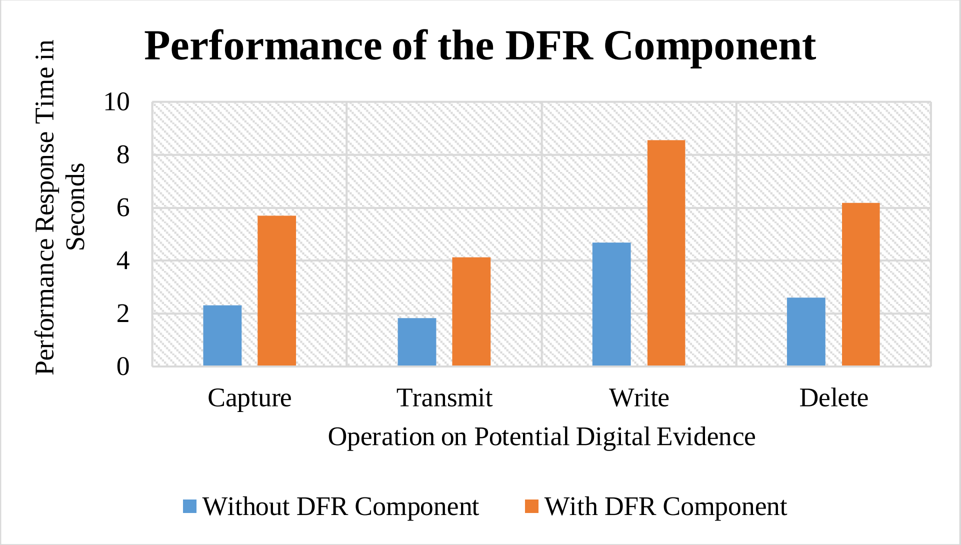

Figure 10.17:- The Performance of the DFR Component

However, this might not necessarily affect the evidence-capturing operation. This is mainly because each client has its own capturing element and adding more, does not affect existing elements. Multiple distributed remote servers would come as a solution to help solve the performance issues.

Moreover, the DFR component was evaluated in terms of its accuracy to correctly classify and flag anomalies. This exercise was conducted over a period of five weeks. Week 1 was used to collect the baseline users' activity data. In the second week, monitoring started to identify possible deviations from the baseline data collected in the first week. The focus of the monitoring was on unusual login location and time. This was linked to the anomalies raised in the strong risk-based MFA. However, on the DFR component, the idea was to identify false positives and the results are reported in Figure 10.18. According to this figure, the false positives start on a high and they decrease with time. This result indicates that the DFR component gets better with time as it is being used. Its accuracy continues to improve over time. For example, at the beginning of the experimentation, i.e. during Week 2, many login attempts were classified as anomalies, yet they were legitimate.

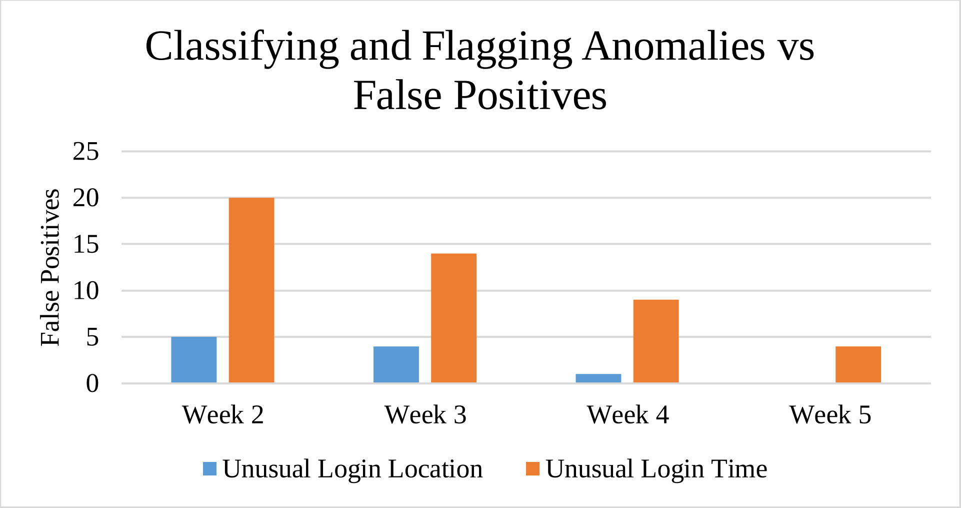

Figure 10.18:- Classification and Flagging of Anomalies against the Number of False Positives

By Week 5, the number of wrongly classified login attempts were reduced by 80%, with only 20% recorded false positives. Surely, in time the DFR component’s accuracy in terms of classification should improve even more.

Furthermore, the DFR component was evaluated on how it preserves the integrity of potential digital evidence. Since preserving the integrity of potential digital evidence has a direct impact on its admissibility in a court of law, it was critical to find the best hashing algorithm. The collected dummy data in the remote server was hashed using CRC 32, CRC 32b, MD5, SHA-1, SHA-256 and SHA-512. For each hashing algorithm, the results were recorded using a PHP script for the time it takes to hash. The remote server had a number of records of potential evidence. However, there were less than 100 records. Figure 10.19 contains a diagram showing the results of the hashing algorithms.

Figure 10.19 shows that CRC 32B is the best and SHA-512 is the worst hashing algorithm. Moreover, the CRC 32B hashing algorithm uses only eight characters as a hash. Compared to SHA-512, which requires 128 characters for each hash index, CRC 32B is far more cost-effective in terms of storage space. Furthermore, CRC 32B seems to be more sensitive than the other algorithms.

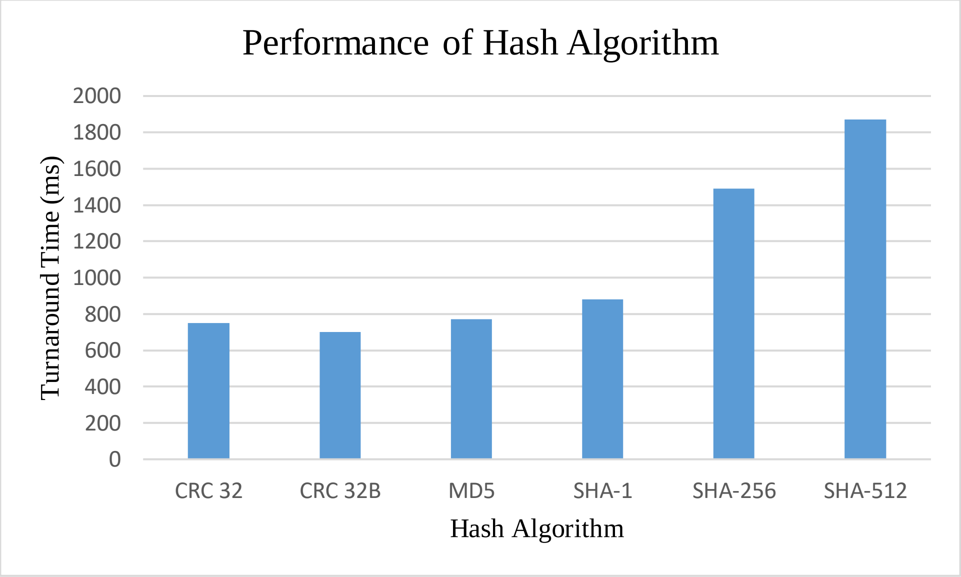

Figure 10.19:- Performance of the Hash Algorithm

A single character change in digital evidence would result in a different hash when CRC 32B is used. Hence, CRC 32B has the best turnaround time and requires less storage than the other hashing algorithms. The integrity of digital evidence in this thesis is preserved using CRC 32B. Although the CRC 32B hashing algorithm has better performance (even only a fraction of a second), there is room for improvement to get to the best encryption and decryption performance. Overall, the DFR component introduces time lags. However, this is not too significant and can be improved with time.

The next section provides a brief overall report of the components in an attempt to sum up the above discussion.

10.5 OVERALL EVALUATION

Overall, using the conflict-aware VM placement solution, it can be concluded that the CTL of the VM instances has a direct effect on the time it takes to do the actual placement. A high

CTL takes longer than a lower one to do the actual placement using the conflict-aware VM placement solution. The same is true about the size of each VM instance. Bigger VM instances add more time lag on the placement. Therefore, the conflict-aware VM placement can be argued to come at a lower performance cost, but it adds time delays to the placement of a VM. The positive is that the solution addresses the threat of inter-VM attacks by effectively isolating the VMs of conflicting tenants. Moreover, it was noted that the simulation results were three times more than the results of actual placement on a real cloud environment. This is a good result for all those who rely on simulations to make decisions on the actual placement. It gives a good indication of what should be factored into simulated results for them to give a true picture of what to expect in a live cloud environment. It would be even better if the author had considered multiple cloud platforms, for then it would be easy to tell if there are variations or a generalisation that could be made. The good part though – from both the simulated (CloudSim) environment and the real cloud (OpenNebula) environment – is that the results are to a certain extent similar. The higher the CTL and bigger the VM, the more time it took to do the placement. The trend is similar in both a simulated and a real cloud environment, which indicates that the author is indeed evaluating ‘apples’ against ‘apples’.

Overall, using the strong risk-based MFA shows that no unauthorised user was granted access, even though they had correct credentials. Moreover, a significant percentage of unauthorised users with the correct credentials failed a global anomaly check. This demonstrates the effectiveness of the solution to keep out criminals even if they have correct credentials. Only a small fraction of the unauthorised users’ login attempts failed the local anomaly check, which is to a certain extent a good result, because it is less than those who failed a global anomaly check. However, the best (an ideal) scenario would be to have all the test user login attempts classified as second-degree outliers for failing a global anomaly check. This would ensure that the solution is totally non-compromising to illegitimate users. Surely, this might be a far-fetched ideal to strive for, but it is indeed the best. However, and despite our solution’s failure to reach an ideal case, it can still be concluded that it is plausible to solve challenges posed by attackers that use compromised user credentials. Future work could add to the work in hand by trying to further reduce the number of test user login attempts that fail the local anomaly check and to improve the results that are classified as a having failed the global anomaly check. The results show that the proposed strong risk-based MFA solution is effective and to a large extent efficient in preventing users who are armed with legitimate credentials from gaining access to systems that they are not authorised to access.

Overall, using incident classifiers, the 3-tier digital forensic readiness solution can proactively prevent data leakage threats prior to their occurrence and help track who accessed what resource, where, how and when. The solution also stands to improve the admissibility of digital evidence in court or at a disciplinary hearing. It does this using its capability to preserve the integrity and confidentiality of potential digital evidence. Inherent in these results is the fact that the solution stands to shorten the digital investigation time because digital evidence is collected and stored prior to an investigation. This is done in a forensically sound manner and therefore able to stand legal scrutiny. The fact that the proposed 3-tier digital forensic readiness solution is based on an industry standard gives it the necessary credibility in the court rooms. Surely, this solution can also be used for corporate investigations. However, modelling it for legal battles in courtrooms would give it the trustworthiness that it deserves.

The classification of incidents also seems to be improving as the solution continues to be used. However, this element of the solution requires a very large dataset for it to be able to accurately perform the classification of incidents on the fly. This classification might not be at the level the author wants it to be, but it proves the concept to be true. Thus, it may be concluded that the 3-tier digital forensic readiness solution has achieved its goal. It must be mentioned though, that while the current study was being carried out, the issue of cybercrime attribution came up a number of times. This is a critical issue today where attackers simply erase their tracks after completing their malicious acts. This study might to a certain extent be moving in the direction of cyberattack attribution. However, it must be categorically stated here, that the proposed model is not directly solving the attribution issue, though it would be an interesting and yet challenging problem to solve. As it is, there are certain things that could still be learnt from this study in an attempt to solve cybercrime attribution, for example, how to keep digital evidence with the highest levels of integrity and confidentiality. This is only possible when the system is able to retrieve the digital evidence before it can be erased by attackers.

The next section concludes this discussion and provides a high-level overview of the contents of Chapter 11.

10.6 CONCLUSION

From the perspective of the conflict-aware VM placement, it can be concluded that lower CTLs coupled with smaller VMs are better to place than higher CTLs coupled with large or extra-large CTLs in terms of their quick turnaround times. From this standpoint, CSPs would be advised to use many smaller VMs in favour of large or extra-large ones. However, this must be balanced with the cost of retrieving data from each of the small instances as compared to one large one. Unfortunately, this research did not cover the data retrieval aspect of VM placement. (This is part of future work.) Moreover, it can be concluded that large to extra-large VM instances that have the highest CTL of four are the worst to place (i.e. worst performance), whereas smaller ones with a CTL of zero are the best to place (i.e. best performance).

From the strong risk-based MFA's perspective, it can be concluded that additional authentication factors on top of the traditional usernames and passwords can help to address the problem of illegitimate users using compromised credentials to access systems that they are not authorised to access. Hence, CSPs could be advised to make use of extra authentication factors to make informed access decisions. Furthermore, it is also advisable to not discard the traditional form of authentication (i.e. username and password), but to use it to monitor keystroke dynamics when the user types these. More factors should be used to make the final decision of whether to grant or deny access. Finally, and to a certain extent, CSPs can use the proposed strong risk-based MFA solution because it has demonstrated its effectiveness and efficiency in preventing unauthorised users who are armed with legitimate credentials from gaining access to systems that they are not authorised to access.

From the digital forensic readiness point of view, it can be concluded that prior collection of digital evidence and its storage in a forensically sound manner can help improve the turnaround time of a digital forensic investigation process. It may be that the classification of incidents is not as perfect as one would want it to be, but this can be improved as the system continues to learn over time. From the performance costs that have been noticed in other components of the proposed solution, it can be concluded that confidentiality and integrity tools that might not be widely used but that have a quick turnaround time, might offer an alternative to tried-and-tested tools that take even longer to produce the desired results.

In concluding this evaluation chapter, it is encouraging to note that the evaluation of results discussion has demonstrated the feasibility and appropriateness of the proposed solution to address data leakages in the cloud on all three fronts, i.e. VM placement, authentication and digital forensic readiness. Hence, the claims that were made in the research hypothesis in Chapter 1 have been validated and found to hold on all three fronts, though performance issues are still pending. Despite the fact that the solution comes with some performance cost in terms of minor delays that it introduces to the turnaround response time, it still performs well to achieve its main target goal – that of preventing data leakage in the cloud. The minor delays in turnaround response times were to a certain extent expected. This is because adding any plugin or application interface programming to an existing solution is often associated with time delays. However, and according to the author, these delays have been kept to the lowest possible figures, most of them as a fraction of a second. Surely, there will always be room for improvement, and future work can be aimed at making these time delays even shorter. Having said this, the next chapter concludes this study and considers further future work.

CHAPTER 11 CONCLUSION

11.1 INTRODUCTION

Chapter 10 provided an overall evaluation and discussion of results of the proposed solution. The main goal was to demonstrate the viability and appropriateness of the proposed solution in addressing data leakages in the cloud from all three fronts, i.e. VM placement, authentication, and digital forensic readiness (as outlined in Chapter 6). Despite the fact that the solution comes with some performance cost in terms of minor delays introduced to the turnaround response time, it still performs well enough to achieve its main goal, that of preventing data leakage in the cloud, from authentication, to VM placement to digital forensic readiness.

Chapter 11 concludes this study and reflects on some of the pending issues that could be covered by future work. The chapter is structured as follows: Section 11.2 re-caps and re-states the problem statement and the main research question. Section 11.3 reflects on how each of the research objectives were achieved. Section 11.4 summarises the contribution of the solution to the body of scholarly knowledge. Section 11.5 highlights and briefly discusses the limitations of the study. Section 11.6 builds on these limitations to provide possible future work that might extend the solution to further issues. Section 11.7 concludes the chapter.

11.2 REVISITING THE PROBLEM STATEMENT

The key problem that the researcher undertook to address, is that of a data leakage threat. At the time of writing this thesis, the researcher noted that the data leakage threat occupied the top three spots in most of the surveys on top security concerns for the cloud. This trend has continued over the past four to seven years, i.e. beginning from around 2012 to 2018. Hence, this study deduced that the data leakage threat can be argued to be the number one and major security challenge that has damaged consumer confidence and resulted in the slow adoption of cloud services by consumers. Considering the number of data leakages that regularly make news headlines – the latest of which was the Equifax debacle (Equifax, 2018) – this is indeed a valid, relevant, timely and worthwhile challenge with which the information security industry is still battling. Its effects are strongly evident in the slow uptake of cloud computing services. It was on these grounds that the author decided to embark on this research to address the prevalent data leakage threats, with a specific focus on cloud computing. This focus is based on the finding that most security surveys show that the data leakage threat is prominent in cloud infrastructures.

Consequently, this study has made a plausible attempt to restore the hampered consumer confidence in and security of cloud services. It started off by arguing that the data leakage threat in the cloud could be holistically addressed from a VM placement perspective, an authentication perspective and a digital forensic readiness perspective. The idea of taking a holistic approach came after it was discovered that most of the existing solutions in the related work focus on only one of these aspects in isolation, without considering the others.

Furthermore, the study argues that conflict-aware VM placement might help prevent data leakage threats through inter-VM attacks by providing a physical separation of cloud-hosted data, based on conflict-of-interest classes between co-residing cloud clients. It must be noted though that for this particular point of view, there is a specific reference to data leakages as a result of inter-VM attacks. This was followed by the argument that a strong risk-based multi-factor authentication solution might help prevent data leakage in the cloud. Lastly, the study argues that digital forensic readiness might help to proactively detect and prevent data leakage threats prior to their occurrence, and that it might help track who accessed what resource, where, how and when.

In order to find supporting evidence for, or to disprove these arguments, the author posed a challenging and thought-provoking research question, i.e. ***how can we restore the damaged consumer confidence and improve the uptake and security of cloud services?***

In order to fully answer this crucial research question, the study identified and tackled three overarching research objectives that are derived from the research question. The next section first re-caps and re-states the research objectives and then shows how each of them was achieved and to what extent.

11.3 ACHIEVEMENT OF RESEARCH OBJECTIVES

Below follows a discussion on how each research objective was achieved and to what extent.

- *Objective 1: Critically analyse current cloud security trends with a specific focus on VM placement, authentication and digital forensic readiness.*

This study critically analysed the trends in security of cloud computing from all three these aspects (i.e. VM placement, authentication and digital forensic readiness). The critical analysis of existing literature was aimed at identifying cloud security trends. These were covered in three chapters, i.e. Chapters 2, 3 and 4. Chapters 2 and 3 are basically background work that ends by uncovering the key trends with regard to security challenges of cloud computing. One of the findings that is reached is that the security of cloud computing is indeed an important issue and its number one concern is data leakage threats.

Furthermore, the critical analysis of existing work helped to uncover numerous digital forensic challenges to look out for when considering investigations in the cloud. This is a reflection on the finding that cloud computing presents a hostile and challenging environment for digital forensic investigators. After a careful and critical analysis of existing work, this study identified that most related work is moving towards a proactive approach for digital forensics in order to facilitate effective and efficient investigations in such a hostile and challenging environment.

Chapters 2 and 3 established the baseline study for Chapter 4. However, Chapter 4 delved deeper to rigorously identify some of the most topical works and their contributions to the body of knowledge. The idea was basically to analyse the trends in the current contributions and solutions to make informed decisions when identifying research gaps. For example, this study found that recent research efforts were now moving towards adaptive authentication systems. Another finding was that the security research community was moving towards making use of contextual data for multifactor authentication purposes. The trend analysis also uncovered that the research community has focused on other aspects of VM placement without really addressing the issue of conflict of interest and data security. In terms of digital forensics, the trends showed an over-dependency on CSPs to provide potential digital evidence. The trends also highlighted a research gap when it comes to securing potential digital evidence. Each of these research gaps were considered in the exercise of requirement elicitation for the high-level conceptual architecture. The formulation of the requirements based on the research gaps was critical to show the main contribution that this study could provide to the body of knowledge. Overall, Objective 1 was fully achieved by the outcomes of Chapters 2, 3 and 4.

- *Objective 2: Investigate, design and develop an innovative architecture that integrates conflict-aware VM placement, cutting-edge authentication and digital forensic readiness to address data leakage threats in the cloud.*

After some investigation and requirement elicitation, this study proposed a high-level conceptual architecture that was designed to meet each of the system requirements raised in Chapter 5. In so doing, it achieved Objective 2, as it models how this research integrates a strong risk-based MFA with 3-tier DFR and conflict-aware VM placement models to address data leakage threats in the cloud. This is indeed a unique contribution of this study in that most of the existing solutions in the related works focus on addressing only one aspect in isolation. The integrated solution as proposed in this thesis provides great value to the curbing of data leakage threats in the cloud. The focus is on preventing data leakages from the underlying CSP’s infrastructure by putting up conflict-aware VM placement algorithms that stand to deter inter-VM attacks. It goes further to ensure that hosted VMs that were placed using conflict-aware VM placement algorithms, are not accessed by unauthorised users who might be equipped with stolen user credentials. This solution authenticates users who use a risk-based approach. Moreover, the proposed solution monitors access and records user activity in order to detect incidents before they happen. This DFR capability proactively gathers and preserves potential digital evidence in a forensically sound manner. This is done prior to a digital forensic investigation. In a nutshell: the research gaps in Chapter 4, which led to the identification of system requirements in Chapter 5 and eventually to the proposed high-level architecture in Chapter 6, together fully achieved Objective 2.

- *Objective 3: Implement and evaluate an innovative model as a proof-of-concept on a real cloud platform to demonstrate its practicability and suitability to prevent data leakage threats.*

The study achieved Objective 3 with the outcomes of Chapters 7, 8, 9 and 10. Chapters 7, 8 and 9 provide the proof-of-concept implementation for each component of the high-level architecture and demonstrate the practicability and suitability of the architecture to prevent data leakage threats. But in general, these three chapters demonstrated how this thesis achieved the ‘implement’ element of Objective 3 in full.

Chapter 10 evaluated and discussed the results of the proof-of-concept implementations in order to show how this study achieved the ‘evaluate’ element of Objective 3. This chapter evaluated the solution in detail to best demonstrate its appropriateness to address the data leakage threat in the cloud. Hence, both parts of this objective were fully achieved by the combination of Chapters 7, 8, 9 and 10.

The fore-going section reflected on how the research objectives of the study were achieved. Achieving the research objectives is an indication that the proposed solution can help restore the hampered consumer confidence and improve the uptake and security of cloud services. Therefore, it can be argued that this study managed to answer the ultimate research question.

The next section discusses the contribution of this thesis to the body of scholarly knowledge in the field. This is the part that separates this study from the rest of the existing work as it ties in nicely with the findings of Chapter 4 on related works. It also demonstrates how this study fills the existing research gaps that were identified in Chapter 4.

11.4 CONTRIBUTION OF THE STUDY

The first contribution of this study comes in the form of the CBVMP model that makes use of varying degrees of conflict, the construct of sphere-of-conflict and sphere-of-non-conflict to provide for the physical separation of VMs belonging to conflicting tenants. According to the author, the focus on conflict-aware VM placement is the first of its kind in the body of cloud computing knowledge. This is also confirmed by the findings in Chapter 4, which assert that existing related work on VM placement focuses on quality of service, optimal utilisation of resources, energy consumption and others, without any consideration for addressing conflict-of-interest issues that arise among tenants sharing the same cloud infrastructure. Hence, it is considered as the first major contribution of this work.

The second contribution of this study relates to the way it handles authentication. The point of departure for this work is that the 'enemy' is already at the gate and armed with the 'secret', i.e. the user credentials. To date, none of the existing literature has started off from this extreme point of departure. Hence, the author argues that the research in hand proposes a solution to a problem that many researchers might not yet know to exist. Although a strong risk-based approach to authentication might already exist in literature, none of the existing implementations combine both contextual and behavioural keystroke dynamics to tackle the problem of weak or inadequate user credentials. This strategy is one major improvement to the body of knowledge. Moreover, none of the existing implementations as identified in Chapter 4 places as much focus on the security of the login data as this solution does with its encryption and secret shares approach.

Thirdly, this study contributes to the body of knowledge by extending the capability of a digital forensic readiness model to proactively detect and potentially prevent security incidents from happening in the first place. According to the author, this is the first of its kind in the digital forensic field. Furthermore, this study's 3-tier digital forensic readiness model is meant to validate the work of the ISO/IEC 27043:2015 standard. This in itself is a substantial contribution to the body of knowledge.

Finally, the major contribution and unique solution of this work comes from integrating all three components of the proposed solution into one comprehensive solution, as illustrated in the high-level conceptual architecture. The integrated solution has an added value beyond that which is independently contributed by each of the individual components. Existing work seems to focus on each component in isolation, without taking an integrated approach. Hence, and according to the author, this is also the first of its kind in the body of knowledge.

This study has made significant contributions to the body of knowledge in terms of its features, how it achieved all the set research objectives and ultimately answered the main research question. However, some research efforts are lacking that could take the study a step further to become a fully-fledged data leakage prevention solution. These limitations are highlighted in the next section.

11.5 LIMITATIONS OF THE STUDY

As a reflexive exercise, this section briefly acknowledges and points out the limitations of the study.

- The first limitation is that the results of this study may be biased, because the results are based solely on a two cloud platforms, i.e. OpenNebula and OwnCloud, without considering others. Moreover, the hierarchical structure of Location, Data Centre, Cluster, Physical Node, Virtual Machine is best suited for OpenNebula. For any other cloud platforms, the proposed CVMP solution might require a change to the underlying cloud infrastructure.

- Secondly, the test data for authentication is limited to a few login data points on laptop devices and is also based on a few keystroke features. Increasing the test data, the features and using different devices could improve the results. On the VM placement front, the conflict-aware VM placement algorithms for the CBVMP model have not been optimised for efficiency in terms of memory and CPU consumption. Optimising these might improve the time lag that is introduced by the current solution.

- A third limitation is that the placement algorithms have not yet been subjected to a formal analysis. A formal analysis of the algorithms would ensure their correctness, completeness and termination.

- Finally, the encryption and decryption solution that is employed in this study is not based on an industry cryptography standard. This might be a problem for users who prefer tried-and-tested security solutions.

The above limitations (among others) reflect the need for further research. Hence, the next section briefly discusses the direction of future work that could build on this study.

11.6 FUTURE WORK

First of all, more experimentation must be conducted for the solution on other open source cloud platforms. Such an ideal case would help to remove biases and provide truly generalised results in order to give an accurate global view on the proposed solution's performance across different platforms. This could be done component by component, before the entire system could be tested as a whole.

Secondly, future research must focus on increasing the test data for user authentication. It is expected that the model would be more effective, and its efficiency would improve with more test data from different devices. This is one point to focus on to improve the resilience of the authentication component. Moreover, future work on authentication must focus on adding more features to the model. This would help to make the solution more sensitive to small changes in the data. This would in turn ensure that an even bigger percentage of unauthorised users would automatically be classified as posing a high risk and being required to authenticate with an OOB token. This is one way to improve the sensitivity and efficiency of the user authentication component. In addition, further studies could tap into conducting research to establish an optimal number of clusters and the best clustering algorithm. However, this is more of an optimisation problem and it can best be solved within the field of artificial intelligence.

Thirdly, future research may also focus on conflict-aware load balancing. This would help CSPs to ensure that load-balancing activities consider conflict-of-interest issues when migrating tenants' VM from one destination to another. Moreover, there is a need for research that would conduct a formal analysis to provide proofs of and verify the correctness, completeness and termination of the proposed VM placement algorithms. It would also be interesting to perform a formal analysis on efficiency to establish the amount of resources (i.e. memory space and CPU time) that is required to execute the proposed algorithms.

Finally, researchers may embark on improving the encryption scheme by making use of blocks and rounds, following the approaches in 3DES and AES cryptography systems. Such improvements must be balanced against the overall performance of both encryption and decryption. Hence, it might be a good approach to start with a small number of rounds or blocks and build upwards.

11.7 CONCLUSION

The data leakage threat is arguably the number one reason why some organisations are still hesitant to fully trust cloud service providers with their confidential data. Hence, this study attributes the slow uptake of cloud services to data leakage threats. The entire information security industry is still battling with this problem.

This study took the initiative to try and address the data leakage problem. Despite some limitations, it has made plausible contributions toward addressing the prevalent data leakage threat in the cloud. This has been achieved by proposing a conceptual architecture that integrates conflict-aware VM placement, a strong risk-based multifactor authentication, and a three-tier digital forensic readiness model in one holistic solution. Each of the three components contributes and provides an added value to the overall solution. With minor customisations and improvements on the limitations, cloud service providers can use the proposed solution as a baseline to come up with a fully-fledged data leakage control system for their cloud offerings.

This research began a couple of years ago and has resulted in a number of publications, some of which are still in the pipeline. Furthermore, this study has produced a number of Computer Science BSc honours and MSc projects that have resulted in several postgraduate students graduating as part of the study's deliverables. Surely, this has been an exciting and at times a frustrating journey. However, it was worth all the effort for the extensive research and priceless postgraduate supervision experience the author gained in the process.

REFERENCES

Ab Rahman, N. H., Glisson, W. B., Yang, Y. and Choo, K.-K. R. (2016). Forensic-by-design framework for cyber-physical cloud systems, *IEEE Cloud Comp*., vol. 3, no. 1, pp.: 50 – 59. DOI: 10.1109/MCC.2016.5

Ab Rahman, N. H., Cahyani, N. D. W. and Choo, K.-K. R. (2017). Cloud incident handling and forensic-by-design: cloud storage as a case study, Special Issue Paper, *Concurrency and Computation: Practice and Experience*, 2017, vol. 29, pp.: 1 – 16. DOI: 10.1002/cpe.3868

Ablon, L. (2018). Data Thieves: the motivation of cyber threat actors and their use and monetization of stolen data, a testimony, *RAND Corporation CT-490*, available online: https://www.rand.org/content/dam/rand/pubs/testimonies/CT400/.../RAND_CT490.pdf, accessed [21/11/2018].

Adams, A. and Sasse, M.A. (1999). Users are not the enemy, *Communications of the ACM, (CACM)*, vol. 42, no. 12, December 1999, ACM, New York, NY, USA, pp.: 40 – 46. DOI: 10.1145/322796.322806

Agarwal, A., Gupta, M., Gupta, S. and Gupta, S.C. (2011). Systematic Digital Forensic Investigation Model, *International Journal of Computer Science and Security (IJCSS)*, vol. 5, no. 1, March/April 2011, pp.:118 – 131. ISSN: 1985 – 1553

Ahamed, F. (2016). Security Aware Virtual Machine Consolidation in Cloud Computing, Master's Dissertation, *Western Sydney University*, Australia.

Ahmed, A.A. and Xue, C. (2018). Analyzing data remnant remains on user devices to determine probative artifacts in cloud environments, *Journal of Forensic Sciences*, vol. 63, no. 1, pp.: 112 – 121. DOI: https://doi.org/10.1111/1556-4029.13506

Ahvanooey, M.T., Li, Q., Hou, J., Mazraeh, H.D. and Zhang, J. (2018). AITSteg: An innovative text steganography technique for hidden transmission of text message via social media, in *IEEE Access*, vol. 6, pp.: 65981 – 65995. DOI: 10.1109/ACCESS.2018.2866063

Ajoy, P.B. (2012). Administrative action and the doctrine of proportionality in India, *IOSR Journal of Humanities and Social Science (JHSS),* vol. 1, no. 6, Sep/Oct 2012, pp. 16 – 23. ISSN: 2279-0837

Alashwali, E.S. and Rusmussen, K. (2018). On the feasibility of fine-grained TLS security configurations in web browsers based on the requested domain name, *The 14th International Conference on Security and Privacy in Communication Networks (SecureComm, 2018),* Singapore, 8 - 10 August 2018.

Alenezi, A., Atlam, H.F. and Wills, G.B. (2019). Experts reviews of a cloud forensic readiness framework for organizations, *Journal of Cloud Computing: Advances, Systems and Applications,* vol. 8, no. 11, pp.: 1 – 14. DOI: 10.1186/s13677-019-0133-z

Alharbi, S.A., Weber-Jahnke, J. and Traore, I. (2011). The proactive and reactive digital forensics investigation process: A systematic literature review, *International Journal of Security and Its Applications,* vol. 5, no. 4, pp.: 59 – 72.

Alharbi, S.A. (2014). Proactive system for digital forensic investigation, PhD Thesis, *University of Victoria*, USA.

Ali, L.M.D., Tappert, C.C., Qui, M. and Monaco, J.V. (2015). Authentication and identification methods used in keystroke biometric systems, in *17th International Conference on High Performance Computing and Communications (HPCC), 2015 IEEE 7th International Symposium on Cyberspace Safety and Security, and 2015 IEEE 12th International Conference on Embedded Software and Systems,* New York, 2015. pp.: 1424 – 1427. DOI: 10.1109/HPCC-CSS-ICESS.2015.66

Ali, S.A., Memon, S. and Sahito, F. (2018). Challenges and solutions in cloud forensics, *in the Proceedings of the 2018 2nd International Conference on Cloud Computing and Big Data*

Computing (ICCBDC'18), Barcelona, Spain, 03 – 05 August 2018, pp.: 6 - 10, ACM, New York, NY, USA, ISBN: 978-1-4503-6474-4, DOI: 10.1145/3264560.3264565

Almutairi, A., Sarfraz, M., Basalamah, S., Aref, W. and Ghafoor, A. (2012). A distributed access control architecture for cloud computing, in *IEEE Software*, vol. 29, no. 2, pp.: 36 – 44, March-April 2012. DOI: 10.1109/MS.2011.153

Alnajdi, S., Dogan, M. and Al-Qahtani, E. (2016). A survey on resource allocation in cloud computing, *International Journal on Cloud Computing: Services and Architecture (IJCCSA),* vol. 6, no. 5, October 2016, pp.: 1 – 11. DOI: 10.5121/ijccsa.2016.6501

Ammann, R. (2012). Network Forensic Readiness: a bottom-up approach for IPv6 networks, PHD Thesis, *Auckland University of Technology,* Australia.

Amri, S., Hamdi, H. and Brahmi, Z. (2017). Inter-VM interference in cloud environments: a survey, in *2017 IEEE/ACS 14th International Conference on Computer Systems and Applications (AICCSA),* Hammamet, 2017, pp.: 154-159. DOI: 10.1109/AICCSA.2017.122

Amza, C., Cox, A.L. and Zwaenepoel, W. (2003). Conflict-aware scheduling for dynamic content applications, in 4th *Proceedings of the USENIX Symposium on Internet Technologies and Systems,* 26 – 28 March 2003, Red Lion Hotel, Seattle, WA, USA. pp.: 71 – 85.

Araiza, A.G. (2011). Electronic discovery in the cloud, *Duke Law & Technology Review,* no. 008, available online: http://scholarship.law.duke.edu/, accessed: [25/10/2012].

Avram, M.G. (2014). Advantages and challenges of adopting cloud computing from an enterprise perspective, *Procedia Technology,* Vol. 12, pp.: 529 – 534. DOI: https://doi.org/10.1016/j.protcy.2013.12.525

Axway (2013). Cloud Research 2013, IDG Research services, available online: https://www.idg.com/tools-for-marketers/cloud-research-2013/, accessed: [06/06/2017].

Bagnasco, S., Vallero, S. and Zaccolo, V. (2017). Fair share scheduling for OpenNebula (FaSS): implementation and performance tests, *CHEP 2018 Conference,* Sofia, Bulgaria, 10 July 2018.

Barreto, L., Scheunemann, L., Fraga, J. and Siqueira, F. (2017). Secure storage of user credentials and attributes in federation of clouds. In *Proceedings of the Symposium on Applied Computing (SAC '17)*. *ACM*, New York, NY, USA, pp.: 364 – 369. DOI: https://doi.org/10.1145/3019612.3019627

Barske, D., Stander, A. and Jordaan, J. (2010). A Digital Forensic Readiness Framework for South African SMEs, *Information Security for South Africa (ISSA),* Johannesburg, South Africa, pp.: 1 – 6. DOI: 10.1109/ISSA.2010.5588281

Barthe, G., Betarte, G., Campo, J.D. and Luna, C. (2011). Formally verifying isolation and availability in an idealized model of virtualization. In: Butler, M., Schulte, W. (eds) FM 2011: Formal Methods. FM 2011. *Lecture Notes in Computer Science*, vol. 6664. Springer, Berlin, Heidelberg, DOI: https://doi.org/10.1007/978-3-642-21437-0_19

Bartok, D. and Mann, Z.A. (2015). A branch-and-bound approach to virtual machine placement, in *Proceedings of the 3rd HPI Cloud Symposium "Operating the Cloud"*, 2015, pp.: 49 – 63.

Bartolini, C., El Kateb, D. Le Traon, Y. and Hagen, D. (2018). Cloud providers viability – how to address it from an IT and legal perspective, Electronic Markets, vol. 28, pp.: 53 – 75. DOI: 10.1007/s12525-018-0284-7

Battaglia, A.J. (2016). Disclosure and discovery under the federal rules of civil procedure, available online: https://judiciary.house.gov/wp-content/uploads/2013/07/Civil2016.pdf, accessed: [05/12/2018].

Baudoin, C. (2010). Cloud computing: fear, hype, reality and pragmatics, *Cutter Consortium, Summit 2010: Strategies for the Road Forward*, Cambridge, MA, 25 – 27

October 2010, available online: www.cutter.com/summit/2010.html, accessed: [08 April 2011].

Baykara, M. Das, R. and Tuna, G. (2017). Cryptology: A new approach to provide log security for digital forensics, *IU-JEEE,* vol. 17, no. 2, pp.: 3453 – 3462.

Bholowalia, P. and Kumar, A. (2014). EBK-Means: A clustering technique based on elbow method and k-means in WSN, *International Journal of Computer Applications,* vol. 105, no. 9, November 2014, pp.: 17 – 24. DOI: 10.5120/18405-9674

Birk, D. (2011). Technical challenges of forensic investigations in cloud computing environments, *Workshop on Cryptography and Security in Clouds, IBM Forum,* 15 – 16 March 2011, Zurich, Switzerland, pp.: 1 – 6.

Birk, D., Panico, M., Alva, A., Jaeger, B., Dykstra, J., Ruan, K., Austin, R., Ginsburg, A., Lichtenauer, B., Santos, L.J.R, Scoboria, E., Scoboria, K. and Yeoh, J. (2013). Mapping Forensic Standard ISO.IEC 27037 to Cloud Computing, *Cloud Security Alliance (CSA) - Incident Management and Forensic Working Group*, June 2013, pp.: 1 – 31.

Birk, D. and Wegener, C. (2011). Technical issues of forensic investigation in cloud computing environments, in *the 2011 Sixth International workshop on Systematic Approaches to Digital Forensic Engineering (SADFE)*, Oakland, CA, 26 May 2011, pp.: 1 – 10. DOI: 10.1109/SADFE.2011.17

Bittencourt, L.F. and Madeira, E.R.M. (2011). HCOC: a cost optimization algorithm for workflow scheduling in hybrid clouds, *Journal of Internet Services and Applications*, Vol. 2, no. 2, pp.: 207 – 227, DOI:10.1007/s13174-011-0032-0

Blackledge, J.M., Bezobrazov, S., Tobin, P. and Zamora, F. (2013). Cryptography using evolutionary computing, *24th IET Irish Signals and Systems Conference (ISSC 2013),* Letterkenny, 2013, pp.: 1 – 8. DOI: 10.1049/ic.2013.0029

Bollo, J. (2017). Mobile forensics must keep up with the times, *Forensics Magazine, Digital Forensic Insider,* Available online: https://www.forensicmag.com/article/2017/06/mobile-forensics-must-keep-times, accessed: [06/12/2018].

Botta, A., de Donado, W., Persico, V. and Pescape, A. (March 2016). Integration of cloud computing and Internet of Things: a survey, *Future Generation Computer Systems (FGCS),* vol. 56, pp.: 684 – 700, DOI: https://doi.org/10.1016/j.future.2015.09.021

Bu-Pasha, S. (2017). Cross-border issues under EU data protection law with regards to personal data protection, *Journal of Information & Communications Technology Law*, vol. 26, no. 3, pp.: 213 – 228. DOI: 10.1080/13600834.2017.1330740

Burke, E.K., Kendall, G. and Whitwell, G. (2009). A simulated annealing enhancement of the best-fit heuristic for the orthogonal stock-cutting problem, *INFORMS Journal on Computing*, vol. 21, no. 3, February 2009. DOI: https://doi.org/10.1287/ijoc.1080.0306

Butterfield, E., Dixon, M., Miller, S. and Schreuders, Z.C. (2018). Automated digital forensics, CARI Project, *Leeds Beckett University and West Yorkshire Police,* available online: eprints.leedsbeckett.ac.uk, accessed: [06/12/2018].

CA Technologies (2014). Authentication Strategy: Balancing Security and Convenience, available online: http://www.ca.com/content/dam/ca/us/files/ebook/intelligent-authentication-balancing-security-and-convenience.pdf, accessed: [08/03/2016].

Caballer, M., Segrelles, D., Molto, G. and Blanquer, I. (2014). A platform to deploy customized scientific virtual infrastructures on the cloud, *2014 6th International Workshop on Science Gateways (IWSG),* Dublin, Ireland, 2014, pp.: 42 – 47. DOI: 10.1109/IWSG.2014.14

Cable & Wireless Worldwide (2011). Ensuring security: the last barrier to cloud adoption, *Cable & Wireless Worldwide*, Whitepaper, March 2011.

Calheiros, R.N., Ranjan, R., Beloglazov, A., De Rose, C. A.F. and Buyya, R. (2010). CloudSim: a toolkit for modelling and simulation of cloud computing environments and evaluation of resource provisioning algorithms, *Software – Practice and Experience, Wiley Online Library,* vol. 41, August 2010, pp.: 23 – 50. DOI: 10.1002/spe.995

Calloway, T.J. (2012). Cloud computing, clickwrap, agreements, and limitation on liability clauses: a perfect storm? *Duke Law & Technology Review,* vol. 11, no. 1, pp.: 163 – 174.

Carrier, B. (2003). Defining digital forensic examination and analysis tools using abstraction layers, *International Journal of Digital Evidence,* vol. 1, no. 4, Winter 2003, pp.: 1 – 12.

Carrier, B. and Spafford, E.H. (2003). Getting physical with the digital investigation process, *International Journal of Digital Evidence,* vol. 2, no. 2.

Casey, E. (2011). Digital Evidence and Computer Crime: Forensic Science, Computers and the Internet, 3rd Edition, USA. ISBN: 978-0-12-374268-1

Casey, E. (2012). Cloud computing and digital forensics, *Digital Investigation,* vol. 9, no. 2012, pp.: 69 – 70. DOI: http://dx.doi.org/10.1016/j.diin.2012.11.001

Catteddu, D. and Hogben, G. (2009). Cloud computing: benefits, risks and recommendations for information security, *European Network and Information Security Agency (ENISA)*, November 2009, available online: www.enisa.europa.eu/act/it/library, accessed: [08/04/2011].

Centrify Corporation (2012). Using Microsoft active directory to address payment card industry (PCI) data security standard requirements in heterogeneous environments, White Paper, *2012 Centrify Corporation,* available online: http://resources.idgenterprise.com/original/AST-0076426_centrify_wp011_active_directory_and_pci.pdf, accessed: [09/12/2012].

Challita, S., Paraiso, F. Merle, P. (2017). A study of virtual machine placement optimization in data centers, in *7th International Conference on Cloud Computing and Services Science, CLOSER 2017,* April 2017, Porto, Portugal, 24 – 26 April 2018.

Cheng, L., Liu, F. and Yao, D. (2017). Enterprise data breach: causes, challenges, prevention, and future directions, *Wires Data Mining and Knowledge Discovery, John Wiley & Sons, Ltd,* vol. 7, pp.: 1 – 14. DOI: 10.1002/widm.1211

Chung, H., Park, J., Lee, S. and Kang, C. (2012). Digital forensic investigation of cloud storage services, *Digital Investigation, Elsevier,* vol. 9, no. 2, November 2012, pp.: 81 – 95. DOI: https://doi.org/10.1016/j.diin.2012.05.015

Cloud Security Alliance (CSA) (2009). Security Guidance for Critical Areas of Focus Cloud Computing, Version 2.1, December 2009, available online: www.cloudsecurityalliance.org/guidance/csaguide.v2.1.pdf, accessed: [18/04/2011].

Cloud Security Alliance (CSA) (2016). The Treacherous 12 – Cloud Computing Top Threats in 2016, available online: https://cloudsecurityalliance.org/group/top-threats/, accessed: [30/11/2018].

Cloud Security Alliance (CSA) (2017). The Treacherous 12 – Top Threats to Cloud Computing + Industry Insights, available online: https://cloudsecurityalliance.org/group/top-threats/, accessed: [30/11/2018].

Cobb, M. (2011). Digital forensic investigation procedure: Form a computer forensic policy, *TechTarget, ComputerWeekly.com,* available online: [http://www.computerweekly.com/tip/Digital-forensic-investigation-procedure-Form-a-computer-forensics-policy], accessed: [19/11/2012].

Computer Forensics and Computer Expert Witness Services (1993). U.S. Supreme Court Daubert v. Merrell Dow Pharmaceuticals, Inc., 509 U.S. 579 (1993), available online: http://euro.ecom.cmu.edu/program/law/08-732/Evidence/Daubert-Dow.pdf, accessed: [06/12/2012].

Cross, D.D. and Kuwahara, E. (2010). E-discovery and cloud computing: control of ESI in the cloud, *EDDE Journal,* vol. 1, no. 2, Spring 2010, pp.: 2 – 12.

D'Orazio, C.J. and Choo, K-K.R. (2018). Circumventing iOS security mechanisms for APT forensic investigation: A security taxonomy for cloud apps, *Future Generation Computer Systems (FGCS),* vol. 79, no. 1, pp.: 247 – 261. DOI: https://doi.org/10.1016/j.future.2016.11.010

Danielsson, J. and Tjøstheim, I. (2004). The need for a structured approach to digital forensic readiness: digital forensic readiness and e-commerce, *IADIS International Conference e-Commerce,* 2004, available online: http://el.trc.gov.om:4000/htmlroot/ENGG/tcolon/e_references/Consolidated/Computer%20Science/Journals/The%20need%20for%20a%20structured%20approach%20to%20digital%20forensic%20readiness.pdf, accessed: [08/10/2013].

Daryabar, F., Dehghantanha, A., Udzir, N.I., Sani, N.F.M., Shamsuddin, S. and Norouzizadeh, F. (2013). A survey about impacts of cloud computing on digital forensics, *International Journal of Cyber-Security and Digital Forensics (IJCSDF),* vol. 2, no. 2, pp.: 77 – 94. ISSN: 2305-0012

Daubner, L. (2018). Effective computer infrastructure monitoring, MSc Thesis, *Faculty of Informatics, Masaryk University*, Czech Republic.

de Ru, W.G. and Eloff, J.H.P. (1997). Enhanced password authentication through fuzzy logic, *IEEE Expert: Intelligent Systems and Their Application,* vol. 12, no. 6, November 1997, pp.: 38 – 45. DOI: http://dx.doi.org/10.1109/64.642960

DFRWS (2001). A road map for digital forensic research: collective work of all DFRWS attendees, DFRWS Technical Report, *The Proceedings of the Digital Forensic Research Conference (DFRWS 2001),* Utica, New York, USA, 7 – 8 August 2001.

Diaz, M., Martin, C. and Rubio, B. (2016). State-of-the-art, challenges, and open issues in the integration of the Internet of things and cloud computing, *Journal of Network and*

Computer Applications, vol. 67, pp.: 99 – 117, DOI: https://doi.org/10.1016/j.jnca.2016.01.010

Dlamini, M.T., Eloff, J.H.P. and Eloff, M.M. (2014). CBAC4C: Conflict Based Allocation Control for Cloud, *9th International Conference for Internet Technology and Secure Transactions (ICITST-2014)*, pp.: 447 – 448.

Dlamini, M.T., Eloff, J.H.P., Venter, H.S., Eloff, M.M., Henha Eyono, R.P.S. and Mosola, N.N. (2017). Behavioural analytics: beyond risk-based MFA, *The 2017 Proceedings of the Annual Southern Africa Telecommunication Networks and Applications Conference (SATNAC 2017)*, Freedom of the Seas Cruise Liner, Royal Caribbean International, Barcelona, Spain, 3 – 10 September 2017. ISBN: 978-0-620-76756-9

Dlamini, M.T., Eloff, M.M., Eloff, J.H.P., Venter, H.S., Chetty, K. and Blackledge, J.M. (2016). Securing cloud computing's blind-spots using strong and risk-based MFA, *Proceedings of the International Conference on Information resources Management (CONF-IRM 2016)*. May 2016, pp.: 1 – 21.

Dlamini, M.T., Venter, H.S., Eloff, J.H.P. and Eloff, M.M. (2011). Security of cloud computing: seeing through the fog, *The 2011 Proceedings of the Annual Southern Africa Telecommunication Networks and Applications Conference (SATNAC 2011)*, East London, South Africa, 4 – 7 September 2011.

Dlamini, M.T., Venter, H.S., Eloff, J.H.P. and Mitha, Y. (2012). Authentication in the cloud: a risk-based approach, *The 2012 Proceedings of the Annual Southern Africa Telecommunication Networks and Applications Conference (SATNAC 2012)*, Fancourt, George, South Africa, 2 – 5 September 2012.

Dlamini, M.T., Venter, H.S., Eloff, J.H.P. and Eloff, M.M. (2014). Requirements for preparing the cloud to become ready for digital forensic investigation, *The Proceedings of the 13th European Conference on Cyber Warfare and Security (ECCWS 2014)*, The

University of Piraeus, Piraeus, Greece, 3 – 4 July 2014, pp.: 242 – 252. ISBN: 978-1-910309-24-7

Dlamini, M.T., Venter, H.S. and Eloff, J.H.P. (2015). An innovative risk-based authentication mechanism for closing the new banking vault, *The Proceedings of the 3rd International Conference on Innovation and Entrepreneurship IICIE 2015)*, Durban, South Africa, March 2015, pp.: 72 – 80. ISBN: 978-1-910309-91-9

Dokras, S., Hartman, B., Mathers, T., Fitzgerald, B., Curry, S., Nystrom, M., Baize, E. and Mehta, N. (2009). The role of security in trustworthy cloud computing, RSA White Paper, *RSA Security Inc., CLOUD 0209*, 2009, available online: www.rsa.com, accessed: [14/042011].

Dostalek, L. (2019). Multi-factor Authentication Modeling, *2019 9th Conference on Advanced Computer Information Technologies (ACIT)*, 5-7 June 2019, Ceske Budejovice, Czech Republic, pp.: 443 – 446. DOI: 10.1109/ACITT.2019.8780068

Du, X., Le-Khac, N.-A. and Scanlon, M. (2017). Evaluation of digital forensic process models with respect to digital forensics as a service, in *The Proceedings of the 16th European Conference on Cyber Warfare and Security (ECCWS 2017)*, University College Dublin, Dublin, Ireland, 29 – 30 June 2017, pp.: 573 – 581.

Duan, Y., Fu, G., Zhou, N., Narendra, N.C. and Hu, B. (2015). Everything as a service (XaaS) on the cloud: origins, current and future trends, in *The 8th IEEE International Conference on Cloud Computing (CLOUD)*, New York, NY, USA, 27 June – 2 July 2015, pp.: 621 – 628, DOI: http://doi.ieeecomputersociety.org/10.1109/CLOUD.2015.88

Dubey, R., Jamshed, M.A., Wang, X. and Batala, R.K. (n.d.). Addressing security issues in cloud computing, Technical Report, *Carnegie Mellon University*.

Duranti, L. and Endicott-Popovsky, B. (2010). Digital records forensics: An interdisciplinary program for forensic readiness, *The Journal of Digital Forensics, Security and Law (JDFSL)*, vol. 5, no. 2, pp.:1 – 12. DOI: https://doi.org/10.15394/jdfsl.2010.1075

Dykstra, J. (2015). Seizing electronic evidence from cloud computing environments, in a book cloud technology: concepts, methodologies, tools and applications, *Information Resources Management Association*, USA. DOI: 10.4018/978-1-4666-6539-2.ch095: ISBN13: 9781466665392. ISBN10: 1466665394. EISBN13: 978146666540

Dykstra, J. and Sherman, A.T. (2012). Acquiring forensic evidence from infrastructure-as-a-service cloud computing: Exploring and evaluating tools, trust and techniques, *Digital Investigation*, vol. 9, supplement, August 2012, pp.: S90 – S98. DOI: https://doi.org/10.1016/j.diin.2012.05.001

Dykstra, J. and Sherman, A.T. (2013). Design and implementation of FROST: Digital forensic tools for the OpenStack cloud computing platform, *Digital Investigation*, vol. 10, August 2013, pp.: S87 – S95. DOI: http://dx.doi.org/10.1016/j.diin.2013.06.010

El-Gazzar, R., Hustad, E. and Olsen, D.H. (2016). Understanding cloud computing adoption issues: A Delphi study approach, *Journal of System and Software*, vol. 118, pp.: 64 – 84. DOI: https://doi.org/10.1016/j.jss.2016.04.061

Elsayed, M. and Zulkernine, M. (2016). IFCaaS: information flow control as a service for cloud security, *2016 11th International Conference on Availability, Reliability and Security (ARES)*, Salzburg, 2016, pp.: 211 – 216. DOI: 10.1109/ARES.2016.27

EMC Corporation (2013). RSA risk-based authentication: For RSA Authentication Manager 8.0, available online: https://www.emc.com/collateral/data-sheet/h11506-rsa-rba-ds.pdf, accessed: [08/03/2016].

Endicott-Popovsky, B. and Frincke, D.A. (2007). Embedding Hercule Poirot in networks: addressing inefficiencies in digital forensic investigations, *In Lecture Notes in Computer Science*, Vol. 4565, No. 2007, Springer. In: *Schmorrow D.D., Reeves L.M. (eds) Foundations of Augmented Cognition. FAC 2007. Lecture Notes in Computer Science*, vol. 4565. Springer, Berlin, Heidelberg, DOI: https://doi.org/10.1007/978-3-540-73216-7_41

Endicott-Popovsky, B., Frincke, D.A. and Taylor, C.A. (2007). A theoretical framework for organizational network forensic readiness, *Journal of Computers*, vol. 2, no. 3, May 2007, pp.: 1 – 11. DOI: 10.4304/jcp.2.3.1-11

Equifax (2018). Equifax releases updated information on 2017 Cybersecurity Incident, available online: https://www.equifaxsecurity2017.com/2018/03/01/equifax-releases-updated-information-2017-cybersecurity-incident/, accessed: [04/12/2018].

Experian, (2018). Data Breach Industry Forecast 2018, *Experian Data Breach Resolution, 2018 Edition*, available online: www.experian.com/.../data-breach/.../2018-experian-data-breach-industry-forecast.pdf, accessed: [21/11/2018].

EZComputer Solutions (2018). 6 things to consider before moving to the cloud, *EZComputer Solutions*, available online: https://www.ezcomputersolutions.com/blog/tips-before-moving-to-cloud/, accessed: [30 August 2018].

Fahmideh, M. and Beydoun, G. (2018). Reusing empirical knowledge during cloud adoption, *Journal of Systems and Software*, vol. 138, April 2018, pp.: 124 – 157, DOI: https://doi.org/10.1016/j.jss.2017.12.11

Farina, J., Scanlon, M., Le-Khac, N. and Kechadi, M. (2015). Overview of the forensic investigation of cloud services, *2015 10th International Conference on Availability, Reliability and Security (ARES)*, Toulouse, France, 24 – 27 August 2015, pp.: 556 – 565. DOI:10.1109/ARES.2015.81

Ferdaus, M.H., Murshed, M., Calheiros, R.N. and Buyya, R. (2017). An algorithm for network and data-aware placement of multi-tier applications in cloud data centers, *Journal of Network and Computer Applications*, vol. 98, November 2017, pp.: 65 – 83. DOI: http://dx.doi.org/10.1016/j.jnca.2017.09.009

Ferdowsi, A. (2011). “Yesterday’s Authentication Bug”, *Dropbox Blog*, available online: https://blogs.dropbox.com/dropbox/2011/06/yesterdays-authentication-bug/, accessed: [08/09/2016].

Fernandes, D.A., Soares, L.F., Gomes, J.V., Freire, M.M. and Inácio P.R. (2014). Security issues in cloud environments: a survey. *International Journal of Information Security* (Int. J. Inf. Secur). vol. 13, no. 2, April 2014, pp.: 113 - 170. Springer-Verlag Berlin, Heidelberg, DOI: http://dx.doi.org/10.1007/s10207-013-0208-7

Filho, M.C.S., Monteiro, C.C., Inacio, P.R.M. and Freire, M.M. (2018). Approaches for optimizing virtual machine placement and migration in cloud environments: A survey, *Journal of Parallel and Distribution Computing*, vol. 111, January 2018, pp.: 222 – 250. DOI: http://dx.doi.org/10.1016/j.jpdc.2017.08.010

Filkins, B.L., Kim, J.Y., Roberts, B., Armstrong, W., Miller, M.A., Hultner, M.L., Castillo, A.P., Ducom, J-C., Topol, E.J., Steinhubl, S.R. (2016). Privacy and security in the era of digital health: what should translational researchers know and do about it? *Am J Transl Res*. 2016; vol.8, no.3, pp.: 1560 – 1580. Published online 15 March 2016. PMCID: PMC4859641

Ford, B. (2012). Icebergs in the clouds: the other risks of cloud computing. *In* Proceedings *of the 4th USENIX conference on Hot Topics in Cloud Computing (HotCloud'12). USENIX Association*, Berkeley, CA, USA, pp.: 2 - 2.

Fortinet (2016). Defining Security for Today’s Cloud Environment: Security Without Compromise, available online: http://public.brighttalk.com/resource/core/138277/ebook-cloud- securitysolution_14707639869609039_244933.pdf, accessed: [12/06/2017].

Fred Cohen & Associates (2009). Analyst Report and Newsletter, *Fred Cohen & Associates,* available online: http://all.net/Analyst/2009-05b.pdf, accessed: [27/11/2012].

Furfaro, A., Garro, A. and Tundis, A. (2014). Towards Security as a service (Secaas): On the modelling of security services for cloud computing, in *The 2014 International Carnahan Conference on Security Technology (ICCST 2014),* Rome, 2014, pp.: 1- 6. DOI: 10.1109/CCST.2014.6986995

Garcia, M.A. (2015). Cloud Computing Forensics, in *Security, Trust, and regulatory Aspect of Cloud Computing in Business Environments,* DOI: 10.4018/978-1-4666-5788-5.ch010

Garfinkel, S.L. (2010). Digital forensics research: The next 10 years, *Digital Investigation,* vol. 7, 2010, pp.: S64 - S73. DOI: 10.1016/j.diin.2010.05.009

Garon, J.M. (2012). Searching Inside Google: Cases, Controversies and the Future of the World's Most Provocative Company, *Shepard Broad College of law,* October 2012, pp.: 1 - 31. DOI: http://dx.doi.org/10.2139/ssrn.1461463

Gemalto (2018). 2017 The Year of Internal Threats and Accidental Data Breaches, Gemalto Report, *Gemalto,* available online: https://www.gemalto.com/breach-level-index-2017-full-report, accessed: [30 August 2018].

Gemalto and Ponemon Institute (2018). The 2018 Global Cloud Data Security Study, *Gemalto,* available online: https://www2.gemalto.com/cloud-security-research/, accessed: [04/12/2018].

General Data Protection Regulation (GDPR) (2016). European (EU) 2016/679 of the European parliament and of the Council of 27 April 2016, *Official Journal of the European Union, L119, vol.59,* ISBN: 1977-0677

Ghazi, Y., Masood, R., Rauf, A, Shibli, M.U. and Hassan, O. (2016). DB-SECaaS: a cloud-based protection system for document-oriented NoSQL databases, *EURASIP Journal on Information Security* December 2016, vol.2016, no.16. DOI: https://doi.org/10.1186/s13635-016-0040-5

Global Research (2018). The Cloud Act is a Dangerous Piece of Legislation. Threatens Civil Liberties and Human Rights in the US and Worldwide, available online: https://www.globalresearch.ca/the-cloud-act-is-a-dangerous-piece-of-legislation-threatens-civil-liberties-and-human-rights-in-the-u-s-and-worldwide/5632751, accessed: [22/03/2018].

Gonzales, D., Kaplan, J.M., Saltzman, E., Winkelman, Z. and Woods, D. (2017). Cloud-Trust—a Security Assessment Model for Infrastructure as a Service (IaaS) Clouds, in *IEEE Transactions on Cloud Computing*, vol. 5, no. 3, 1 July-Sept. 2017, pp.: 523 – 536. DOI: 10.1109/TCC.2015.2415794

Goode, A. (2015). How Risk-based Authentication can Improve the Authentication Experience, available online: https://blogs.rsa.com/risk-based-authentication-can-improve-authentication-experience/, accessed: [08/03/2016].

Gornaik, S., Ikonomou, D., Saragiotis, P., Askoxylakis, I., Belimpasakis, P., Broda, M., Buttyan, L., Clemo, G., Kijewski, P., Merle, A., Mitrokotsa, K., Munro, A., Popov, O., Probst, C.W., Romano, L., Siaterlis, C., Siris, V., Verbauwhede, I. and Vishik, C. (2010). Priorities for Research on Current and Emerging Network Technologies, European Network and Information Agency (ENISA) European Union Agency) Report, 20 April 2010, available online: www.enisa.europa.eu/act/it/library/deliverables/procent, accessed: [08/04/2011].

Goyal, T., Singh, A. and Agrawal, A. (2012). Cloudsim: simulator for cloud computing infrastructure and modelling, in *Procedia Engineering*, vol. 38, no. 2012, pp.: 3566 – 3572. DOI: 10.1016/j.proeng.2012.06.412

Gregg, M. (2011). 10 Security Concerns for Cloud Computing, *Global Knowledge, Expert Reference Series of White Papers*, available online: www.globalknowledge.com, accessed: [18/04/2011].

Grispos, G., García-Galán, J., Pasquale, L. and Nuseibeh, B. (2017). Are you ready? Towards the engineering of forensic-ready systems, in *2017 11th International Conference on Research Challenges in Information Science (RCIS)*, Brighton, 2017, pp.: 328-333. DOI: 10.1109/RCIS.2017.7956555

Grispos, G., Storer, T. and Glisson, W. B. (2011). Calm before the storm: The emerging challenges of cloud computing in digital forensics. August 2011. [.pdf], available online: http://www.dcs.gla.ac.uk/~tws/papers/grispos11calm-rev2425.pdf, accessed: [09/09/2013].

Grobler, M. (2010). Digital Forensic Standards: International Progress, in *Proceedings of the South African information Security Multi-Conference (SAISMC 2010),* pp.: 261 – 271.

Grobler, T. and Louwrens, B.(2007), New Approaches for Security, Privacy and Trust in Complex Environments, in *IFIP International Federation for Information Processing*, vol. 232, , eds. Venter, H., Eloff, M., Labuschagne, L., Eloff, J.,von Solms, R., (Boston: Springer), pp.: 13 - 24. ISBN: 9780387723662

Grobler, C.P., Louwrens, C.P.and von Solms, S.H. (2010a). A framework to guide the implementation of Proactive Digital Forensics in organizations, in *2010 International Conference on Availability, Reliability and Security, Appears in IEEE Computer Society,* pp.: 677 – 682. DOI: 10.1109/ARES.2010.62

Grobler, C.P., Louwrens, C.P.and von Solms, S.H. (2010b). A multi-component view of digital Forensics, in *the Proceedings of the 2010 International Conference on Availability, Reliability and Security, appears in IEEE Computer Society*, pp.: 647 – 652. DOI:10.1109/ARES.2010.61

Gupta, U. (2011). Forensics in the Cloud: 5 Hot Skills; Today's Virtual Environment Puts New Demands on Investigators, *Bank Information Security Articles*, available online: http://www.bankinfosecurity.com/articles.php?art_id=4091&opg=1, accessed: [04/10/2013].

Gurary, J., Zhu, Y., Alnahash, N. and Fu, H.(2016) Implicit Authentication for Mobile Devices Using Typing Behavior, in *4th International Conference, Human Aspects of Information Security, Privacy, and Trust*, in In: *Tryfonas T. (eds) Human Aspects of Information Security, Privacy, and Trust. HAS 2016. Lecture Notes in Computer Science,* vol. 9750. Springer, Cham, Toronto, June 2016, pp.: 25-36. DOI: https://doi.org/10.1007/978-3-319-39381-0_3

Han, Y., Alpcan, T., Chan, J. and Leckie, C. (2013). Security Games for Virtual Machine Allocation in Cloud Computing, *in* 4th International Conference on Decision and Game Theory for Security, *vol. 8252* (GameSec 2013), Sajal Das, Cristina Nita-Rotaru, and Murat Kantarcioglu (Eds.), vol. 8252. *Springer-Verlag* New York, Inc., New York, NY, USA, pp.: 99 - 118. DOI: http://dx.doi.org/10.1007/978-3-319-02786-9_7

Han, Y., Alpcan, T., Chan, J. and Leckie, C. (2014). Virtual Machine Allocation Policies against Co-resident Attacks in Cloud Computing, in *Proceeding IEEE Conference on Communications (ICC 2014)*, Sydney, NSW, 2014, pp.: 786 – 792. DOI: 10.1109/ICC.2014.6883415

Han, Y. Chan, J. Alpcan, T., Leckie, C. (2015). Using Virtual Machine Allocation Policies to Defend against Co-resident Attacks in Cloud Computing, in *IEEE Transaction on Dependable and Secure Computing*, pp.: 1 – 14. DOI 10.1109/TDSC.2015.2429132

Han, Y., Alpcan, T., Chan, J., Leckie, C. and Rubinstein, B.I.P. (2016). A Game Theoretical Approach to Defend Against Co-Resident Attacks in Cloud Computing: Preventing Co-Residence Using Semi-Supervised Learning, in *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 3, March 2016, pp.: 556 - 570. DOI: 10.1109/TIFS.2015.2505680

Haque, M.A., Khan, Z.N. and Khatoon, G. (2015). Authentication through Keystrokes: What You Type and How You Type, in *2015 IEEE International conference on Research in Computational Intelligence and Communication Networks (ICRCICN)*, Kolkata, 2015, pp.: 257 -261. DOI: 10.1109/ICRCICN.2015.7434246

Hare-Brown, N. and Douglas, J. (2011). Digital investigations in the Cloud, *QCC Information Security Ltd, Digital Forensics Laboratories Whitepaper*, available online: www.qccis.com/whitepapers, accessed: [22/08/2011].

Harrington, S.L. (2012). Collaborating with a Digital Forensics Expert: Ultimate Tag-team or Disastrous Duo? *William Mitchell Law Review*, vol. 38, no. 1, January 2012, pp.: 353 – 396.

Hartigan, J.A. and Wong, M.A. (1979). A *K*-Means Clustering Algorithm, in *Journal of the Royal Statistical Society. Series C. (Applied Statistics)*, vol. 28, no. 1, pp.: 100 – 108.

Hashem, I.A.T., Yaqoob, I. Anuar, N.B., Mokhtar, S., Gani, A. and Khan, U. (January 2015). The rise of "big data" on cloud computing: Review and open research questions. In *Information Systems*, vol. 47, pp.: 98 – 115. DOI: https://doi.org/10.1016/j.is.2014.07.006

Hashemi, S.M. and Ardakani, M.R.M. (2012). Taxonomy of the Security Aspects of Cloud Computing, in *International Journal of Applied Information Systems (IJAIS), Foundation of Computer Science*, vol. 4, no. 1, September 2012. ISSN: 2249-0868

Hashizume, K., Rosado, D.G., Fernandez-Medina, E. and Fernandez, E.B. (2013). An analysis of security issues for cloud computing, in *Journal of Internet Services and Applications*, *a Springer Open Journal*, 2013, vol. 4, no. 5, pp.: 1 - 13. DOI: https://doi.org/10.1186/1869-0238-4-5

Hatch, O., Coons, C., Graham, L. and Whitehouse, S. (2018). Clarifying Lawful Overseas Use of Data Act or CLOUD Act, *Authenticated U.S. Government Information GPO*, available online: https://www.gpo.gov/fdsys/pkg/BILLS-115s2383is/pdf/BILLS-115s2383is.pdf, accessed [22/03/2018].

Hickey, A. (2018). Ahead of US v. Microsoft, CLOUD Act offers glimpse at proposed legal changes for data storage, available online: https://www.ciodive.com/news/ahead-of-us-v-microsoft-cloud-act-offers-glimpse-at-proposed-legal-change/517690/, accessed [22/03/2018].

Hook, S.A. (2018). Electronic Discovery in 2018: Current Challenges and Helpful Resources. in *American Bankruptcy Trustee Journal*. vol. 34, no. 1, Winter 2018, pp.: 14 - 19. DOI: http://hdl.handle.net/1805/15644

IDG Enterprise (2016). 2016 Cloud Computing Survey, available online: http://core0.staticworld.net/assets/2016/11/03/cloud_exec_summ_2016.pdf, accessed: [06/06/2017].

INFOSEC Institute (2016). A brief introduction to Forensic Readiness, available online: https://resources.infosecinstitute.com/a-brief-introduction-to-forensic-readiness/#gref, accessed: [08/09/2016].

Irwin, R. (March 2012). NHS Lancashire Forensic Readiness Policy, available online: http://www.centrallancashire.nhs.uk/Library/Documents/policies/information-governance-policies/NHS%20Lancashire%20Forensic%20Readiness%20Policy%20March%202013.pdf, accessed: [31/10/2012].

Ismail, S., Hassen, H.R., and Zantout, H. (2016). Open challenges in security of cloud computing. In *Proceedings of the International Conference on Big Data and Advanced Wireless Technologies, BDAW 2016,* Blagoevgrad, Bulgaria, 10 – 11 November 2016. *[a62] Association for Computing Machinery (ACM).* DOI: 10.1145/3010089.3016025

ISO/IEC 27043 (2015). Information technology – Security techniques – Incident investigation principles and processes, in *International Organization for Standardization, DIN German Institute for Standardization,* Berlin, available online: https://www.iso.org/standard/44407.html, accessed: [08/09/2016].

Iqbal, S., Kiah, L.M., Dhaghighi, B., Hussain, M., Khan, S., Khan, M.K. and Choo, K-K. R. (2016). On cloud security attacks: A taxonomy and intrusion detection and prevention as a service, in *Journal of Network and Computer Applications,* vol. 74, no. 2016, pp.: 98 – 120. DOI: http://dx.doi.org/10.1016/j.jnca.2016.08.016

Jabareen, Y. (2009). Building a Conceptual Framework: Philosophy, Definitions, and Procedure, in *International Journal of Qualitative Methods, International Institute for Qualitative Methodology (IIQM), University of Alberta,* vol. 8, no. 4, pp.: 49 – 62.

Jansen, W. and Grance, T. (2011). Guidelines on Security and Privacy in Public Cloud, NIST Draft Special Publication 800 – 144, January 2011, *National Institute of Standards and Technology* (*NIST*), Department of Commerce, USA.

Jula, A, Sundararajan, E and Othman, Z. (15 June 2014). Cloud computing service composition: A systematic literature review. *Expert Systems with Applications*, vol. 41, no. 8, pp.: 3809 – 3824, DOI: https://doi.org/10.1016/j.eswa.2013.12.017

Kassner, M. (2011), Digital forensics: The Science behind 'who done it', *TechRepublic*, available online: https://www.techrepublic.com/blog/it-security/digital-forensics-the-science-behind-who-done-it/, accessed: [25/01/2019].

Kaur, C. (2017). Cloud deployment Models: Public, Private or Hybrid. *Cynosure Solution Blog*, 11 January 2017, available online: http://www.cynosure-solutions.com/blog/tag/externally-hosted-private-cloud/, accessed: [19/06/2017].

Kaur, G. and Bhardwaj, V. (2016). A Review on VM Placement Strategies, in *International Journal of Advanced Research in Computer Science and Software Engineering (IJARCSSE)*, vol. 6, no. 5, May 2016, pp.: 521 – 526. ISSN: 2277 128X

Kebande, V.R. and Venter, H.S. (2016): On digital forensic readiness in the cloud using a distributed agent-based solution: issues and challenges, in *Australian Journal of Forensic Sciences (online), Taylor & Francis Group*, pp.: 1 – 30, DOI: 10.1080/00450618.2016.1194473

Kennedy, D., Kanjee, S., Schraader, D., Ward, E., Nyangaya, J.A., Gibson, C., Sing, N., Alberts, M., Kock, J. and Simon, T. (2013). Enemy at the gates: What executives need to know about the keys to the organization's front door. *2013 Deloitte & Touche*, available online [http://deloitteblog.co.za/wp-content/uploads/2013/08/Enemy-at-the-gates.pdf], accessed: [10/10/14].

Kent, K., Chevalier, S., Grance, T. and Dang, H. (2006). Guide to integrating forensic techniques into incident response. SP800-86. Gaithersburg: *US department of Commerce*, 2006.

Kernighan, B.W. and Lin, S. (1970). An efficient heuristic procedure for partitioning graphs, in *The Bell System Technical Journal*, vol. 49, no. 2, Feb. 1970, pp.: 291 – 307. DOI: 10.1002/j.1538-7305.1970.tb01770.x

Kesidis, G., Shan, Y., Jain, A., Urgaonkar, B., Khamse-Ashari, J. and Lambadaris, I. (2018). Scheduling Distributed resources in Heterogeneous Private Clouds, in *2018 IEEE 26th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, Milwauke, WI, pp.: 102 – 108. DOI: 10.1109/MASCOTS.2018.00018

Khan, M.N.A and Ullah S.W. (July 2017), A log aggregation forensic analysis framework for cloud computing environments, in *Computer Fraud & Security Journal*, vol. 2017, no. 7, pp.: 11 – 15. DOI: https://doi.org/10.1016/S1361-3723(17)30060-X

Khan, N. and Al-Yasiri, A. (2016). Identifying Cloud security Threats to Strengthen Cloud Computing Adoption Framework, in *the 2nd International Workshop on Internet of Things: Networking Applications and Technologies (IoTNAT' 2016), Procedia Computer Science*, vol. 94, no. 2016, pp.: 485 – 490. DOI: https://doi.org/10.1016/j.procs.2016.08.075

Khan, N. and Al-Yasiri, A. (2018). Cloud Security Threats and Techniques to Strengthen Cloud Computing Adoption Frameworks, *IGI Global*, DOI: 10.4018/978-1-5225-5634-3.ch016

Khosla, P.K. and Kuar, A. (2018). Big Data Security Solutions in Cloud. in *Mittal, M. Balas, V.E., Hemanth, D.J. and Kumar, R. Data Intensive Computing Applications for Big Data*, First Edition, *IOS Press, 2018,* Armstedam, Netherlands, pp.: 89, DOI:10.3233/978-1-61499-814-3-80. ISBN:1614998132 9781614998136

Kigwana, I. and Venter, H.S. (2018). A Digital Forensic Readiness Architecture for Online Examinations. in *South African Computer Journal,* vol. 30, no. 1, pp.: 1 – 39. https://doi.org/10.18489/sacj.v30i1.466

Knoll, T. (2018). Adapting Kerckhoff's principle: CPU Attacks leading a path from cryptography to open-source-hardware, A Technical Report, in *the 4th Wiesbaden Workshop on Advanced Microkernel Operating Systems (WAMOS 2018),* July 2018, Germany, pp.: 93 – 97.

Krishnamoorthy, S., Rueda, L., Saad, S. and Elmiligi, H. (2018). Identification of User Behavioral Biometrics for Authentication Using Keystroke Dynamics and Machine Learning. in *Proceedings of the 2018 2nd International Conference on Biometric Engineering and Applications (ICBEA '18). ACM,* New York, NY, USA, pp.: 50-57. DOI: https://doi.org/10.1145/3230820.3230829

Kulkarni, A.K. and Annappa, B. (2019). Context Aware VM Placement Optimization Technique for Heterogeneous IaaS Cloud, *IEEE Access Multidisciplinary, Rapid Review, Open Access Journal,* vol. 7, pp.: 89702 – 89713. DOI: 10.1109/ACCESS.2019.2926291

Kumar, A., Patwari, A. and Sabale, S. (2014). User Authentication by Typing Pattern for Computer and Computer based Devices, in *International Journal of Advanced Research in Computer and Communication Engineering,* vol. 3, no. 10, October 2014, pp.: 8132 - 8134.

Kumar Raju, B.K.S.P., Gosala, N.B. and Geethakumari, G. (2017). CLOSER: applying aggregation for effective event reconstruction of cloud service logs. in *Proceedings of the 11th International Conference on Ubiquitous Information Management and Communication (IMCOM '17), ACM,* New York, NY, USA, Article 62, 8 pages. DOI: https://doi.org/10.1145/3022227.3022288

Kumari, T. and Singh, K. (2018). A Review on Information Hiding Methods, in *International Journal of Engineering Science and Computing, (IJESC),* May 2018, vol. 8, no. 5, pp.: 17474 – 17476.

Kwiat, L., Kamhoua, C.A., Kwiat, K.A., Tang, J. and Martin, A. (2015). Security-Aware Virtual Machine Allocation in the Cloud: A Game Theoretic Approach, in *2015 IEEE 8th International Conference on Cloud Computing (CLOUD)*, New York City, NY, USA, 2015, pp.: 556 - 563. DOI:10.1109/CLOUD.2015.80

Leavitt, N. (2013). Hybrid Cloud Move to the Forefront, in *IEEE Computer Society, Technology News*, pp.: 15 -18. DOI: 10.1109/MC.2013.168Lewis, G.A. (May 2017). Cloud Computing, *IEEE Computer Society*, vol. 50, No. 5, pp.: 8 – 9. DOI: http://doi.ieeecomputersociety.org/10.1109/MC.2017.141

Levitin, G. Xing, L. and Dai, Y. (2018). Co-residence based data vulnerability vs. security in cloud computing system with random server assignment, in *European Journal od Operational Research*, vol. 267, no. 2, 2018, pp.: 676 – 686. DOI: https://doi.org/10.1016/j.ejor.2017.11.064

Leyden, J. (8 June 2017). Forcing digital forensics to obey 'one size fit all' crime lab standard is stupid and expensive. *The Register*, available online: https://www.theregister.co.uk/2017/06/08/digital_forensics_standards_push/, accessed: [11/10/2017].

Linthicum, D.S. (2017). PaaS Death Watch?, in *IEEE Cloud Computing*, vol. 4, no. 1, pp.: 6 – 9. DOI: 10.1109/MCC.2017.1

Litke, P. and Stewart, J. (2014). Cryptocurrency Stealing Malware Landscape, *Dell SecureWorks Counter Threat Unit^{TM} Threat Intelligence*, *Threats and Defenses*, available online: https://www.secureworks.com/research/cryptocurrency-stealing-malware-landscape, accessed: [23/11/2017].

Llorente, I.M. (2013). Eucalyptus, CloudStack, OpenStack and OpenNebula: A tale of two Cloud Models, available online: https://opennebula.org/eucalyptus-cloudstack-openstack-and-opennebula-a-tale-of-two-cloud-models/, accessed: [11/10/2017].

Loock, M., Eloff, J.H.P. (2005). A New Access Control Model based on the Chinese Wall Security Policy Model, in *Proceedings of the 5th Annual International Information Security South Africa (ISSA) conference*, Johannesburg, South Africa, July 2005. ISBN 1-86854-625X

Lopez, E.M., Moon, S.Y. and Park, J.H. (2016). Scenario-Based Digital Forensics Challenges in Cloud Computing, in *Symmetry 2016, MDPI*, vol. 8, no. 107, pp.: 1 – 20. DOI:10.3390/sym8100107

Ma, J. (14 December 2015). Top 10 Security Concerns for Cloud-Based Services, *Imperva Incapsula*, available online: https://www.incapsula.com/blog/top-10-cloud-security-concerns.html, accessed: [30/11/2018].

Ma, X. (2009). On the feasibility of data loss insurance for personal cloud storage, in *Proceedings of the 6th USENIX Syposium on Networked Systems Design and Implementation*, 22-24 April 2009, Boston, Massachusetts, USA, pp.: 2 - 2.

Madhusudhan, H.S. and Satish, K.T. (2017) Resource Optimization using Virtual Machine Placement: A Survey, in *Proceeding of the International Conference on Intelligent Computing Systems (ICICS 2017)*, San College of Technology, Salem, Tamilnadu, India, 15th – 16th December 2017, pp.: 157 – 164.

Mainstay (2016). An Economic Study of the Hyperscale Data Center, A White Paper, available online: http://cdn2.hubspot.net/hubfs/1931790/Content-Offers/Economic-Study-Hyperscale-Datacenter.pdf, accessed: [13/06/2017].

Maistry, A. (2015). Digital Forensic Readiness for the Cloud, Honours Project, *University of KwaZulu Natal*, Westville Campus, Durban, South Africa.

Majmudar, K.B. (1993). Daubert v. Merrell Dow: A Flexible Approach to the Admissibility of Novel Scientific Evidence, in *Harvard Journal of Law & Technology*, vol. 7, no. 1, Fall 1993, pp.: 187 – 205.

Mariani, L., Pezzè, M. and Zuddas, D. (2018). Augusto: Exploiting Popular Functionalities for the Generation of Semantic GUI Tests with Oracles. in *ICSE '18: ICSE '18: 40th International Conference on Software Engineering*, May 27-June 3, 2018, Gothenburg, Sweden. ACM, New York, NY, USA. DOI: https://doi.org/10.1145/3180155.3180162

Marshall, A,M. and Paige, R. (2018). Requirements in digital forensics method definition: Observations from a UK study. in *Digital Investigation,* DOI: 10.1016/j.diin.2018.09.004

Martim, E. Dlamini, M.T. van Greunen, D., Eloff, J.H.P. and Herselman, M. (2009). Is buying and transacting online easier and safer than down town?: an emerging economy perspective, in *4th International Conference on Information Warfare and Security*, Cape Town, South Africa, March 2009. DOI: http://hdl.handle.net/10204/3718

Martini, B. And Choo, K.R. (2012). An integrated conceptual digital forensic framework for cloud computing, in *Digital Investigation*, vol. 9, No. 2, November 2012, pp.: 71 – 80. DOI: https://doi.org/10.1016/j.diin.2012.07.001

Masdari, M., Nabavi, S.S. and Ahmadi,V. (2016) An Overview of virtual machine placement schemes in cloud computing, in *Journal of Network and Computer Applications,* vol. 66, 2016, pp.: 106 – 127. DOI: http://dx.doi.org/10.1016/j.jnca.2016.01.011

Mashayekhy, L., Nejad, M.M. and Grosu, D. (2014). A Framework for Data Protection in Cloud Federation, in *2014 43rd International Conference on Parallel Processing,* Minneapolis, 2014, pp.: 283 - 290. DOI 10.1109/ICPP.2014.37

Mason, S. and George, E. (2011). Digital evidence and 'cloud' computing, in *Computer Law & Security Review, Elservier Ltd*, vol. 27, no. 5, September 2011, pp.: 524 - 528. DOI: https://doi.org/10.1016/j.clsr.2011.07.005

Matsakis, L. (2018). Microsoft's Supreme Court Case Has Big Implications For Data, available online: https://www.wired.com/story/us-vs-microsoft-supreme-court-case-data/, accessed: [22/03/2018].

McKemmish, R. (1999). What is forensic computing? *Trends & Issues in Crime and Criminal Justice*, *Australian Institute of Criminology*, no. 118, pp.: 1 – 6. ISSN: 0817-8542

Mcleod, A.K. (2000). Is Frye Dying or Is Daubert Doomed? Determining the Standard of Admissibility of Scientific Evidence in Alabama Courts, in *Alabama Law Review*, vol. 51, no. 2, pp.: 883 – 905.

Mehmi, S., Verma, H.K. and Sangal, A.L. (2017). Simulation modelling of cloud computing for smart grid using CloudSim, in *Journal of Electrical Systems and Information technology*, vol. 4, October 2017, pp.: 159 – 172. DOI: http://dx.doi.org/10.1016/j.jesit.2016.10.004

Mell, P. and Grance, T. (2009). Effectively and Securely Using the Cloud Computing Paradigm, *22nd Annual Conference: "Awareness, Training and Education: The Catalyst for Organisational Change, Federal Information Systems Security Educator's Association" (FISSEA 2009)*, NIST, IT Lab, March 24 – 26, 2009, available online: csrc.nist.gov/.../fissea/.../fissea09-pmell-day3_cloud-computing.pdf, accessed: [18/04/2011].

Mell, P. and Grance, T. (2011). The NIST Definition of Cloud, *National Institute of Standards and Technology*, vol. 53, no. 6, pp.: 50, available online: http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc, accessed: [18/04/2011].

Mendoza, A., Kumar, A., Midcap, D., Cho, H. and Varol, C. (2015). BroStEx: A tool to aggregate browser storage artifacts for forensic analysis, in *Digital Investigation*, vol. 14, 2005, pp.: 63 - 75. DOI: http://dx.doi.org/10.1016/j.diin.2015.08.001

Mink, D., Yasinsac, A., Choo, K.-K. R. and Glisson, W. (2016). Next Generation Aircraft Architecture and Digital Forensic, in *AMCIS*, *Information Systems Security and Privacy (SIGSEC)*, San Diego, California, August 2016.

Miteva, A. (2018). 8 cloud adoption statistic that reveal the future of the cloud, *Cloud Worldwide Services,* available online: https://www.cloudworldwideservices.com/en/cloud-adoption-statistics-cloud-future/, accessed [04/12/2018].

Misbahuddin, M., Bindhumadhava, B.S. and Dheeptha, B. (2017). Design of a risk based authentication system using machine learning techniques, in *2017 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computed, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI),* san Francisco, CA, 2017, pp.: 1 – 6. DOI: 10.1109/UIC-ATC.2017.8397628

Miyachi, C. (208). What is "Cloud"? It is time to update the NIST definition? in *IEEE Cloud Computing,* May/June 2018, vol. 5, no. 3, pp.: 6 – 11. DOI: 10.1109/MCC.2018.032591611

Monteleone, S. and Puccio, L. (2017). From Safe Harbour to Privacy Shield: Advances and shortcomings of the new EU-US data transfer rules, in *European Parliament Research Service (EPRS), European Union*, January 2017, pp.: 1 – 36. DOI: 10.2861/09488

Morioka, E. and Sharbaf, M.S. (2016). Digital forensics research on cloud computing: An investigation of cloud forensics solutions, in *2016 IEEE Symposium on Technologies for Homeland Security (HST),* Waltham, MA, 2016, pp.: 1-6. DOI: 10.1109/THS.2016.7568909

Mosenia, A. and Jha, N. K. (2017). A Comprehensive Study of Security of Internet-of-Things," in *IEEE Transactions on Emerging Topics in Computing*, vol. 5, no. 4, pp.: 586-602, 1 Oct.-Dec. 2017. DOI: 10.1109/TETC.2016.2606384

Moshirnia, A. (2018). No Security Through Obscurity: Changing Circumvention Law to Protect our Democracy Against Cyberattacks, in *Brooklyn Law Review,* vol. 83, no. 4, pp.: 1279 – 1344.

Mosola, N.N., Dlamini, M.T., Blackledge, J.M., Eloff, J.H.P. and Venter, H.S. (2017). Chaos-based encryption keys and neural key-store for cloud-hosted data confidentiality, in *the Proceedings of the 20th Southern Africa Telecommunication Networks and Applications*

Conference (SATNAC 2016), Freedom of the Seas Cruise Liner, Royal Caribbean International, Barcelona, Spain, 3 – 10 September 2017. ISBN: 978-0-620-76756-9

Mosola, N.N., Dlamini, M.T., Eloff, J.H.P. Venter, H.S. and Eloff, M.M. (2016). Evolutionary Neural Crypto-System for Cloud-bound Data, in *Southern Africa Telecommunication Networks and Applications Conference (SATNAC) 2016*, Fancourt, George, Western Cape, South Africa, 4-7 September 2016.

Mouhtaropoulos, A., Li, C.-T. and Grobler, M. (2014). Digital Forensic Readiness: Are We There Yet? in *Journal of International Commercial Law and Technology (JICLT)*, vol. 9, no. 3, 2014, pp.: 173 - 179.

Moussa, A.N., Ithnin, N. and Zainal, A. (2018). CFaaS: bilaterally agreed evidence collection, *Journal of Cloud Computing: Advances, Systems and applications*, vol. 7, no. 1, pp.: 1 – 19. DOI: DOI 10.1186/s13677-017-0102-3

Mouton, F. and Venter, H.S. (2011). Requirements for Wireless Sensor Networks in Order to Achieve Digital Forensic Readiness, in *Proceedings of the 6th International Workshop on Digital Forensics & Incident Analysis (WDFIA 2011)*, Kingston University, London, UK, 7 – 8 July 2011.

Narwal, P., Kumar, D., and Sharma. M. (2016). A Review of Game-Theoretic Approaches for Secure Virtual Machine Resource Allocation in Cloud. *In Proceedings of the 2nd International Conference on Information and Communication Technology for Competitive Strategies (ICTCS 2016). ACM*, New York, NY, USA, Article 93, 5 pages. DOI: http://dx.doi.org/10.1145/2905055.2905152

Ngobeni, S., Venter, H.S. and Burke, I. (2012). The Modelling of a Digital Forensic Readiness Approach for Wireless Local Area Networks, in *Journal of Universal Computer Science (J.UCS)*, vol. 18, no. 12, 2012, pp.: 1721 – 1740. DOI: 10.3217/jucs-018-12-1721

Noor, T.H., Zeadally, S., Alfazi, A. and Sheng, Q.Z. (2018). Mobile cloud computing: Challenges and future research directions, in *Journal of Network and Computer Applications*, vol.115, pp.: 70 – 85. DOI: https://doi.org/10.1016/j.jnca.2018.04.018

Nuñez, D. and Agudo, I. (2014). BlindIdM: A privacy-preserving approach for identity management as a service. in *International Journal of Information Security*, vol. 13, no. 2, April 2014, pp.: 199 - 215. DOI:http://dx.doi.org/10.1007/s10207-014-0230-4

Nutonian (2018). Eureqa: The A.I.-Powered Modelling Engine, available online: https://www.nutonian.com/products/eureqa/, accessed: [12/03/2018].

Odeh, M., Hauer, T., Mcclatchey, R. and Solomonides, T. (2004). A Use-Case Driven Approach in Requirements Engineering: the MammoGrid Project, in *the 7th IASTED International Conference on Software Engineering Applications*, *CoRR*, Marina del Rey, USA November 2003, arXiv:cs.DB/0402008, http://arxiv.org/abs/cs.DB/0402008, Feb 2004.

Oliveira, T., Thomas, W. and Espadanal, M. (July 2014). Assessing the determinants of cloud computing adoption: An analysis of the manufacturing and services sectors, in *Information & Management,* vol. 51, no.5, pp.: 497 – 510. DOI: https://doi.org/10.1016/j.im.2014.03.006

Orofino, S. (1996). Daubert v. Merrell Dow Pharmaceuticals, Inc.: The Battle Over Admissibility Standards for Scientific Evidence in Court, in *Journal Undergraduate Science, History of Science,* vol. 3, Summer 1996, pp.: 109 - 111.

Orton, I., Alva, A. and Endicott-Popovsky, B. (2015). Legal process and requirements for Cloud Forensic Investigations, in *Cloud Technology: Concepts, Methods, Tools, and Applications,* Book Chapter, DOI: 10.4018/978-1-4666-6539-2.ch016

Pandey, U.N. (2018). Data Breach Statistics 2017: See What is the Status of Cloud Security? *LetToKnow,* available online: https://lettoknow.com/data-breach-statistics-2017-status/, accessed: [30 August 2018].

Pandya, M.K., Homayoun, S. and Dehghantanha, A. (2018). Forensics Investigation of OpenFlaow-Based SDN Platforms, *in A. Dehghantanha, M., conti and T. Dergahi (ets), Cyber Threat Intelligence, Advances in Information Security 70, Springer International Publishing AG, part of Springer Nature 2018,* Switzerland, pp.: 284. DOI: https://doi.org/10/1007/978-3-319-73951-9_14

Panko, R. (2017). 2017 Consumer Cloud Security Survey, *Clutch,* available online: https://clutch.co/cloud/storage/resources/consumer-security-survey-2017, accessed: [23/10/2019]

Park, S., Akatyev, N., Jang, Y., Hwang, J., Kim, D., Yu, W., Shin, H., Han, C. and Kim, J. (2018). A comparative study on data protection legislations and government standards to implement digital forensic readiness as a mandatory requirement, in *Digital Investigation*, vol. 24, no. 2018, pp.: 93 – 100. DOI: https://doi.org/10.1016/j.diin.2018.01.012

Park, S., Kim,Y. Park, G., Na, O. and Chang, H. (2018). Research on Digital Forensic Readiness Design in a Cloud Computing-Based Smart Work Environment, in *Sustainability 2018, MDPI,* vol. 10, No. 1203, pp.: 1 – 24. DOI: 10.3390/su10041203.

Patrascu, A. and Patriciu, V.-V. (2013). Beyond Digital Forensics. A Cloud Computing Perspective Over Incident Response and Reporting, in *Proceedings of the 8th International Symposium on Applied Computational Intelligence and Informatics (SACI),* Timisoara, Romania, May 2013, pp.: 455 – 460. DOI: 10.1109/SACI.2013.6609018

Pearson, S. and Charlesworth, A. (2009). Accountability as a Way Forward for Privacy Protection in the Cloud. in*: Jaatun M.G., Zhao G., Rong C. (eds) Cloud Computing. CloudCom 2009. Lecture Notes in Computer Science (LNCS)*, vol. 5931. Springer, Berlin, Heidelberg, DOI: 10.1007/978-3-642-10665-1_12

Perez-Botero, D., Szefer, J. and Lee, R.B. (2013). Characterizing hypervisor vulnerabilities in cloud computing servers. in *Proceedings of the 2013 international workshop on Security* *in cloud computing (Cloud Computing '13). ACM*, New York, NY, USA, pp.: 3-10. DOI: https://doi.org/10.1145/2484402.2484406

Perloff-Giles, A. (2018). Transnational Cyber-Offenses: Overcoming Jurisdictional Challenges, in *HeinOnline, 43 Yale J. International L.* vol. 191, no. 2018, pp.: 191 – 226.

Pham, T. (2014). Answer to OTP Bypass: Out-of-Band Two-Factor Authentication, available online: https://duo.com/blog/answer-to-otp-bypass-out-of-band-two-factor-authentication, accessed: [23/11/2017].

Pires, F.L. and Baran, B. (2015). Virtual Machine Placement Literature Review, *Polytechnic School National University of Asuncion Tech. Rep*, arXiv:01506.01509v1 [cs.DC], June 2015.

Pisani, P.H., Lorena, A.C., and de Carvalho, A.C.P.F.L. (2015). Ensemble of Adaptive Algorithms for Keystroke Dynamics, in *2015 Brazilian Conference on Intelligent Systems (BRACIS)*, Natal, Brazil, 2015, pp.: 310 - 315. DOI: 10.1109/BRACIS.2015.29

Ponemon Institute (2010). Flying Blind in the Cloud: The State of Information Governance, A Research Report, *Ponemon Institute LLC, sponsored by Symantec*, April 2010, available online: http://eval.symantec.com/mktginfo/enterprise/white_papers/b-ponemon_institute_flying_blind_in_the_cloud_WP.en-us.pdf, accessed: [08/05/2011].

Ponemon Institute (2014). 2014 Cost of Data Breach Study: United States, Ponemon Institute Research Report, *Ponemon Institute LLC*, May 2014, available online: *https://www.ponemon.org/blog/2014-cost-of-data-breach-united-states, accessed:* [08/09/2016].

Pooe, A. and Labuschagne, L. (2012). A conceptual model for digital forensic readiness, in *Information Security for South Africa (ISSA 2012)*, Johannesburg, South Africa, 15 – 17 August 2012, pp.: 1 – 8.

Quan, L., Wang, Z. and Ren, F. (2017). An RTT-Aware Virtual Machine Placement Method, in *Information, MDPI,* vol. 9, no. 4, 2018, pp.: 1 - 15. DOI: 10.3390/info9010004

Quick, D. and Choo, K.-K.R. (2018). Digital Forensic Data and Intelligence. in *Big Digital Forensic Data. Springer Briefs on Cyber Security Systems and Networks. Springer*, Singapore, DOI: https://doi.org/10.1007/978-981-13-0263-3_3

Quick, D. and Choo, K.-K.R. (2013a). Digital droplets: Microsoft SkyDrive forensic data remnants, in *Future Generation Computer Systems (FGCS),* vol. 29, 2013, pp.: 1378 – 1394. DOI: 10.1016/j.future.2013.02.001

Quick, D. and Choo, K.-K.R. (2013b). Dropbox analysis: Data remnants on user machines, in *Digital Investigation,* vol. 10, 2013, pp.: 3 – 18. DOI: http://dx.doi.org/10.1016/j.diin.2013.02.003

Quick, D. and Choo, K.-K.R. (2013c). Forensic collection of cloud storage data: Does the act of collection result in changes to the data or its metadata? in *Digital Investigation,* vol. 10, 2013, pp.: 266 – 277. DOI: http://dx.doi.org/10.1016/j.diin.2013.07.001

Quick, D. and Choo, K.-K.R. (2014). Google Drive: Forensic analysis of data remnants, in *Journal of Network and Computer Applications,* vol. 40, October 2014, pp.: 179 – 193. DOI: http://dx.doi.org/10.1016/j.jnca.2013.09.016

Quick, M., Hollowood, E., Miles, C. and Hampson, D. (2017). World's Biggest Data Breaches, in *Information is beautiful,* available online: http://www.informationisbeautiful.net/visualizations/worlds-biggest-data-breaches-hacks/, accessed: [06/06/2017].

Raghavan, S. (2013). Digital forensic research: current state of the art, in *CSI Transaction on ICT*, *Springer-Verlang,* vol. 1, no. 1, pp.: 91 – 114. DOI: https://doi.org/10.1007/s40012-012-0008-7

Random.org (2017). Random.org: True Random Number Service, available online: https://www.random.org/, accessed: [18/12/2017].

Raphiri, T.V., Dlamini, M.T., Venter, H.S. (2015). Strong Authentication: Closing the Front Door to Prevent Unauthorized Access to Cloud Resources, in *Proceedings of the 10th International Conference on Cyber Warfare and Security (ICCWS 2015)*, Kruger National Park, South Africa, March 2015, pp.: 252 - 260.

Ratsoma, M.S., Dlamini, M.T., Eloff, J.H.P. and Venter, H.S. (2015). A Conflict-aware Placement of Client VMs in Public Cloud Computing, *Proceedings of the 10th International Conference on Cyber Warfare and Security (ICCWS 2015)*, Kruger National Park, South Africa, March 2015, pp.: 502 - 509.

Reilly, D., Wren, C. And Berry, T. (2011). Cloud Computing: Pros and Cons for Computer Forensic Investigations, in *International Journal Multimedia and Image Processing (IJMIP)*, vol.1, no.1, March 2011, pp: 26-34.

Reith, M., Carr, C. and Gunsch, G. (2002). An Examination of Digital Forensic Models, in *International Journal of Digital Evidence*, vol. 1, no. 3, 2002, pp.: 2 – 12.

Rimal, B.P., Choi, E. and Lumb, I. (2009). A Taxonomy and Survey of Cloud Systems, in *The 5th International Joint Conference on INC, IMS and IDC, IEEE Computer Society*, Seoul, 2009, pp.: 44-51. DOI 10.1109/NCM.2009.218

Ristenpart, T., Tromer, E., Shacham, H. and Savage, S. (2009). Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. in *Proceedings of the 16th ACM conference on Computer and Communications Security (CCS '09)*. ACM, New York, NY, USA, pp.: 199 - 212. DOI: http://dx.doi.org/10.1145/1653662.1653687

Ritchey, R. (2011). Cloud Computing Security: Five key Considerations, Presentation Slides, DOE IMC, March 2011, available online: http://cio.energy.gov/documents/Tuesday_Marquis_4_1355_Ritchey.pdf, accessed: [08/05/2011].

Rowlingson, R. (2004). A Ten Step Process for Forensic Readiness, in *International Journal of Digital Evidence*, vol. 2, no. 3, Winter 2004, pp.: 1 - 28.

Ruan, K., Carthy, J., Kechadi, T. and Crosbie, M. (2011). Cloud Forensics: An Overview, available online: http://cloudforensicsresearch.org/publication/Cloud_Forensics_An_Overview_7th_IFIP.pdf, accessed: [10/09/2011].

Ruan, K., Carthy, J., Kechadi, M.T. and Baggili, I.M. (2013). Cloud forensics definitions and criteria for cloud forensic capability: An overview of survey results, in *Digital Investigation*, vol. 10, no. 1, pp.: 34 – 43. DOI: https://doi.org/10.1016/j.diin.2013.02.04

Sabi, H.M., Uzoka, F.-M.E, Langma, K. and Njeh, F.N. (2016). Conceptualizing a model for adoption of cloud computing in education, *International Journal of Information Management*, vol. 36, no. 2016, pp.: 183 – 191. DOI: http://dx.doi.ord/10.1016/j.ijinfomgt.2015.11.010

Saif, I., Siebenaler, R. and Mapgaonkar (2013). Risk-based Authentication: A Primer, *Deloitte*, available online:http://deloitte.wsj.com/cio/2013/10/30/risk-based-authentication-a-primer/, accessed: [08/03/2016].

Sailer, R., Jaeger, T., Valdez, E., C'aceres, R., Perez, R., Berger, S., Griffin, J.L. and van Doon, L. (2005). Building a MAC-based security architecture for the Xen open-source hypervisor. in *21st Annual Computer Security Applications Conference (ACSAC 2005)*, Tucson, AZ, December 2005, pp.: 10 pp.: – 285. DOI: 10.1109/CSAC.2005.13

Schmitt, V. and Jordaan, J. (2013). Establishing the Validity of MD5 and SHA-1 Hashing in Digital Forensic Practice in Light of Recent research Demonstrating Cryptographic Weakness in these Algorithms, in *International Journal of Computer Applications (0975 – 8887)*, vol. 68, no. 23, April 2013, pp.: 40 – 43. DOI: 10.5120/11723-7433

Shamsi, J.A., Zeadally, S., Sheikh, F. and Flowers, A. (2016). Attribution in cyberspace: techniques and legal implications, Special Issue Paper, in *Security and Communication*

Networks, Wiley Online Library, vol. 9, no. 15, April 2016, pp.: 2886 - 2900. DOI: 10.1002/sec.1485

Sharma, P.K. and Rajni (2012). Analysis of Image Watermarking using Least Significant Bit Algorithm, in *International Journal of Information Sciences and Techniques (IJIST)*, vol. 2, no. 4, July 2012, pp.: 95 – 101. DOI : 10.5121/ijist.2012.2409

Shen, G., Yang, F. and Zhou, L. (2018). Usable Security of Online Password Management with Sensor-Based Authentication, *Microsoft Corporation, US Patent*, Patent No.: US 9,858,402 B2, 2nd January 2018.

Shetty, S., Yuchi, X. and Song, M. (2016). Security-aware virtual machine placement in cloud data center, in *Moving Target Defense for Distributed Systems. Wireless Networks. Springer*, April 2016, pp. 13--24. DOI: https://doi.org/10.1007/978-3-319-31032-9_2

Si, Y., Xiaolin, G., Jiancai, L., Feng, T., Jianqiang, Z., and Min, D. (2014). A Security-Awareness Virtual Machine Management Scheme Based on Chinese Wall Policy in Cloud Computing, in *The Scientific World Journal*, vol. 2014, Article ID 805923, 12 pages, 2014. DOI: https://doi.org/10.1155/2014/805923

Sibiya, G., Venter, H.S. and Fogwill, T. (2012). Digital Forensic Framework for a Cloud Environment, in *IST-Africa 2012 Conference Proceedings, Paul Cunningham and Miriam Cunningham (Eds), IIMC International Information Management Corporation*, 9th May 2012, Tanzania. ISBN: 978-1-905824-34-2

Sibiya, G., Venter, H. S. and Fogwill, T. (2015). Digital forensics in the Cloud: The state of the art, in *2015 IST-Africa Conference*, Lilongwe, 2015, pp.: 1 - 9. DOI: 10.1109/ISTAFRICA.2015.7190540

Singh, J. (2014). Cyber-Attacks in Cloud Computing: A Case Study, in *International Journal of Electronics and Information Engineering*, vol. 1, no. 2, pp.: 78 – 87, December 2014.

Someswar, G.M. and Kalaskar, H. (2016). Design and Development of a Computational Model using Virtualization and Multi-tenancy Technologies for Cloud Computing Architecture, in *International Journal of Scientific Research in Science, Engineering and Technology (IJSRSET)*, vol. 2, no. 1, Jan/Feb 2016, pp.: 369 – 381. ISSN: 2394-4099

Stevens, G. (2012). Data Security Breach Notification Laws, in *Federal Information Security and Data Breach Notification Laws, Congressional Research Service 7-5700 - Report for Congress*, pp. 1-20, available online: www.crs.gov, accessed: [06/06/2017].

Stone, A. (2015). Chain of Custody: How to Ensure Digital Evidence Stands Up In Court, available online: https://www.govtechworks.com/chain-of-custody-how-to-ensure-digital-evidence-stands-up-in-court/#gs.EDJPelw, accessed: [09/01/2018].

Strom, D. (2015). Taking a Risk-Based Approach to Fraud Prevention, Vasco Whitepaper, *Vasco*, available online: https://www.vasco.com/Images/Taking%20a%20Risk-based%20approach%20to%20Fraud%20VASCO%20WP%20v4.pdf, accessed: [20/01/2016].

Su, K., Xu, C., Chen, C., Chen, W. and Wang, Z. (2015). Affinity and Conflict-Aware Placement of Virtual Machines in Heterogeneous Data Centers, *2015 IEEE 12th International Symposium on Autonomous Decentralized Systems,* Taichung, pp.: 289 – 294. DOI: 10.1109/ISADS.2015.42

Subramanian, N. and Jeyaraj, A. (2018). Recent security challenges in cloud computing, in *Computers and Electrical Engineering*, vol.71, pp.: 28 – 42, DOI: https://doi.org/10.1016/j.compeleceng.2018.06.006

Sule, D. (2014). Importance of Forensic Readiness, in *ISACA Journal*, vol.1, 2014

Sy, E., Burkert, C., Mueller, T., Federrath, H. and Fischer, M. (2019). QUICker connection establishment with out-of-band validation tokens, *Cornell University*, arXiv:1904.06228

Symantec (2013). Internet Security Threat Report, *Symantec Corporation*, available online: www.symantec.com/content/.../b-istr_main_report_v18_2012_21291018.en-us.pdf, accessed: [06/06/2017].

Takabi, H., Joshi, J.B.D. and Ahn, G.J. (2010). Security and Privacy Challenges in Cloud Computing Environments, in *IEEE Security & Privacy*, vol. 8, no. 6, pp.: 24 - 31, Nov.-Dec. 2010. DOI: 10.1109/MSP.2010.186

Tan, J. (2001). Forensic readiness, *Cambridge*, MA: Stake; 2001.

Taylor, K. (2012). Cloud Consideration: E-Discovery, *K&L Gates*, available online: http://www.legalcloudcentral.com/2012/10/articles/in-the-courts/cloud-considerations-ediscovery/#more, accessed: [27/11/2012].

Taylor, C., Endicott-Popovsky, B. and Frincke, D.A. (2007). Specifying digital forensic: A forensic policy approach, in *Digital Investigation*, Vol. 4, Supplement, September 2007, pp.: 101 – 104. DOI: https://doi.org/10.1016/j.diin.2007.06.006

Taylor, M., Haggerty, J., Gresty, D. and Hegarty, R. (2010). Digital Evidence in Cloud Computing Systems, in *Computer Law & Security Review Journal, Elsevier Ltd*, vol. 26, no. 3, May 2010, pp.: 304 – 308. DOI: https://doi.org/10.1016/j.clsr.2010.03.002

Taylor, M., Haggerty, J., Gresty, D. and Lamb, D. (2011). Forensic investigation of cloud computing systems, in *Network Security*, vol. 2011, no. 3, March 2011, pp.: 4 – 10. DOI: https://doi.org/10.1016/S1353-4858(11)70024-1

Thales and Ponemon Institute (2019). Protecting Data In The Cloud 2019 Thales Cloud Security Study, Global Cloud data Security Study – 2019 Report, *Thales* and *Ponemon Institute*, available online: https://www.thalesesecurity.com/2019/cloud-security-research, accessed: [08/10/2019].

Thomas J. and Goudar, R.H. (2018). Multilevel Authentication using QR code based watermarking with mobile OTP and Hadamard transformation, *2018 International*

Conference on Advances in Computing, Communications and Informatics (ICACCI), Bangalore, pp.: 2421 – 2425. DOI: 10.1109/ICACCI.2018.8554891

Tieng, D.G. and Nocon, E. (2016). Some Attacks on Shamir's Secret Sharing Scheme by Inside Adversaries, in *Proceedings of the DLSU Research Congress 2016,* vol. 4, De La Salle University, Manila, Philippines, March 2016. ISSN:2449-3309

Teing, Y-Y., Dehghantanha, A., Choo, K.-K.R. and Yang, L.T. (2017). Forensic investigation of P2P cloud storage services and backbone for IoT networks. in *Computers & Electrica Engineering,* vol. 58, (February 2017), pp.: 350 - 363. DOI: https://doi.org/10.1016/j.compeleceng.2016.08.020

Thilakarathne, A. and Wijayanayake, J.I. (2014). Security of Cloud Computing, in *International Journal of Scientific & Technology Research,* vol. 3, no. 11, November 2014, pp.: 200 – 203. ISSN: 2277-8616

Thulo, M. and Eloff, J.H.P. (2017). Towards Optimized Security-aware (O-Sec) VM Placement Algorithms, in *Proceedings of the 3rd International Conference on Information Systems Security and privacy (ICISSP 2017),* pp.: 411 – 422. DOI: 10.5220/0006206504110422

Tout, S. (2018). The Growing Issue of Compromised Credentials, *Forbes,* available online: https://www.forbes.com/sites/forbestechcouncil/2018/10/12/the-growing-issue-of-compromised-credentials/#3dfb1d4d2434, accessed: [21/11/2018].

Tsai, T.-H.., Chen. Y., Huang, H.-C., Huang, P.-M. and Chou, K.-S. (2011). A Practical Chinese Wall Security Model in Cloud Computing. in *The 13th Asia-Pacific Network Operations and Management Symposium (APNOMS): Managing Clouds, Smart Networks and Services*, Taipei, Taiwan, 2011; pp.: 1 - 4. DOI: 10.1109/APNOMS.2011.6076992

Usmani, Z. and Singh, S. (2016). A Survey of Virtual Machine Placement Techniques in a Cloud Data Center, in *Procedia Computer Science,* vol.78, 2016, pp.: 491 – 498. DOI: https://doi.org/10.1016/j.procs.2016.02.093

Valjarevic, A., Venter, H., Petrovic, R. (2016). ISO/IEC 27043:2015 - role and application. in *2016 24th Telecommunications Forum, TELFOR 2016,* Belgrade, 2016, pp.: 1 – 4. DOI: 10.1109/TELFOR.2016.7818718

Van Wankle, W.W. (2012). External Private Clouds: What, When and Why, *ITPRO Real-World Business Technology,* 26 March 2012, available online: http://www.tomsitpro.com/articles/private_cloud-resource_pooling-rapid_elasticity-cloud_computing-external_private_cloud,2-279.html, accessed: [19/06/2017].

Vasiljeva, T., Shaikhulina, S. and Kreslins, K. (2017). Cloud Computing: Business perspectives, benefits and Challenges for Small and Medium Enterprises (Case of Latvia), in *Procedia Engineering,* vol. 178, no. 2017, pp.: 443 – 451. DOI: https://doi.rog/10.1016/j.proeng.2017.01.087

Walsh, J.T. (1998). The Evolving Standards of Admissibility of Scientific Evidence, in *General Practice, Solo & Small Firm Division*, vol. 2, no. 1, Spring 1998.

Wang, B., Chen, C., He, L., Gao, B., ren, J., Fu, Z., Fu, S., Hu, Y. and Li, C-T. (2018). Modelling and developing conflict-aware scheduling on large-scale data centres, in *Future Generation Computer Systems (FGCS),* vol. 86, no. 2018, pp.: 995 – 1007. DOI: http://dx.doi.org/10.1016/j.future.2017.07.043

Wang, L., Li, M., Zhang, Y., Ristenpart, T. and Swift, M. (2018). Peeking Behind the Curtains of Serverless Platforms, in *Proceedings of the 18th 2018 USENIX Annual Technical Conference (USENIX ATC' 18),* 11-13 July 2018, Boston, Massachussetts, USA, pp.: 133 - 145. ISBN: 978-1-931971-44-7

Wang, Y., Anokhin, O., Anderl, R. (2017). Concept and Use Case Driven Approach for Mapping IT Security Requirements on System Assets and Processes in Industry 4.0, in *the 50th CIRP Conference on Manufacturing Systems,* vol. 63, no. 2017, pp.: 207 – 212. DOI: https://doi.org/10.1016/j.procir.2017.03.142

Wang, Z., Sun, K., Jajodia, S. and Jing, J. (2012). Disk Storage Isolation and Verification in Cloud, in *2012 IEEE Global Communication and Information System Security Symposium (Globecom)*, Anaheim, CA, 3 - 7 December 2012; pp.: 771 - 776. DOI: 10.1109/GLOCOM.2012.6503206

Webroot (2014). Secure Mobile Banking: Protecting Your Customers and Your Bottom Line, available online: [http://f6ce14d4647f05e937f4-4d6abce208e5e17c2085b466b98c2083.r3.cf1.rackcdn.com/secure-mobile-banking-protecting-your-customers-your-bottom-line-pdf-4-w-1141.pdf], accessed: [10/10/14].

Welch, C.H. (2006). Flexible Standards, Deferential Review: Daubert's Legacy of Confusion, in *Harvard Journal of Law & Public Policy*, vol. 29, no. 3, pp.: 1085 – 1105.

Wu, L., Du, X., Wang, W. and Lin, B. (2018). An Out-of-band Authentication Scheme for Internet of Things Using Blockchain Technology, *2018 International Conference on Computing, Networking and Communications (ICNC),* Maui, HI, pp.:769 – 773. DOI: 10.1109/ICCNC.2018.8390280

Wu, R., Ahn, G.-J., Hu, H. and Singhal, M. (2010). Information Flow Control in Cloud Computing, in *the 6th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom 2010),* Chicago, IL, October 2010, pp.: 1 - 7.

Yadav, S., Ahmad, K. and Shekhar, J. (2011). Analysis of Digital Forensic and Investigation, in *International Conference on High Performance and Grid Computing (HPAGC 2011),* vol. 1, no. 3, pp.: 435 – 441.

Yeh, K., Su, C., Chiu, W. and Zhou, L. (2018). I Walk, Therefore I Am: Continuous User Authentication with Plantar Biometrics, in *IEEE Communications Magazine*, vol. 56, no. 2, Feb. 2018, pp.: 150 - 157. DOI: 10.1109/MCOM.2018.1700339

Zatyko, K. (2007). Commentary: Defining Digital Forensics, in *Forensic Magazine,* Feb/March 2007, pp.: 1 - 5.

Zawoad, S., Dutta, A.K. and Hasan, R. (2013). SecLaaS: Secure Logging-as-a-Service for Cloud Forensics, in *the 8th Proceedings of the ACM Symposium on Information, Computer and Communications Security (ASIACCS 2013),* Hangzhou, China, May 7-10, 2013, pp.: 219 – 230. DOI: 10.1145/2484313.2484342

Zhang, Y., Juels, A., Reiter, M.K. and Ristenpart, T. (2012). Cross-VM Side Channels and Their Use to Extract Private Keys. in *The 19th ACM Conference on Computer and Communications Security,* 16 – 18 October 2012, Sheraton Raleigh Hotel, Raleigh, North Carolina, USA.

Zimmerman, S. and Glavach, D. (2011). Cyber Forensics in the Cloud, in *IA Newsletter for Information Assurance Technology Professionals*, Winter 2011, vol. 14, no. 1, pp: 4 – 7.

Zuikeviciute, V. and Pedone, F. (2008). Conflict-Aware Load-Balancing Techniques for Database Replication, in *Proceedings of the 2008 ACM symposium on Applied computing (SAC '08). ACM,* New York, NY, USA, pp.: 2169 - 2173. DOI: https://doi.org/10.1145/1363686.1364205

APPENDIX ABSTRACTS OF PUBLICATION OUTPUTS

A.1 SECURITY OF CLOUD COMPUTING: SEEING THROUGH THE FOG

Abstract: Cloud computing is a new computing paradigm for the provisioning, delivery and consumption of IT resources and services on the Internet. This computing paradigm comes with huge benefits such as cost savings, increased resilience and service availability, improved IT operations efficiency and flexibility. However, most research cites security concerns as one of the biggest challenges for most of these organizations. This has led to fallacy or misconception about security challenges of the ‘cloud’ which needs to be clarified. This is a call for more research to separate reality from the hype. Hence, this paper aims to separate justified security concerns from the hype, fear of the unknown and confusion that currently prevails within cloud computing. This paper aims to advance the current discussions on cloud computing security in order to clear the ‘foggy cloud’ hovering over such a promising technology development. It seeks to inform and make decision makers aware of the real pertinent and justified security issues within cloud computing.

Full Paper:
http://www.satnac.org.za/proceedings/2011/papers/Internet_Services_and_Applications/178.pdf

Dlamini, M.T., Venter, H.S., Eloff, J.H.P., and Eloff, M.M. 2011. Security of Cloud Computing: Seeing through the Fog, *The Southern Africa Telecommunication Networks and*

Applications Conference (SATNAC 2011), East London, South Africa, 4 – 7 September 2011.

A.2 AUTHENTICATION IN THE CLOUD: A RISK-BASED APPROACH

Abstract: Most companies are moving their data and applications to the cloud in order to exploit the numerous benefits that this computing paradigm presents. Yet, there is still insufficient research on how user authentication is to be handled on cloud computing environments. Cloud computing challenges the way people think about authentication and how to manage user identity across multiple domains. Hence, this paper outlines the requirements for user authentication and handling identity in the cloud. It goes further to discuss real world scenarios that illustrates the multi-faceted nature of handling authentication within a cloud environment. The main contribution of this paper is our proposed cloud-based authentication architecture. Our architecture makes a proposal on how to provide flexible, robust and scalable authentication by taking a risk-based approach to user authentication on cloud environments.

Full Paper:
http://www.satnac.org.za/proceedings/2012/papers/8.Data_Centre_Cloud/108.pdf

Full Reference:

Dlamini, M.T., Venter, H.S., Eloff, J.H.P. and Mitha, Y. (2012). Authentication in the cloud: a risk-based approach, *The 2012 Proceedings of the Annual Southern Africa Telecommunication Networks and Applications Conference (SATNAC 2012)*, Fancourt, George, South Africa, 2 – 5 September 2012.

A.3 REQUIREMENTS FOR PREPARING THE CLOUD TO BECOME READY FOR DIGITAL FORENSIC INVESTIGATION

Abstract: Some research work claims that the adoption rates for cloud computing have not scaled as anticipated. This is mainly due to security concerns. Hence, most research efforts have been directed at the cloud security challenges. In the meantime, researchers have rarely focused on an equally important issue: that of digital forensics investigation in the cloud. It is on this premise that this paper makes an attempt to understand the challenges and complexities of conducting an effective digital forensics investigation in the cloud. The contribution is to provide a better understanding of the implications that cloud computing poses for digital forensics investigators. This paper proposes digital forensic readiness as a solution to help prepare the cloud for an effective investigation. The paper goes further to explain and motivate why the authors view digital forensic readiness as the most appropriate solution. It also discusses the major factors that are influencing the move towards digital forensic readiness for cloud computing. As a work-in-progress, this paper ends by proposing necessary requirements that must be considered in order to make the cloud become ready for an effective investigation. It also provides a glimpse to the proposed three-tier digital forensic readiness model for the cloud which is grounded on these fundamental requirements.

Full Paper:
https://www.researchgate.net/profile/Andrew_Liaropoulos/publication/264337838_Proceedings_of_the_13th_European_Conference_on_Cyber_Warfare_and_Security/links/53d904af0cf2e38c6331db58/Proceedings-of-the-13th-European-Conference-on-Cyber-Warfare-and-Security.pdf#page=256

Full Reference

Dlamini, M.T., Venter, H.S., Eloff, J.H.P. and Eloff, M.M. (2014). Requirements for preparing the cloud to become ready for digital forensic investigation, *The Proceedings of the 13th European Conference on Cyber Warfare and Security (ECCWS 2014)*, The

University of Piraeus, Piraeus, Greece, 3 – 4 July 2014, pp.: 242 – 252. ISBN: 978-1-910309-24-7

A.4 CBAC4C: CONFLICT BASED ALLOCATION CONTROL FOR CLOUD

Abstract: Cloud infrastructures are vulnerable to serious data leakage threats. Tenants with conflicting interests, residing on a shared cloud infrastructure, can potentially view the data of other potentially conflicting tenants' by means of inter-VM attacks. This paper discusses an innovative solution to overcome this data leakage problem by proposing the requirements for a so-called Conflict-Based Allocation Control for Cloud (CBAC4C) model.

Full Paper: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7038854

Full Reference:

Dlamini, M. T., Eloff, J. H. P. and Eloff, M. M. (2014). CBAC4C: Conflict Based Allocation Control for Cloud, *The 9th International Conference for Internet Technology and Secured Transactions (ICITST-2014)*, London, 2014, pp.: 447-448. DOI: 10.1109/ICITST.2014.7038854

A.5 AN INNOVATIVE RISK-BASED AUTHENTICATION MECHANISM FOR CLOSING THE NEW BANKING VAULT

Abstract: Technology breakthroughs such as online, mobile banking and e-wallet services have empowered today's clients with a full self-service banking platform to access their accounts and transact from anywhere, at any time and using any device. Such innovations have caused a paradigm shift that has seen the new front door of financial banking shifts from physical access branch doors to online banking logical doors. Consequently, this has created a new banking vault which must be secured by strong authentication measures. As banks innovate on their products, the information security community must also innovate around security systems to protect the bank's new services and help prevent unauthorized access. This paper has employed a design thinking methodology to help innovate around authentication measures. The end result is a risk-based multi-factor authentication architecture that enforces strong passwords and provides several other factors such as the variable SurePhrase, geo-location, SIM card serial number, PC serial number etc to authenticate users based on their risk levels. A proof-of-concept is used to demonstrate the actual features of the proposed solution.

Full Paper:

https://books.google.co.za/books?hl=en&lr=&id=-RB2BwAAQBAJ&oi=fnd&pg=PA72&ots=Mj8ZumEJLw&sig=AR1wdl1w5q_pecDNbOR-kqFTzfI#v=onepage&q&f=false

Full Reference:

Dlamini, M.T., Venter, H.S. and Eloff, J.H.P. (2015). An innovative risk-based authentication mechanism for closing the new banking vault, *The Proceedings of the 3rd International Conference on Innovation and Entrepreneurship (ICIE 2015)*, Durban, South Africa, March 2015, pp.: 72 – 80. ISBN: 978-1-910309-91-9

A.6 A CONFLICT-AWARE PLACEMENT OF CLIENT VMS IN A PUBLIC CLOUD

Abstract: The usage and adoption of cloud computing as a private deployment model is continuously improving, regardless of its susceptibility to security issues. This can be attributed to the huge benefits that the cloud provides such as pay-per-use model, quick deployment, turn-around times, huge cost saving, flexible and on-demand self-service provision to cloud users. Since public cloud makes use of virtualization technology, VMs belonging to clients who are in competition may be placed within the same physical infrastructure. This raises the issue around hosting VMs from clients who might be in direct conflict on the same physical infrastructure. Malicious clients could exploit and launch inter-VM attacks to leak confidential information with a competitive advantage. A lot could happen once confidential data is illegally disclosed to unauthorized users. This work makes an attempt to eliminate the confidential data leakage threat posed by inter-VM attacks within the cloud. Hence, it sets itself up to investigate and determine an approach to physically separate potentially conflicting client VMs within the cloud in order to mitigate the confidential data leakage threat posed by inter-VM attacks. In this paper we propose a conflict-aware VM allocation and placement architecture that is implemented with an algorithm modelled using a Chinese Wall Security Policy for physical separation of VMs. The solution is abstracted and applied to different levels of conflict and different levels of the cloud; the data centers, clusters and physical nodes, hence optimizing allocation in terms of conflict of interest. This solution focuses on optimally allocating computing space to client VMs depending on their CoI which then determines the separation distances between conflicting clients' WM. This guarantees that clients who are in direct conflict have their VMs placed very far from each other and VMs belonging to clients that are not in conflict may be placed within the same physical node.

Full Paper:
https://books.google.co.za/books?hl=en&lr=&id=piikBwAAQBAJ&oi=fnd&pg=PA252&ots=EXsPuhAD6A&sig=nELvFcHKqctmtRQg-uK_II2Siuk#v=onepage&q&f=false

Full Reference:

Ratsoma, M.S., Dlamini, M.T., Eloff, J.H.P. and Venter, H.S. (2015). A Conflict-aware Placement of Client VMs in Public Cloud Computing, *Proceedings of the 10th International Conference on Cyber Warfare and Security (ICCWS 2015)*, Kruger National Park, South Africa, March 2015, pp.: 502 - 509.

A.7 STRONG AUTHENTICATION: CLOSING THE FRONT DOOR TO PREVENT UNAUTHORIZED ACCESS TO CLOUD RESOURCES

Abstract: Cloud computing is a computing paradigm where IT resources such as applications, software and hardware are made available over the Internet. However, inadequate authentication in the cloud is one of the major contributing factors to identity theft, leakage of sensitive data and security problems in general. Identity theft and leakage of sensitive data comes with a high risk for breached cloud customers. Such customers could suffer embarrassment, huge financial loss, bankruptcy or even lose their competitive edge. Hence, from the cloud customers' perspective, it is important that the cloud providers have the ability to protect their login credentials, prove their authenticity and prevent unauthorized access to their cloud-based IT resources through world class and proven authentication models, tools and architectures. Hence, this paper proposes strong authentication architecture to mitigate the leakage of sensitive information due to inadequate authentication in the cloud. The proposed architecture seamlessly authenticates end users with a number of attributes such as device ID, geo-location, time of access etc. Our findings reflect that even though strong authentication has been around it is a better solution to prove identity and authenticity of cloud customers, its adoption has been slow. However, current trends indicate that the adoption of strong authentication is on the rise and will become the dominant and best practice authentication scheme in the future for cloud computing. With proper cost, risk and benefit analysis strong authentication can be applied to most cloud providers or any other organizations. The main objective is to ultimately prevent unauthorized access to customer data in the cloud through strong authentication.

Full paper:
https://books.google.co.za/books?hl=en&lr=&id=piikBwAAQBAJ&oi=fnd&pg=PA252&ots=EXsPuhBz9v&sig=JlITxtqW9kEzFBGTsd3cAx75C8Y#v=onepage&q&f=false

Full Reference:

Raphiri, T.V., Dlamini, M.T., Venter, H.S. (2015). Strong Authentication: Closing the Front Door to Prevent Unauthorized Access to Cloud Resources, in *Proceedings of the 10th*

International Conference on Cyber Warfare and Security (ICCWS 2015), Kruger National Park, South Africa, March 2015, pp.: 252 - 260.

A.8 INDUSTRIAL ESPIONAGE: CORPORATE DATA CONTINUES TO LEAK

Abstract: A press release from the White House, dated 25 September 2015, states that the USA and China have agreed to collaborate in a number of global challenges. An agreement was reached that the two respective governments will not knowingly support theft of confidential business information. Access to advanced integrated IT infrastructures such as cloud has empowered many businesses economically but unfortunately also made them vulnerable to threats such as the leakage or theft of any confidential information, including business information. The problem is that it is nowadays much easier for confidential information to leak especially when businesses that are in competition with each other, share the same cloud infrastructure. Furthermore, it is difficult for Cloud Service Providers (CSP's) to control the sharing of resources such as competitors sharing the same physical node in a cloud. The key contributions of this paper are the identification of security challenges, requirements for preventing leakage of confidential business information and an approach to show how the physical placement of a business's data in the cloud can be controlled based on conflict-of-interest classes.

Full Paper:
https://books.google.co.za/books?hl=en&lr=&id=XD7QCwAAQBAJ&oi=fnd&pg=PA90&ots=myp2L7ZIwD&sig=h5WZF-YXaTZcoe4VEPQ6aTwG_9E#v=onepage&q&f=false

Full Reference:

Dlamini, M.T., Eloff, J.H.P. and Eloff, M.M. (2016). Industrial Espionage: Corporate Data Continues to Leak, *The Proceedings of the 11th International Conference on Cyber Warfare and Security (ICCWS 2016)*,Boston University, USA, 17-18 March 2016, pp.: 90 – 97, Edited by Zlateva, D.T. and Greiman, V.A, Boston University, USA. ISBN: 978-1-910810-82-8

A.9 SECURING CLOUD COMPUTING'S BLIND-SPOTS USING STRONG AND RISK-BASED MFA

Abstract: Cloud computing presents an innovative technique to deliver computing as a set of services that can be consumed and utilized over the Internet. This technology breakthrough opens new doors for cybercriminals. The process of authenticating a new breed of cloud users who connect to cloud resources from anywhere, using any Internet-enabled devices and at any time complicates things more. This paper presents an innovative strong and risk-based authentication system. This is to deal with the rising issue related to unauthorized access of cloud hosted resources as a result of inadequate or weak authentication or even stolen user credentials. The proposed solution makes use of a risk engine which monitors user behaviour in order to authenticate users based on specific risk indicators. Furthermore, this paper also makes use of an innovative encryption algorithm that takes chaotic random noise as input to generate encryption algorithms to help encrypt user authentication data at rest, in use and transit. The encrypted authentication data is then stored in multiple storage locations to improve its resiliency to cyber-attacks. This forms a key part of our contribution. The usage of multi-factor authentication systems can be argued to improve the overall system security by monitoring users interacting with a system and modifying defensive strategies to handle malicious behaviour in a proactive manner. Our main goal for this architecture is to try and reduce cyber-criminal activities that are normally caused by weak authentication or stolen user credentials. The proposed solution has already been implemented as a proof-of-concept prototype to evaluate its applicability and suitability to achieve its goals. Despite a 7% of reported false positives or negatives, this solution is guaranteed to take strong risk-based multi-factor authentication in the cloud to the next level.

Full Paper:
https://pdfs.semanticscholar.org/0fd6/ee9480e267d26faf105f9d9cbba159e74ce1.pdf

Full Reference

Dlamini, M.T., Eloff, M.M., Eloff, J.H.P., Venter, H.S., Chetty, K. and Blackledge, J.M. (2016). Securing cloud computing’s blind-spots using strong and risk-based MFA, Proceedings of the International Conference on Information resources Management (CONF-IRM 2016). May 2016, pp.: 1 – 21.

A.10 BEHAVIOURAL ANALYTICS: BEYOND RISK-BASED MFA

Abstract: This paper investigates how to effectively stop an attacker from using compromised user credentials to gain authorized entry to systems that they are otherwise not authorised to access. The proposed solution extends previous work to move beyond a risk-based multi-factor authentication system. It adds a behavioural analytics component that uses keystroke dynamics to grant or deny users access. Given the increasing number of compromised user credential stores, we make the assumption that criminals already know the user credentials. Hence, to test our solution, users were given authentic user credentials and asked to login to our proof-of-concept. Despite the fact that all illegitimate users in our test cases were given the correct user credentials for legitimate users, none of these were granted access by the system. This demonstrates zero-tolerance to false positives. The results demonstrate the uniqueness of keystroke dynamics and its use to prevent users with stolen credentials from accessing systems they are not authorized to access.

Full Paper:
https://pdfs.semanticscholar.org/eb3a/0a26fd9f30a9233846be15897648798a03a9.pdf

Full Reference:
Dlamini, M.T., Eloff, J.H.P., Venter, H.S., Eloff, M.M., Henha Eyono, R.P.S. and Mosola, N.N. (2017). Behavioural analytics: beyond risk-based MFA, *The 2017 Proceedings of the Annual Southern Africa Telecommunication Networks and Applications Conference (SATNAC 2017)*, Freedom of the Seas Cruise Liner, Royal Caribbean International, Barcelona, Spain, 3 – 10 September 2017. ISBN: 978-0-620-76756-9