# Demystifying the risks of public cloud computing

Christopher J. Hodson

## Technical Report

RHUL–ISG–2018–2

2 March 2018

Information Security Group
Royal Holloway University of London
Egham, Surrey, TW20 0EX
United Kingdom

**Student Number: 140134520**

**Christopher J Hodson**

# DEMYSTIFYING THE RISKS OF PUBLIC CLOUD COMPUTING

**Royal Holloway University of London**

**Egham, Surrey, TW20 0EX United Kingdom**

**Supervisor: Dr. Geraint Price**

**Submitted as part of the requirements for the award of the MSc in Information Security at Royal Holloway, University of London.**

**March 2017**

# DECLARATION

I declare that this dissertation is all my own work, and that I have acknowledged all Quotations from the published or unpublished works of other people.

I declare that I have also read the statement on plagiarism in the General Regulations for Awards at Graduate and Masters Levels for the MSc in Information Security and in accordance with it I submit this project report as my own work.

Please sign here to show that you have read the above:

_______________________________________________________ Date: ________________

Christopher John Hodson

# ACKNOWLEDGEMENTS

Tuesday 14th July 2016

Before a word is figuratively penned of this dissertation, I want to start with the most important part. The sincere and unequivocal thanks to those closest to me. Alexandra, Matilda, Mabel and Ralph – you have made enormous sacrifices of your time and energies to support my professional and educational endeavours. You have tolerated hours locked away in my office, at work and / or overseas and I am forever indebted to you all for such a generous, understanding approach to things. I could never dream of being so selfless and I think you are all beautiful. Thank you does not begin to cover it, but it is a good place to start.

## TABLE OF CONTENTS

## TABLE OF FIGURES

## TABLE OF TABLES

## ABSTRACT

Public cloud computing platforms are bringing the benefits of scale, flexibility and cost-effectiveness to organisations of all sizes. Cloud adoption continues to grow in all industry verticals, with technology vendors offering multitenant solutions for infrastructure and line of business applications which were previously only available at the customer's datacentre using physical infrastructure.

There is a perception by many that cloud computing introduces risks to the enterprise. These "risks" being further compounded through the use of public services with tenants from different organisations. In this paper, I will present a thorough analysis of the components of cloud computing with a focus on public cloud. I will characterise the constituent parts of a cloud environment and, through a study of preeminent industry material, identify if and where cloud architecture introduces unique vulnerabilities which lead to new or exacerbated risk.

I will define "information risk" and objectively apply a risk assessment methodology across each of the threat vectors (events) identified as relevant to public cloud. To apply context and qualify a prioritised approach to risk management, I will apply an identical methodology to other common threats which exploit vulnerabilities in people, process and technology.

My findings identify that vulnerabilities exist within the construction of virtualised, multitenant architectures relevant to public cloud; however, most of these vulnerabilities manifest themselves in any technology deployment leveraging contemporary datacentre platforms. Whilst hypervisors and virtualisation introduce technical vulnerabilities, the application of a pragmatic risk methodology identifies that exploitation of these vulnerabilities is unlikely when compared with attacks focused on users and applications.

Public cloud adoption requires an organisation to amend working practices and re-evaluate how security operational and assurance processes are applied and validated. The shared responsibility model, which public cloud inevitably introduces, requires organisations to fully understand "who does what" in relation to security operations. Often cloud providers supply technology but the customer retains operational responsibility.

This thesis includes recommendations for any organisation embarking on a public cloud deployment. I propose a cloud risk metamodel which is the output of my research activity into the components of a cloud risk management framework. I draw on industry-recognised data lifecycle processes and highlight their applicability for public cloud.

I have identified that the public cloud computing does not introduce new types of risks. Cloud architecture carries inherent vulnerabilities but these exist mainly in people and process. Information security controls should be applied commensurate with the sensitivity of the data being stored or transmitted. At no point is organisational accountability outsourced with public cloud adoption.

**KEYWORDS: RISK | PUBLIC CLOUD | MULTITENANCY| VIRTUALISATION | VULNERABILITIES | CONTROLS**

## 1. INTRODUCTION

“Cloud computing” as a concept is as fluffy as its meteorological namesake. “Cloud” means many things to many people. What is the cause of this confusion? Is cloud nomenclature esoteric and incomprehensible?

For decades, technological paradigms have focused on building applications and services locally to the enterprise. We call this model “on-premise”. Such approaches have inherent capital and operational overheads as organisations procure hardware and invest time and effort supporting their infrastructure through software upgrades, configuration and patching. Until recently, organisations had to accept these costs as a practical alternative wasn’t available.

Our contemporary world places less emphasis on the value of ownership; leasing and service-based models are pervasive in all aspects of our lives. Our users are demanding convenience and frictionless interactions with technology. In recent years DVD and CD sales have plummeted, replaced with online streaming services. In 2016, for the first time, Britons spent more on online video than DVDs A spending trend expected to continue over the next five years (**Figure 1-1**).

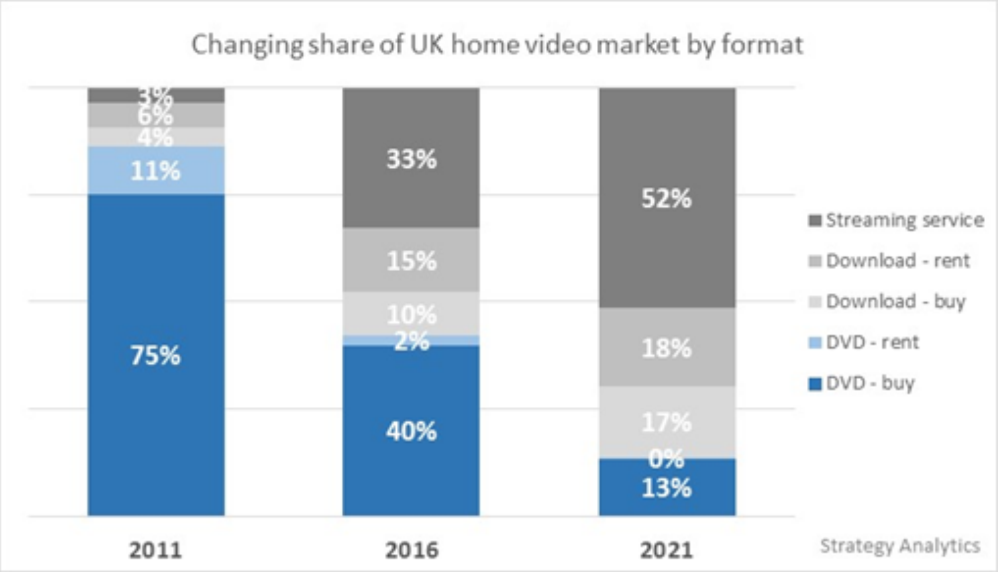

**Figure 1-1 UK Video download growth market [1]**

Text messaging and email usage is declining as real-time messaging platforms are exploding into the marketplace [2]. Established companies are being ousted from their leadership positions by disruptive players who leverage distributed architectures to deliver economies of scale, flexibility and innovation. Companies who stagnate and rest on their laurels are being displaced. We now live in a world where the most valuable public companies are all technology providers [3] who are adopting agile methodologies and reducing time-to-market; cloud plays a significant role in their journey.

Ownership matters little. Uber and Airbnb have transformed the taxi and hotel business respectively. Uber, a company purportedly valued at $66 billion[1] [4] owns not a single taxi car. Airbnb is active in 34,000 cities spanning 191 countries and doesn't own a hotel [5]. In a world where "anything as a service" is ubiquitous, cloud computing has risen to service the cost and efficiency needed to deliver solutions.

Technology is changing our world and cloud computing is at the heart of this revolution. The "as-a-service" model intrinsically associated with cloud computing is a palatable one for business leaders and budget holders: Pay-as-you-go and only for what you use. A potent combination. *"Public cloud computing can avoid capital expenditures because no hardware, software, or network devices need to be purchased. Cloud usage is billed on actual use only and is therefore treated more as an expense. In turn, usage-based billing lowers the barrier to entry because the upfront costs are minimal"* [6]. These lower barriers to entry make public cloud appealing to organisations of all sizes.

The Chief Information Officer (CIO) is under pressure to minimise IT expenditure. The "do more with less" mantra is being articulated from executive boards. This view has been supported by industry reports [7, 8], cloud computing is a high priority for CIOs and so is security.

It is important that cloud computing is understood to be an architectural model and not a panacea to all enterprise IT challenges. *"Cloud is a new way of delivering resources, not a new technology"* [9]. By 2020, more compute power will have been sold by IaaS and PaaS cloud providers than sold and deployed into enterprise data centres, [10] it is, therefore, imperative that we get cloud security right.

Security is a crucial consideration in cloud implementation [11]. Cloud adoption is growing rapidly [12], as our line-of-business applications migrate to the cloud, they are taking the organisational security perimeter with them. Cloud isn't just a new way of deploying infrastructure and applications; it requires a data-centric approach to information and cyber security.

In this thesis, we will review risk lexicon and methodologies associated with information risk management with a view to the application of these methods to the threats and vulnerabilities associated with public cloud. The objective of this activity is to understand if the risks related to public cloud undermine the purported efficiencies and cost savings.

There are inherent benefits for the organisation through cloud adoption but there are also several risks which we will research. As an industry, we frequently speak about the "people, process and technology" of information security [13]. Are any of these more-or-less important when considering cloud? A significant proportion of our peer group is made up of IT professionals and technologists; It is, therefore, natural to focus on the bits and bytes without considering the people and process. My hypothesis being that as organisations introduce cloud-like technologies, their established processes and procedures are altered introducing vulnerabilities.

## 1.1. MOTIVATION

I have spent 18 years working in and around information security. Across this time, I have seen first-hand how our reliance on technology has dramatically changed. Businesses in all industry verticals now rely on computers

[1] Uber's valuation has been the focus of scrutiny from the financial services industry [235, 236]. Whilst $66BN is considered by many to be overinflated; the fact remains that Uber has transformed the automotive industry through their disruptive approach.

for the completion of critical business processes. As our dependency on computers, and the data they produce increases, so does the need to protect the confidentiality, integrity and availability of information.

The IT landscape is always evolving and we, as security professionals, need to adapt. We need to support our businesses in the achievement of their goals. Security has moved from a retrospective tick-box; on the periphery and only consulted when something goes wrong, to an integral service offering business value. We need to establish repeatable methods to assess risk rather than relying on gut-feel and emotion.

Per Gartner [14] "By 2020, a Corporate "No-Cloud" Policy Will Be as Rare as a "No-Internet" Policy Is Today". To support this acceptance of cloud, the security professional needs to rethink how we apply controls. The conventional network perimeter was devised in a time of n-tier architecture and local datacentres (**Figure 1-2**).

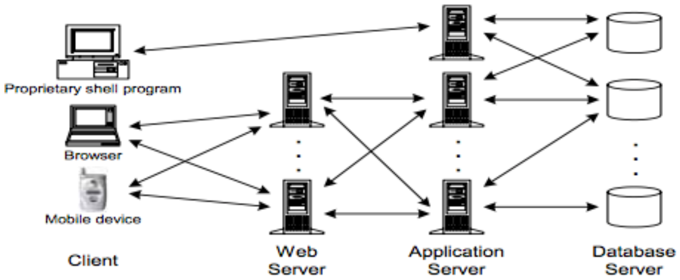

Figure 1-2 N-Tier Architecture [15]

The approach was analogous to a castle and moat defence strategy. As a new threat was discovered, we built a bigger wall or a deeper moat. Cloud facilitates on-demand, ubiquitous access to information. Data centre security worked well when our data was exclusively in locations we controlled. It no longer is [16].

I believe that the subject of cloud security is relevant and of the time. The security community must stop applying the security defences of yesterday to combat sophisticated, user-focused security threats. This thesis will cut through the rhetoric of cloud and offer actionable, pragmatic recommendations for the secure adoption of cloud.

## 1.2. OBJECTIVES

I will provide the reader with a balanced view of the security considerations associated with public cloud computing. At a macro-level, the thesis is written to document the benefits, risks and good practice guidelines when dealing with this new model of computing. This thesis will act as a single repository of information relating to the risks and benefits of public cloud adoption.

The objectives herein provide the reader with direction and allow for a measurable set of goals. As this paper is a comprehensive analysis of public cloud, I will include signposting within the document, using "research questions", which tie-back to the aims outlined in this section.

At a micro-level, a collection of specific objectives has been documented:

| Number | Objectives | Chapter |
|---|---|---|
| 1 | **Understanding Cloud:**<br>a) To establish a common lexicon for cloud computing through an exploration into cloud actors, service/deployment models and reference architectures.<br>b) To understand why cloud computing has become so pervasive in all industry verticals.<br>c) To decide if benefits exist to public cloud adoption, both from a business and a security perspective.<br>d) To assess the impact that public cloud has had on network architecture. | **Chapter 2** |
| **2** | **To understand the security risks associated with cloud:**<br>a) To define the threat actors, (threat) events and vulnerabilities which are associated with public cloud computing.<br>b) To clarify what is meant by the term "risk" in relation to information security and cloud computing. | **Chapter 3**<br>**Chapter 4** |
| **3** | **To perform an analysis of public cloud vulnerabilities:**<br>a) To establish if the risks associated with public cloud computing are unmanageable and/or unpalatable to an enterprise. My hypothesis being that multitenancy is a core consideration for organisations adopting public cloud although most cited cloud vulnerabilities exist across existing computing paradigms.<br>b) To clarify if the risks associated with resource isolation are significantly greater than existing threats and vulnerabilities inherent across enterprise computing.<br>c) To understand if public cloud can be implemented securely for the enterprise.<br>d) To establish if security equivalence can be achieved between on-premise and public cloud architectures. | **Chapter 4**<br>**Chapter 5** |
| **4** | **To understand if established risk management methodologies appropriately cater for public cloud**<br>a) To analyse if respected information risk management frameworks suitably accommodate public cloud.<br>b) To present recommendations, where appropriate, for organisations applying risk management processes to public cloud. | **Chapter 4**<br>**Chapter 5** |

Table 1-1 Research Objectives

## 1.3. SCOPE

This thesis will focus on the security considerations for the adoption of public cloud computing in an enterprise environment. I will identify the service and deployment models associated more generally with cloud although when I address the risks and benefits, “cloud” will be referring to public cloud. “Risk” in the context of this thesis relates to the information assets of an enterprise and the associated confidentiality, integrity and availability (CIA) considerations.

Cloud computing has brought with it many benefits to end-users. Whilst several benefits of cloud apply to the enterprise and the consumer, the user-based benefits have not been analysed specifically. The term “customer” in this paper refers to an organisation consuming public cloud services.

For completeness, the study will identify all cloud service and deployment models although the focus of the study will be around public cloud platforms as these are generally considered to evidence the broadest range of benefits and risks [17]. All service models will be acknowledged and discussed, Infrastructure-as-a-service will be the focus of any technical analysis, the rationale being that a cloud service becomes more vulnerable as more freedom is given to the users since any one of them may be an attacker [18].

With the time constraints of the study, virtualisation and hypervisor discussions will be limited to type 1 and type 2 hypervisor technology. Browser-based sandboxing, containment and virtualisation will not be discussed. Type 2 hypervisors will not be explored in detail as they are generally deployed to support end-user virtualisation requirements and not enterprise hosting [19].

To evidence the maturity and breadth of security controls and compliance capabilities of public cloud providers, I have selected Amazon Web Services (AWS) [20] as a case study. AWS provides a globally distributed framework of cloud services, included IaaS. Based on my experience, AWS have also made the most progress in providing equivalency of capability offering when comparing on-premise and cloud security. AWS have significantly more Virtual Machines (VMs) deployed in production than any other provider [21], something we return to in **Chapter 4**. AWS also leverages a (customised) Xen hypervisor [22] making it a strong candidate for understanding applicable threats and vulnerabilities.

My thesis is outside the length recommendations within Royal Holloway University Project Guide [23]. This was a conscious decision given the research subject I selected. The paper contains regular signposting for the reader and a clear, concise set of objectives have been defined. I assert that one of the major reasons cloud security is such a misunderstood discipline is because very few sources exist which holistically document requisite information within a single paper. This thesis is, therefore, thorough and necessarily longer than most papers at MSc level.

## 1.4. METHODOLOGY AND STRUCTURE

### 1.4.1. METHODOLOGY

This study is based on a comprehensive literature review. I will use various sources including:

- Books
- Research publications
- The Internet
- Vendor literature: surveys, whitepapers, websites
- Conference proceedings

I will draw on my own experience gained through 18 years of working in and around information and cyber security. As a Chief Information Security Officer (CISO), I have extensive experience of the application of risk management frameworks and analysis methods. Where opinion of the author is used, it will be supported with industry and/or other academic sources. Wherever possible, I will support an argument referencing several sources.

I had originally decided to leverage my extensive network of security peers and conduct a survey and interview process uncovering the motivations, priorities and constraints associated with public cloud. Upon reflection, and after assessment, it was decided that similar studies [24, 25] have been conducted across a much broader cross-section of professionals with results which support my preliminary hypotheses surrounding cloud concern and adoption rates. More specifically:

- Adoption of public cloud is growing [26, 27, 28].
- Organisations are increasingly using public cloud for the storage of sensitive information [17].
- Utilisation of "private only" cloud is decreasing in all industry verticals [25].
- Cloud Service Providers are providing commensurate security capabilities to those found across on-premise solutions [29].

Practical application of a selected risk methodology will be conducted in **Chapter 4**.

### 1.4.2. CHAPTER STRUCTURE

This paper is intended to be read from beginning to end and follows a sequential structure where the understanding of a chapter assumes knowledge and context obtained from previous sections. The subject of information risk has a history shrouded in esotericism; something this paper wishes to address. **Chapter 5** can be read in isolation although the reader will benefit from context in earlier chapters. **Chapter 6** completes my thesis and whilst it could be read in isolation, the rationale to support several my assertions will not be included.

Each chapter will begin with a chapter introduction which will detail the theme and structure of that chapter. **Chapters 2-4** will conclude with a conclusion section which aims to provide closure on the sub-plot of the chapter and a definitive view of the author based on his research. **Chapter 5** covers recommendations and a conclusion is not required. **Chapter 6** is my thesis evaluation and close.

In **Chapter 2**, the definition of cloud will be investigated. Based on a literature review including industry and academia, this section will define a consistent cloud parlance which will be used across subsequent chapters. The objective of the chapter is to establish a common lexicon for cloud computing and to understand "what is cloud?" Cloud architecture will be analysed; the chapter will review the essential characteristics of cloud computing. I will compare the various service and deployment models associated with cloud computing.

**Chapter 3** begins with a look at what "risk" means from several lenses across society. We will explore definitions of information risk before deciding upon an agreed definition to take into an assessment of cloud risks. The chapter will proceed to analyse established information security risk frameworks with a view to understanding their suitability for cloud. This chapter will aim to decompose the essential characteristics of an information risk equation including threat actors, (threat) events, vulnerabilities and controls. An important deliverable of this chapter is the documentation of "cloud-centric" vulnerabilities and threat events.

In **Chapter 4**, my research is focused on understanding what is meant by the term "multitenancy". Multitenancy is a foundational component of cloud and is a fundamental shift-away from ownership. In this chapter, we will apply a qualitative risk analysis to the threat events associated with multitenancy and finish by applying the same risk analysis to some common threat events and vulnerabilities with a view to contextualising the likelihood of multitenancy vulnerabilities being exploited.

**Chapter 5** will provide a series of recommendations for an organisation considering public cloud for their applications and/or infrastructure. In this chapter, I will present a "*cloud risk metamodel*" which details the actors and requisite components of a public cloud risk model. I include observations to improve the secure adoption of public cloud. An aim of this chapter is to understand if security conservation can be achieved between security services on-premise and within public cloud; this directly aligns to the objectives of this thesis. I will close the chapter with a "*security conservation process flow*".

**Chapter 6** will form the conclusion to this thesis and is entitled *"Conclusions: Public Cloud is not a Technology Problem"*. This chapter is a review of my research findings and, to close, a presentation of a "*Top 10 Public Cloud Considerations*".

### 1.4.3. DOCUMENT FORMATTING

This thesis has been constructed in line with Royal Holloway University London (RHUL) requirements [23]. As both electronic and physical versions will be reviewed, I have endeavoured to ensure readability and simplified document navigation across both formats.

The following formatting will be used throughout the paper:

**Bold:** Used for explicit cross-referencing throughout the paper. **Chapters**, **figures**, **tables, questions** and **objectives** will appear in **bold**. Title sections in tables will be **bold**.

*Italics:* Direct quotations will be presented inline within quotation marks and with italic notation.

Underline: Areas of emphasis or attention will be identified through underlined text.

Hyperlinks: Hyperlinks will only be included to assist with document navigation. All Figures and tables will have bold, hyperlinked, inline references. Chapter references will contain hyperlinked numbers[2].

[2] Microsoft Word 2016 creates a formatting error when adding bold hyperlinks if exported to .pdf. If you are reading this paper in .pdf form, figures and tables will not have bolded numbers. This is not an oversight and is unavoidable.

# 2. CLOUD COMPUTING

*"If someone asks me what cloud computing is, I try not to get bogged down with definitions. I tell them that, simply put, cloud computing is a better way to run your business" [30]*.

**Marc Benioff, Founder, CEO and Chairman of Salesforce**

What is cloud and how does it differ from previous computing models? This chapter will explore what I call "Cloud DNA"; the metaphorical organisms that combine to make a cloud-based eco-system. I will start by researching the definition of cloud computing before reviewing the datacentre architectures which have brought us to where we are today. In this chapter, I will study the core components of cloud computing. Cloud architecture is reviewed in this chapter along with a look at the available academic research in the space of cloud taxonomy. The chapter will conclude with a look at how cloud computing is driving changes in computer networking.

## 2.1. DEFINITION OF CLOUD

There are many different definitions of cloud computing. A recent analogy I used in a speaking engagement [31] was that clouds are like roads; they facilitate getting to your destination. Be that destination a network location, an application or a development environment. No one would enforce a single, rigid set of rules and regulations for all roads - many factors come into play: volume of traffic, the likelihood of an accident, safety measures, requirements for cameras. If all roads carried a 30 mile an hour limit, you might reduce fatal collisions but freeways and motorways would cease to be efficient. Equally, if you applied a 70 mile an hour limit to a pedestrian precinct, unnecessary risks would be introduced. Context is very important, imperative in fact. The same goes for cloud computing. If we are to establish the risks, benefits and suitability of cloud, we need to better define the various cloud deployment, operational and service models in use today and looking forward.

A universally accepted definition of cloud is not available although some sources are considered authoritative and appear more frequently. Cloud is a disruptive[3] and pervasive technology movement; it is therefore understandable that consensus on definition is not possible; academia, business and cyber security vendors are all approaching cloud computing from different perspectives.

Across my research, the four most referenced definitions are that of the National Institute of Standards and Technology (NIST), the Cloud Security Alliance (CSA), European Union Agency for Network and Information Security (ENISA) and Gartner. NIST [32] defines cloud computing as: *"…a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction".*

The CSA is *"the world's leading organisation dedicated to defining and raising awareness of best practices to help ensure a secure cloud computing environment"* [33]. The CSA takes the NIST definition of cloud although interestingly elaborates to include the benefits of cloud: *"Cloud computing is a disruptive technology that has*

[3] In a technology sense, disruption is brought about through the creation of a new market and the challenging of existing ways of working.

*the potential to enhance collaboration, agility, scaling, and availability, and provides the opportunities for cost reduction through optimized and efficient computing"* [34].

ENISA [35] define cloud as: *"…an on-demand service model for IT provision, often based on virtualization and distributed computing technologies".* Gartner [36] goes with: *"A style of computing where scalable and elastic IT-related capabilities are provided 'as-a-service' to external customers using Internet technologies"*.

Cloud adoption continues to grow [37], as it does, such an explicit delineation of cloud and on-premise will not be necessary. Is the world of commodity computing displacing traditional datacentre models to such an extent that soon all computing will be elastic, distributed and based on virtualisation? Will all computer access be service-based and ubiquitous? In **Section 2.5**, the fundamental components of cloud are defined. These features are becoming increasingly common across on-premise and public cloud deployments rendering our existing definitions of "cloud" confusing and problematic.

As technology moves '"to the cloud", the definition of what constitutes a computer becomes opaque; some suggesting that computers will become "invisible" [38]; users will not realise they're accessing billion-dollar computer networks in remote data centres, technology is becoming intrinsically integrated into our lives and firms are striving to deliver frictionless, intuitive experiences. Artificial Intelligence (AI) platforms are providing overlays and support mechanisms for the completion of daily tasks; the cloud is key to the delivery of these services although neither the client or server components fall into existing definitions of computers.

## 2.2. THE EVOLUTION OF INFRASTRUCTURE AND JOURNEY TO CLOUD

The Greek philosopher Heraclitus once said: "*life is flux*" [39] which translated into contemporary language means "*the only constant is change*". This is certainly true of the information technology world where our modern digital ecosystems are barely recognisable from those of yesteryear. **Figure 2-1** visualises the timeline provided by SiliconANGLE [40] and shows the transformational journey our IT landscape has undergone.

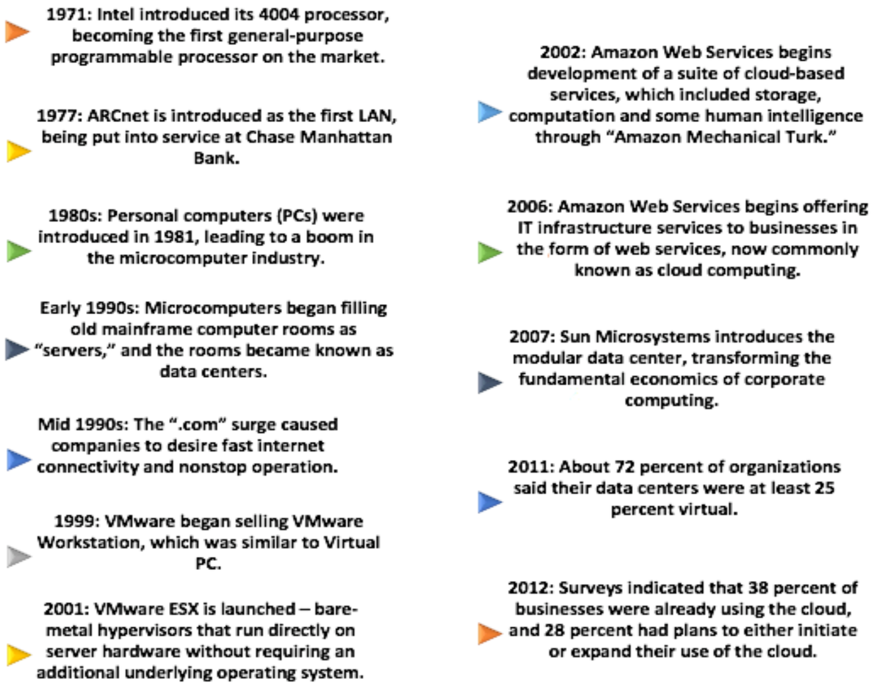

1971: Intel introduced its 4004 processor, becoming the first general-purpose programmable processor on the market.

1977: ARCnet is introduced as the first LAN, being put into service at Chase Manhattan Bank.

1980s: Personal computers (PCs) were introduced in 1981, leading to a boom in the microcomputer industry.

Early 1990s: Microcomputers began filling old mainframe computer rooms as "servers," and the rooms became known as data centers.

Mid 1990s: The ".com" surge caused companies to desire fast internet connectivity and nonstop operation.

1999: VMware began selling VMware Workstation, which was similar to Virtual PC.

2001: VMware ESX is launched – bare-metal hypervisors that run directly on server hardware without requiring an additional underlying operating system.

2002: Amazon Web Services begins development of a suite of cloud-based services, which included storage, computation and some human intelligence through "Amazon Mechanical Turk."

2006: Amazon Web Services begins offering IT infrastructure services to businesses in the form of web services, now commonly known as cloud computing.

2007: Sun Microsystems introduces the modular data center, transforming the fundamental economics of corporate computing.

2011: About 72 percent of organizations said their data centers were at least 25 percent virtual.

2012: Surveys indicated that 38 percent of businesses were already using the cloud, and 28 percent had plans to either initiate or expand their use of the cloud.

Figure 2-1 Computing Paradigm Shift

NIST asserts that cloud computing is a model – a way of working, an enabler [41]. The sharing of a configurable set of resources. While cloud is the latest step on the data centre transformation journey, the sharing of resources is something which has been achieved through many varied paradigms of the past.

Cloud computing is unlike any technological model before it; a salient consideration being the location of information. Historically, all approaches have relied on the assurances, valid or otherwise, of the physical data centre residency. Organisations could apply a combination of logical and physical security controls to protect their data assets.

Cloud computing requires an institution to apply the concept of trust, allowing a third-party to manage data on their behalf. At first glance, this unfamiliar approach sounds radical and dangerous although pragmatism and context suggest otherwise. Companies have been relying on third parties to manage information for centuries; the difference with cloud computing is that the information is in digital form. The subject of multitenancy will be discussed extensively throughout this thesis but it is important to remember that multitenancy has existed across all phases of supply chain in perpetuity.

## 2.3. BENEFITS OF CLOUD

### 2.3.1. BUSINESS BENEFITS

I feel it is import to cover the business benefits of cloud. This will help to explain the meteoric rise of cloud computing. IT services are there to support business operations. Cloud adoption should not be a unilateral technology decision.

The NIST Cloud Computing Reference Architecture (NCCRA) [41] states that *"a primary focus of the cloud computing model is on the economic benefits of shared use that can provide higher-quality and faster services at a lower cost to users."* These economic benefits are not solely financial. Cloud computing provides businesses with economies of scale – computing resources are acquired for the task in hand. Organisations are not forced to provision expensive, redundant infrastructure to accommodate periods of peak activity which may only occur sporadically. The concept of "capital expenditure" (organisations procuring hardware and suffering depreciation) can be removed from the IT equation although "operational expenditure'" is a consideration as cloud models incur periodical service costs.

Cloud computing allows organisations to adopt flexible technologies which lower time-to-market. Virtualised and centralised environments can be provisioned and logically segregated to create eco-systems which used to be possible only through complex and expensive physical infrastructures. Cloud computing enables development and testing teams to create "crash and burn[4]" components in an expedient fashion, limiting the need for rogue-deployments and proliferation of "Shadow IT[5]". "Failing fast" is a term used extensively in the modern IT world with Gartner asserting that *"…smart organisations will embrace fast and frequent project failure in their quest for agility"* [42].

Cloud brings some significant user experience benefits. User can access applications from anywhere around the globe through commodity Internet connections. As covered in my introduction, the service-based consumption

[4] Environments that can be provisioned swiftly and in a ring-fenced ecosystem removing the fears of test and development environments impacting the confidentiality, integrity or availability of production systems.

[5] Shadow IT is a general term covering unsanctioned IT applications and environments.

model of cloud has completely transformed markets such as music and video where Amazon [43], Spotify [44] and Netflix [45] have cornered an industry previously dominated by CD purchases and DVD rental. Cloud computing is an essential vehicle for the delivery of our "Internet of Things" (IoT) world. Our homes and offices are run, managed and secured via Smart Devices and sensors which communicate to centralised infrastructure on the Internet. Either directly or indirectly, IoT and cloud computing are opening doors to advancements in our business efficiency [46], our health [47, p. 8] and our transportation needs [48].

Business continuity and disaster recovery processes can be improved through cloud adoption. Organisations can scale highly-redundant, geographically-dispersed architectures without the need for capital expenditure. Failing over between cloud instances removes the need for cold / warm standbys and allows an enterprise to pay for what they need, when they need it. Multitenant architectures, strategically-positioned around the global can improve the reliability and availability of service for the enterprise customer.

The green factor should not be overlooked. Cloud adoption brings with it several environmental benefits. A 2013 study by the Global E-Sustainability Initiative [49] identifies that found if 80% of enterprises adopted cloud computing we would see approximately a 4.5 megaton reduction in greenhouse gas emissions (based on conservative estimates. To put this figure in perspective, a 4.5 megaton reduction in greenhouse gas emissions is roughly equivalent to taking 1.7 million cars off-the-road.

### 2.3.2. SECURITY BENEFITS

*"Cloud computing is often far more secure than traditional computing, because companies like Google and Amazon can attract and retain cyber-security personnel of a higher quality than many governmental agencies" [50].*

**Vivek Kundra, Former Federal CIO of the United States of America**

The business benefits of cloud appear to be well understood [51]. What seems paradoxical to some are the inherent security benefits of cloud computing. It is natural to think that because an IT team procures and manages their infrastructure and applications, they will do a better job that an outsourced provider. Across my research, it has become clear that the areas of security benefit draw direct parallels with the business benefits of cloud computing: Scale, flexibility, cost savings and improved performance.

In the interests of avoiding scope creep, this section will not include a detailed analysis into the security benefits of cloud although several of these are becoming increasingly important with the ever-evolving threat landscape. A precis of these will be included below:

#### 2.3.2.1. RISK AVERSION

Public cloud computing providers apply "paranoia by default" [52]. A publicised data breach or exploited vulnerability could destroy a cloud provider overnight. Their appetite for risk is very low. Risk appetite is *"…the amount of risk, on a broad level, an organisation is willing to accept in pursuit of value…"* [53]. For a Cloud Service Provider (CSP), customer confidence and assurance is of paramount importance. The customer must have trust and confidence that the CSP is protecting their most sensitive assets. Reputation is key: "*Trust is between two entities; but the reputation of an entity is the aggregated opinion of a community towards that entity"* [54]. Cloud service providers need to instil and retain positive reputation across the general population. This societal reputation metric enforces a low tolerance to risk. The risk tolerance threshold of a CSP hosting thousands of customers is much lower than any individual organisation. A breach at a client who, for example, trades stocks or sells groceries, can have a negative effect on stock prices and customer confidence although effects are often short term [55] . For a CSP, a loss in consumer confidence can have an irreparable impact.

#### 2.3.2.2. SCALE

Through public cloud adoption, organisations benefit from the economies of scale afforded to a global, multitenant, service-based consumption model. An often-overlooked benefit is the application of the same principles to security services.

As covered in **Section 2.3.2.1**, CSPs have a low tolerance for information risk. Patch management, virtualisation, hypervisor and network security operations management are all a top priority for all mature CSPs. Customers benefits from the security measures that these companies deploy globally.

The protection inherently available to protect the operation of the cloud[6] significantly assist the customer organisations in their operational security strategy. Services which, as appliances, are cost prohibitive and require expensive and onerous deployment models, can be deployed in all locations, often simply with the click of a button.

For the protection of applications and individual VMs, CSPs offer virtualised, elastic security services which would not be possible without the breadth, scale and capacity of public cloud. AWS Marketplace [56] provides organisations with a comprehensive set of security capabilities to prevent and detect cyber-attacks and mitigate the risk of a data breach.

Redundancy and availability of data are core tenants of information security which are improved through public cloud. Mature providers offer the ability for customers to provision infrastructure and applications in multiple geographies (availability), often with several cloud tenants in the same geographical location (redundancy).

In our "post-Snowden" world, the demand for consumer privacy is at an all-time high. Privacy and the need for encryption are intrinsically linked; the former US Director of National Intelligence asserting that "*As a result of the Snowden revelations, the onset of commercial encryption has accelerated by seven years"*. [57] The use of encryption-in-transit[7], and more specifically, Transport Layer Encryption (TLS) [58] for web communication is on the rise. A recent report from the network provider Sandvine [59] suggesting that in Europe, approximately two-thirds of all traffic (both fixed line and mobile) is encrypted (**Figure 2-2**). A trend supported by Mozilla [60] who recently announced that, on average, over 50% of their Firefox browser traffic was encrypted.

Encryption has undeniable benefits for the user although, in almost all instances, it presents a significant overhead in terms of performance and cost for the security department. On-premise appliances were deployed at a time when encryption was reserved for rare use cases and regulatory compliance. With encryption being ubiquitously deployed and malware actors increasingly leveraging encrypted channels for attacks [61], the need to inspect encrypted traffic is undeniable. Multitenant, public cloud architecture can significantly relieve the cost and performance constraints of encrypted traffic inspection. Cloud security providers are offering solutions which can perform TLS inspection at micro-second latency without the need to locally deploy appliances [62]. Reports suggest that encrypted traffic is continuing to rise with over 75% of web traffic predicted to be encrypted by 2019 [63]. Cloud-based solutions provide the scale, coverage and performance benefits that allow organisations to retain visibility of their sensitive information.

[6] A reference to the delineation of responsibility between a customer and a cloud provider.

[7] Ensuring confidentiality (and optional integrity) of data communicated between an endpoint and a server.

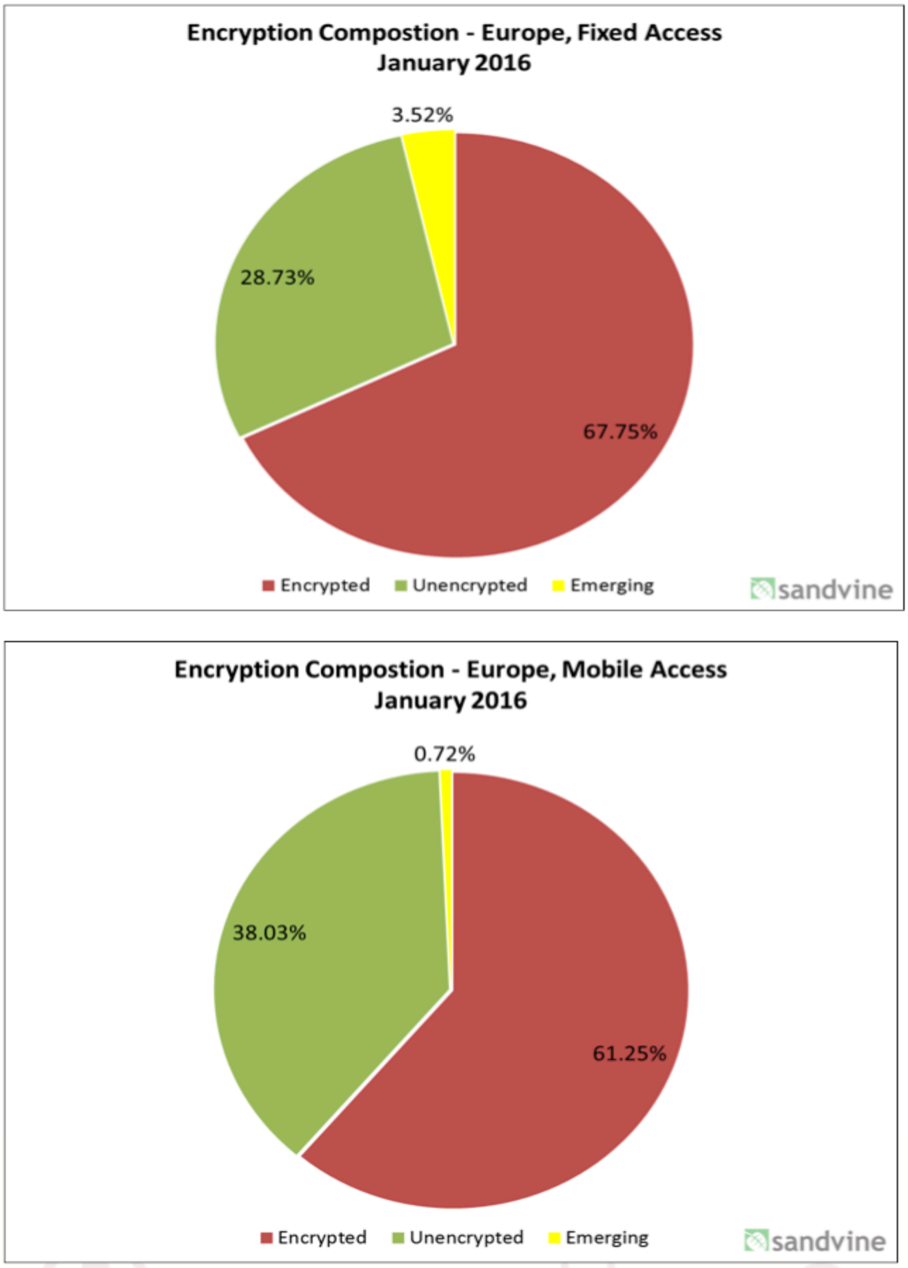

**Figure 2-2 European Encrypted Traffic Figures 2016 [59]**

### 2.3.2.3. VISIBILITY AND COVERAGE

In almost all situations, public cloud adoption brings with it an improved visibility of the threats and traffic patterns traversing the Internet. For example; Zscaler, the world's largest security-as-a-service cloud provider, reports seeing approximately 30 billion HTTP transactions daily [62]. When a threat is detected, it is blocked for the originating organisation but also for all other subscribers to the service.

The malware of today is increasingly polymorphic [8] and sophisticated. Organisations need to be leveraging a globally dispersed network of threat intelligence information to provide immediate protection against threats that are programmed to change construction in seconds rather than weeks [64]; rendering conventional, on-premise antivirus engines ineffective in detecting and responding to sophisticated malware campaigns. Cloud coverage mitigates these threats by provides global distribution of threat intelligence information which simply isn’t possible with an insular, organisationally-maintained cyber defence structure.

#### 2.3.2.4. REGULATORY AND LEGAL COMPLIANCE

Legal and compliance concerns are often cited as “cloud risks” [65]. It is also worth considering the benefits that public cloud computing brings to the compliance and legal conversation.

Developed cloud platforms which are serving enterprise customers understand the need to provide environments capable of achieving cross-industry compliance against common frameworks. Amazon Web Services [20] (AWS) provides comprehensive accreditation against industry-standard compliance and legal frameworks (**Figure 2-3**).

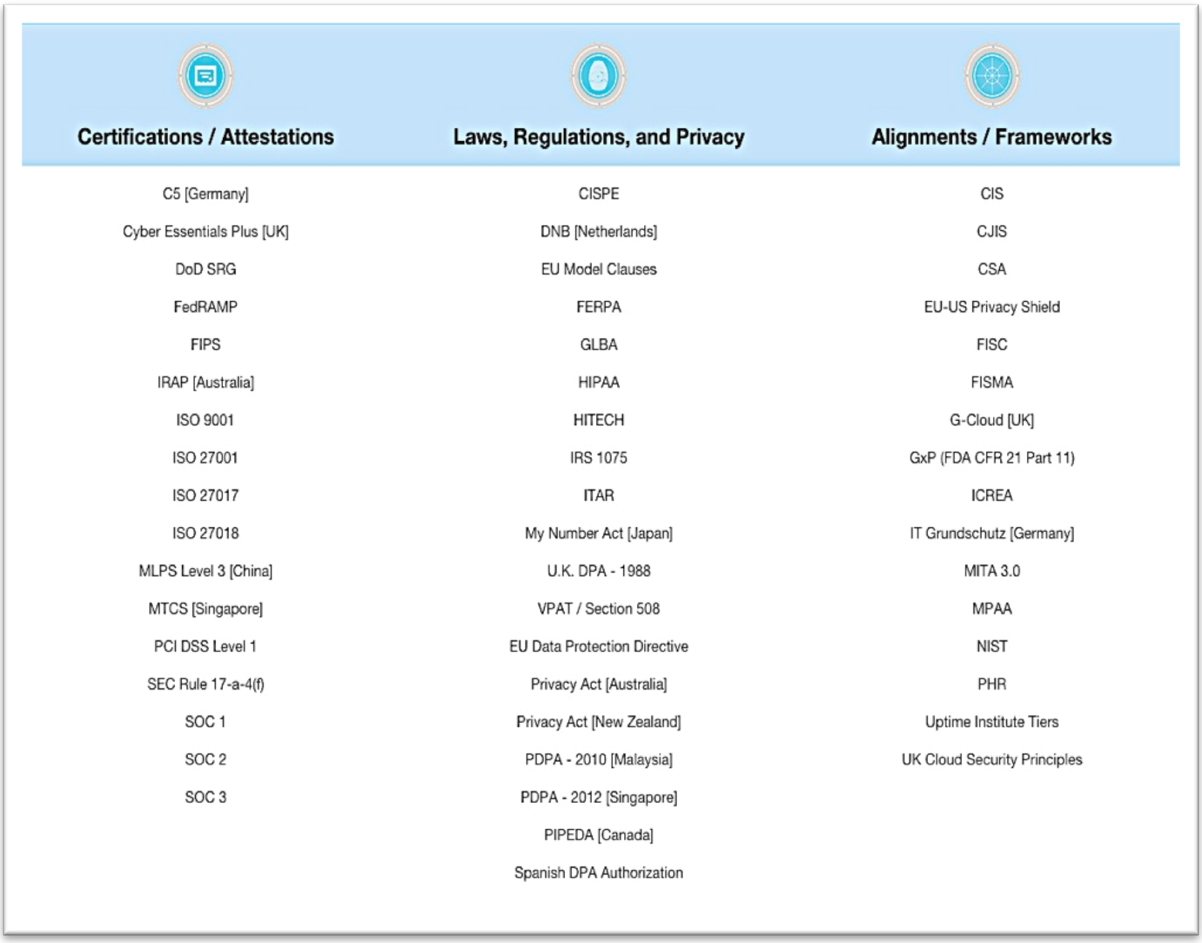

| Certifications / Attestations | Laws, Regulations, and Privacy | Alignments / Frameworks |
|---|---|---|
| C5 [Germany] | CISPE | CIS |
| Cyber Essentials Plus [UK] | DNB [Netherlands] | CJIS |
| DoD SRG | EU Model Clauses | CSA |
| FedRAMP | FERPA | EU-US Privacy Shield |
| FIPS | GLBA | FISC |
| IRAP [Australia] | HIPAA | FISMA |
| ISO 9001 | HITECH | G-Cloud [UK] |
| ISO 27001 | IRS 1075 | GxP (FDA CFR 21 Part 11) |
| ISO 27017 | ITAR | ICREA |
| ISO 27018 | My Number Act [Japan] | IT Grundschutz [Germany] |
| MLPS Level 3 [China] | U.K. DPA - 1988 | MITA 3.0 |
| MTCS [Singapore] | VPAT / Section 508 | MPAA |
| PCI DSS Level 1 | EU Data Protection Directive | NIST |
| SEC Rule 17-a-4(f) | Privacy Act [Australia] | PHR |
| SOC 1 | Privacy Act [New Zealand] | Uptime Institute Tiers |
| SOC 2 | PDPA - 2010 [Malaysia] | UK Cloud Security Principles |
| SOC 3 | PDPA - 2012 [Singapore] | |
| | PIPEDA [Canada] | |
| | Spanish DPA Authorization | |

**Figure 2-3 AWS Compliance and Legal Position [66]**

[8] Malware which can “morph” making it difficult to detect through signature-based anti-malware technologies

A message that I will return to throughout this paper is that organisations do not outsource accountability when adopting public cloud. As Amazon state: *"While AWS manages security of the cloud, security in the cloud is the responsibility of the customer. Customers retain control of what security they choose to implement to protect their own content, platform, applications, systems and networks, no differently than they would in an on-site data centre."* [66].

IT compliance programmes are expensive and time-consuming. Through the adoption of public cloud, organisations can inherit environments which are already certified against compliance and legal frameworks although it is important that the applications and services deployed by the customer satisfy the requirements of individual laws and regulations. We will further explore Amazon's "Shared Responsibility Model" in **Chapter 5**.

## 2.4. CLOUD ACTORS

Whilst many varied definitions of cloud computing exist, little has been written in terms of defining the stakeholders involved across a cloud computing eco-system. The NCCRA [41] defines five actors with actor defined as *"…an entity (a person or an organization) that participates in a transaction or process and/or performs tasks in cloud computing."*

**Table 2-1** outlines the actors defined by NIST, names and descriptions are taken verbatim:

| **Cloud Actor** | **Description** |
|---|---|
| Consumer | A person or organization that maintains a business relationship with, and uses service from, *Cloud Providers*. |
| Provider | A person, organization, or entity responsible for making a service available to interested parties. |
| Broker | An entity that manages the use, performance and delivery of cloud services, and negotiates relationships between Cloud Providers and Cloud Consumers. |
| Carrier | An intermediary that provides connectivity and transport of cloud services from Cloud Providers to Cloud Consumers. |
| Auditor | A party that can conduct independent assessment of cloud services, information system operations, performance and security of the cloud implementation. |

**Table 2-1 Cloud Actors [32]**

Actors in an ecosystem of any description need to communicate. NIST provides a communication flow which identifies relationships between cloud actors. It is important to understand communication flows if we are to be able to define roles and responsibilities accordingly. **Figure 2-4** identifies communication paths between cloud actors.

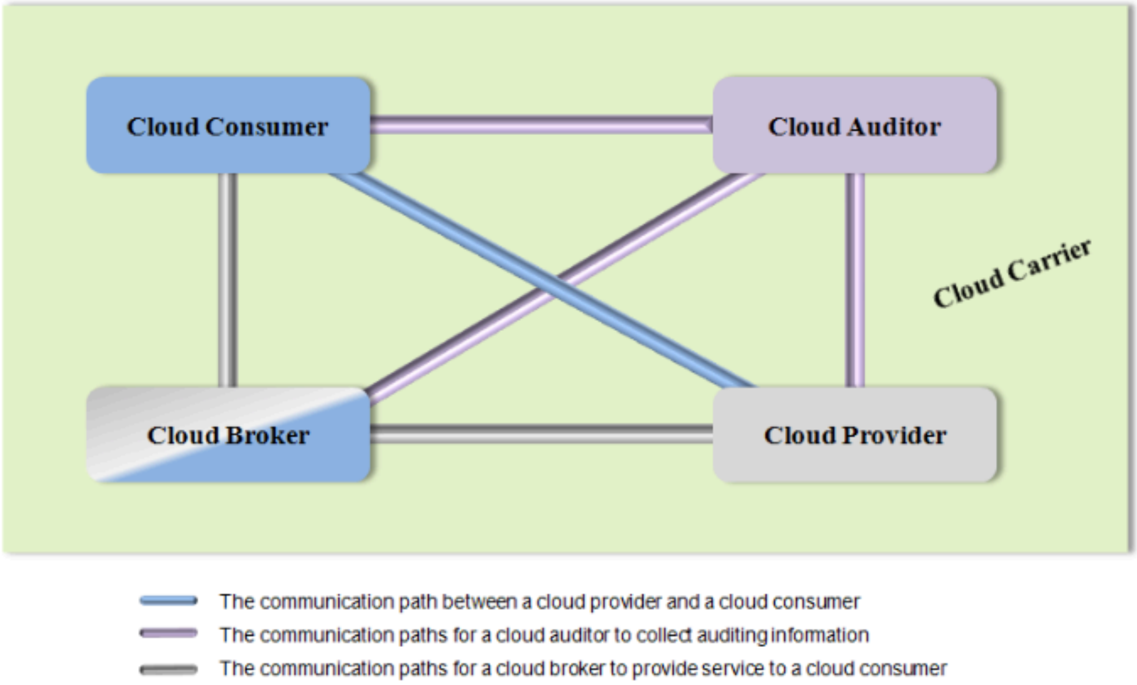
Cloud Consumer
Cloud Auditor
Cloud Carrier
Cloud Broker
Cloud Provider
The communication path between a cloud provider and a cloud consumer
The communication paths for a cloud auditor to collect auditing information
The communication paths for a cloud broker to provide service to a cloud consumer

Figure 2-4 Actor Communication Paths [32]

The blue line in **Figure 2-4** identifies the flow between the two most common entities in a cloud model: a consumer (or customer) talking to a provider (CSP). Brokers play an increasingly important role in cloud adoption. The Cloud Access Security Broker (CASB) is a term defined by Gartner [67] and describes a form of security enforcement between a customer and a CSP. The term is interpreted in different ways across our industry although CASB functionality often includes:

- Data Loss Prevention (DLP)
- Cloud Application Visibility
- Application API Integration
- Identity Management

CASB functionality is used as a security control to mitigate the risks associated with "shadow IT" (discussed in **Section 3.9**).

The auditing of cloud services can be performed through similar means to on-premise deployments and take the form of technology, legal & regulatory and process controls.

It is helpful to see these interactions across a logical architecture. The Open Security Architecture (OSA) produced a "cloud computing pattern" [68] which includes identical actors to those defined by NIST. The OSA extends the schema through the inclusion of architect, end-user, developer, IT Manager actors; variations of the "consumer" class and "Auditor", an instance of the Auditor category. (**Figure 2-5**). The OSA Cloud Security Pattern assists in defining actors but also the required security capabilities and services for public cloud adoption.

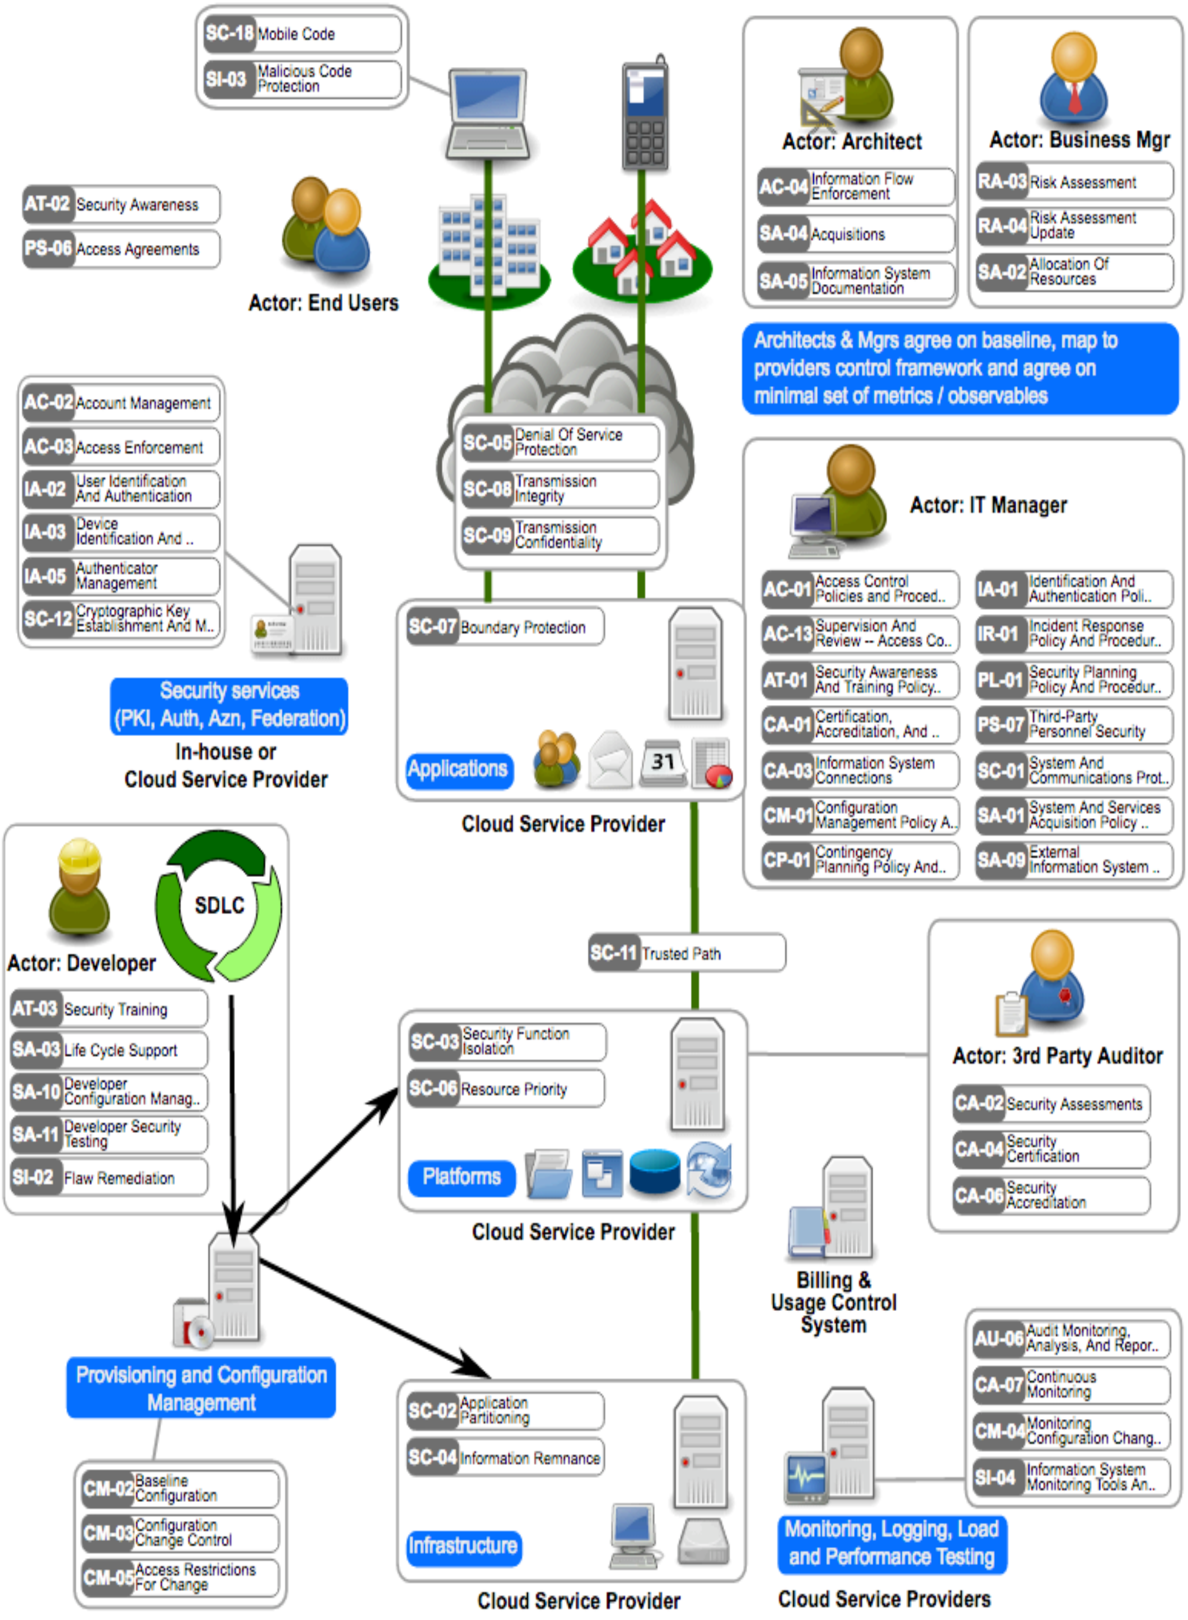
SC-18 Mobile Code
SI-03 Malicious Code Protection
AT-02 Security Awareness
PS-06 Access Agreements
Actor: End Users
Actor: Architect
AC-04 Information Flow Enforcement
SA-04 Acquisitions
SA-05 Information System Documentation
Actor: Business Mgr
RA-03 Risk Assessment
RA-04 Risk Assessment Update
SA-02 Allocation Of Resources
Architects & Mgrs agree on baseline, map to providers control framework and agree on minimal set of metrics / observables
AC-02 Account Management
AC-03 Access Enforcement
IA-02 User Identification And Authentication
IA-03 Device Identification And ..
IA-05 Authenticator Management
SC-12 Cryptographic Key Establishment And M..
Security services (PKI, Auth, Azn, Federation)
In-house or Cloud Service Provider
SC-05 Denial Of Service Protection
SC-08 Transmission Integrity
SC-09 Transmission Confidentiality
SC-07 Boundary Protection
Applications
31
Cloud Service Provider
Actor: IT Manager
AC-01 Access Control Policies and Proced..
AC-13 Supervision And Review -- Access Co..
AT-01 Security Awareness And Training Policy..
CA-01 Certification, Accreditation, And ..
CA-03 Information System Connections
CM-01 Configuration Management Policy A..
CP-01 Contingency Planning Policy And..
IA-01 Identification And Authentication Poli..
IR-01 Incident Response Policy And Procedur..
PL-01 Security Planning Policy And Procedur..
PS-07 Third-Party Personnel Security
SC-01 System And Communications Prot..
SA-01 System And Services Acquisition Policy ..
SA-09 External Information System ..
SDLC
Actor: Developer
AT-03 Security Training
SA-03 Life Cycle Support
SA-10 Developer Configuration Manag..
SA-11 Developer Security Testing
SI-02 Flaw Remediation
SC-11 Trusted Path
SC-03 Security Function Isolation
SC-06 Resource Priority
Platforms
Cloud Service Provider
Actor: 3rd Party Auditor
CA-02 Security Assessments
CA-04 Security Certification
CA-06 Security Accreditation
Billing & Usage Control System
Provisioning and Configuration Management
CM-02 Baseline Configuration
CM-03 Configuration Change Control
CM-05 Access Restrictions For Change
SC-02 Application Partitioning
SC-04 Information Remnance
Infrastructure
Cloud Service Provider
AU-06 Audit Monitoring, Analysis, And Repor..
CA-07 Continuous Monitoring
CM-04 Monitoring Configuration Chang..
SI-04 Information System Monitoring Tools An..
Monitoring, Logging, Load and Performance Testing
Cloud Service Providers

Figure 2-5 OSA Cloud Security Pattern [68]

## 2.5. ESSENTIAL CHARACTERISTICS

NIST [41] define five essential characteristics for cloud computing. These are not specific to public cloud and should be considered across all deployment models. The NIST definitions were drafted in 2011 and whilst still relevant in 2017, technological advances mean that a more contemporary lens should be applied to certain definitions. Where necessary, I will explicitly reference these in the following subsections:

### 2.5.1. ON DEMAND, SELF-SERVICE

The ability for the user to unilaterally provision resources: compute, memory, storage, network configuration or user access. The unilateral consideration is key and carries with is an enormous user benefit but also a significant security consideration. On demand, self-service significantly reduces time to market. The consumer no longer requires the engagement of the service provider for provisioning. Change tickets and helpdesk calls are a thing of the past[9]. Expediency is provided, processes are streamlined. Whilst the benefits are clear, we posit that this characteristic carries with it one of most significant objection to cloud. Unilateral provisioning is not always carried out by approved IT or security teams. Anyone with a credit card and an Internet connection can procure resources and as the colloquialism goes "spin them up in seconds". Such expediency carries with it concerns around data leakage and "shadow IT" architectures.

### 2.5.2. BROAD NETWORK ACCESS

Ubiquitous connectivity is a key security benefits of cloud computing. Broad network access is the ability for users to connect to their resources anywhere, anytime. Cloud computing needs repeatable, standards-based methods of network access which support heterogeneous endpoints. NIST breaks these systems into thin and thick clients [32].

In a world of the consumerisation of IT and the pervasiveness of the "Internet of Thing"' (IoT) [69], we are seeing a myriad of devices connecting to cloud-based resources. In my opinion, the explicit delineation of "thick" and "thin" requires revision; historically, there was an implicit assumption of client-based security. A set of controls, rudimentary or otherwise, to protect the confidentiality, integrity and availability of data. IoT carries no such assurances with vulnerabilities resident in most IoT hardware [70]. IoT is the first technology to be considered by many "insecure by default" [71] and with over 20 billion IoT devices estimated to be in production by 2020 [69], we must consider broad network access to consider Smart Devices, IoT and mobility.

By NIST's definition, for a solution to be considered "cloud", it must support broad network access; this is a view that I share although "broad" and "ubiquitous" should not be confused with "unfettered". As we will see in later chapters, cloud solutions must retain an ability to enforce access control accordance with the OSA pattern outlined in **Section 2.4**.

### 2.5.3. RESOURCE POOLING

We take the NIST definition verbatim:

*"The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned per consumer demand" [32].*

[9] In relation to provisioning services.

Resource pooling is intrinsically-associated with the ability for Cloud Service Providers (CSPs) to provide high-performance, cost-effective computing; consumers reap the benefits of scale and elasticity. The concept of resource sharing is a significant hurdle in cloud adoption – as covered in **Section 2.2**, cloud is the first computing model whereby it is both conceivable and likely that different organisations will be housing data on the same physical server hardware.

It is my hypothesis that resource pooling not only forms an essential characteristic of cloud computing but it is also the most significant security consideration. I will aim to prove across this thesis that the vulnerabilities associated with resource sharing can be suitably mitigated through the application of controls.

### 2.5.4. RAPID ELASTICITY

Amazon [72] presents the limitations of traditional infrastructure provisioning in their best practice guide. These are broken-down into two deployment options:

**Scaling Up:** This method ignores a scalable application architecture and scaling is achieved through investment in larger, more powerful computers to accommodate load. This approach is heavily capex dependent and depending on application usage, demand could quickly outgrow capacity.

**Scaling Out:** A horizontally-scaled architecture. Application components are shared across multiple physical components and following a service-orientated design. A requirement to understand workload capacity still exists with this model resulting in infrastructure being deployed and underutilised.

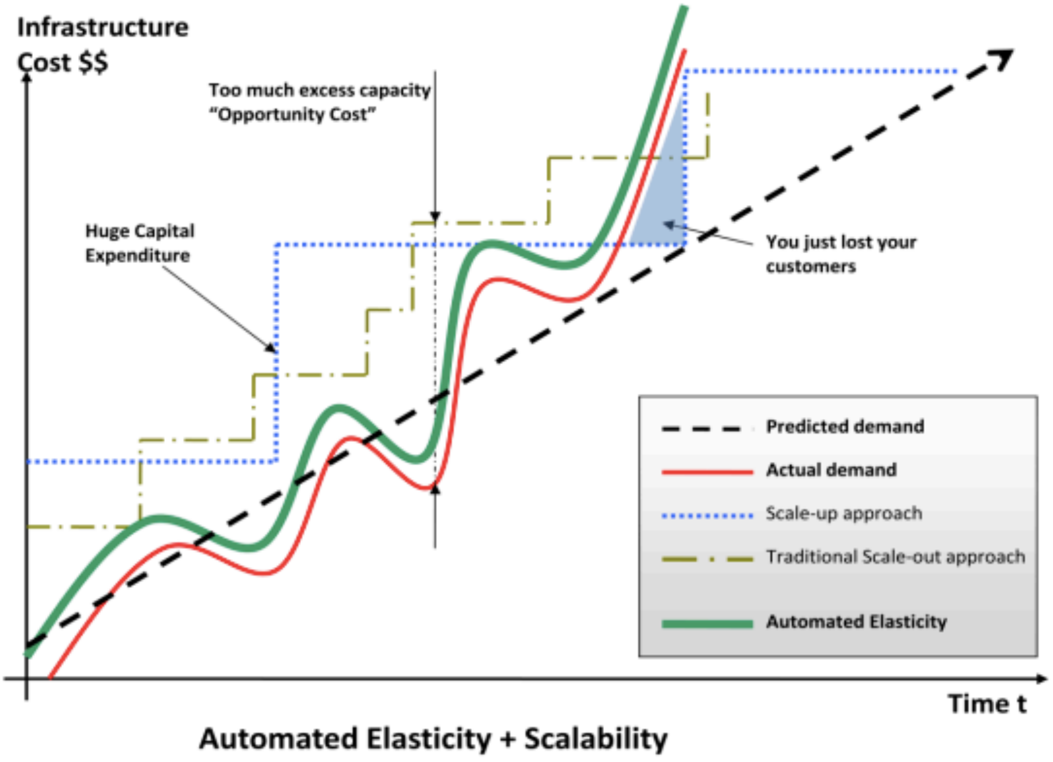

Figure 2-6 Elasticity and Demand Scaling [72]

Cloud Elasticity is *"the degree to which a system is able to adapt to workload changes by provisioning and de-provisioning resources in an autonomic manner, such that at each point in time the available resources match the current demand as closely as possible"* [73].

The service-based model of cloud is critically important for businesses of all sizes. On-demand scaling removes the requirement for an enterprise to procure server computing resources in anticipation of peak periods. This approach was accepted as necessary in computing paradigms of the past.

### 2.5.5. MEASURED SERVICE

Organisations in all industry verticals are required to deliver cost-effective, technology platforms. Measured service allows the enterprise to pay for what they use. NIST [32] says "*Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service*". Measured service provides the IT department, and subsequently budget holders, the visibility of precisely what services are being paid for.

### 2.5.6. MULTITENANCY AND VIRTUALISATION

Whilst NIST does not explicitly define multi-tenancy as an essential characteristic, the Cloud Security Alliance (CSA) [34] refers to multitenancy as *"…in its simplest form implies use of same resources or application by multiple consumers that may belong to same organization or different organization"*. Resource pooling is said to leverage a "multitenant" model. Whilst the statement doesn't define multitenancy as a mandatory requirement, it acknowledges the importance of the characteristic. Although multitenancy and virtualisation are not required characteristics per NIST, several high-profile vulnerabilities are as a direct result of multitenancy. **Chapter 4** will discuss multitenancy and its associated flavours.

## 2.6. SERVICE MODELS

As stated in the introduction to this thesis, cloud means many things to many people. An objective of this paper is to be very clear on the service models available for cloud computing. Customer functionality and security requirements generally drive the decision of which service model is adopted. There are three established services models in existence today with a fourth, catch-all, model appearing in recent years.

### 2.6.1. SOFTWARE AS A SERVICE

Software as a service (SaaS) is the service model which offers the least flexibility to the customer but conversely is generally the easiest to deploy. With SaaS, an organisation selects a CSP to provide an application function to users. The CSP is responsible for the provisioning of infrastructure (compute, memory, network) and for the operational activities which were previously tasked to the organisation when their application lived on-premise.

The list of SaaS applications appears endless. Two very successful examples of SaaS applications are Microsoft Office 365 [74] and SalesForce.com [75]. Microsoft financial results for 2016 [28] report that monthly commercial users of the platform are now over 85 million which is a 40% year-on-year rise. The number of purchased seats and revenue all significantly up on the previous year. Adoption of Office 365 suggests a profound shift in the acceptance of enterprise public cloud. FY 2016 Annual Reports [26] show that Salesforce.com achieved $6.67 billion in revenue with a 24% revenue growth for the year.

The growth of Salesforce and Office 365 is relevant to this study because both platforms deal with information which is considered commercially-sensitive to any organisation. At its core, Salesforce is a customer records management (CRM) platform. It houses personally identifiable information (PII) and data pertaining to orders, quotes and other commercially sensitive records. Office 365 is being adopted by organisations as a replacement for on-premise email, file storage and collaboration. Customers are making a conscious shift to public cloud for the storage of their most sensitive information [17].

### 2.6.2. PLATFORM AS A SERVICE

Platform as a Service (PaaS) is defined by NIST [32] as *“the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services and tools supported by the provider”*.

PaaS affords the consumer more control and flexibility than a SaaS implementation. The consumer is responsible for selecting the applications deployed to the platform and the associated configuration of the environment. The CSP retains responsibility for manging the hardware and OS configuration.

PaaS has been an integral part of the DevOps [10] world we find ourselves in today. Development and Innovation teams are leveraging the agility benefits of PaaS and relying on automated configuration management systems such as Chef [76] and Puppet [77] which both treat system, application and security configurations as code enabling them to be deployed consistently and rapidly from a central console.

PaaS should be considered as an enterprise development and deployment ecosystem: *“Like IaaS, PaaS includes infrastructure – servers, storage and networking – but also middleware, development tools, business intelligence (BI) services, database management systems and more. PaaS is designed to support the complete web application life cycle: building, testing, deploying, managing, and updating”* [78].

### 2.6.3. INFRASTRUCTURE AS A SERVICE

In an IaaS deployment, compute storage and network resources are owned and hosted by a service provider and delivered to the customer. Of the service models available today, Infrastructure as a Service (IaaS) affords the customer the greatest flexibility but with this comes the highest level of responsibility for security. The consumer is responsible for the provisioning of resources for their virtual machines: storage, compute, network and memory. This is not to be confused with the physical installation of hardware. In an IaaS deployment, the customer selects the amount of resource she requires for the workload in question.

The operating system management is the responsibility of the consumer and with it comes obligation for patching and account/service management. When selecting service models, it is important that a mutual understanding of responsibilities exists between the service provider and the customer (**Figure 2-7**).

[10] DevOps is a way of working that emphasises the benefit of cross-team and cross-functional working for the delivery of IT services and platforms. Agile, fluid methodologies are embraced. Security should be part of the cross-functional team.

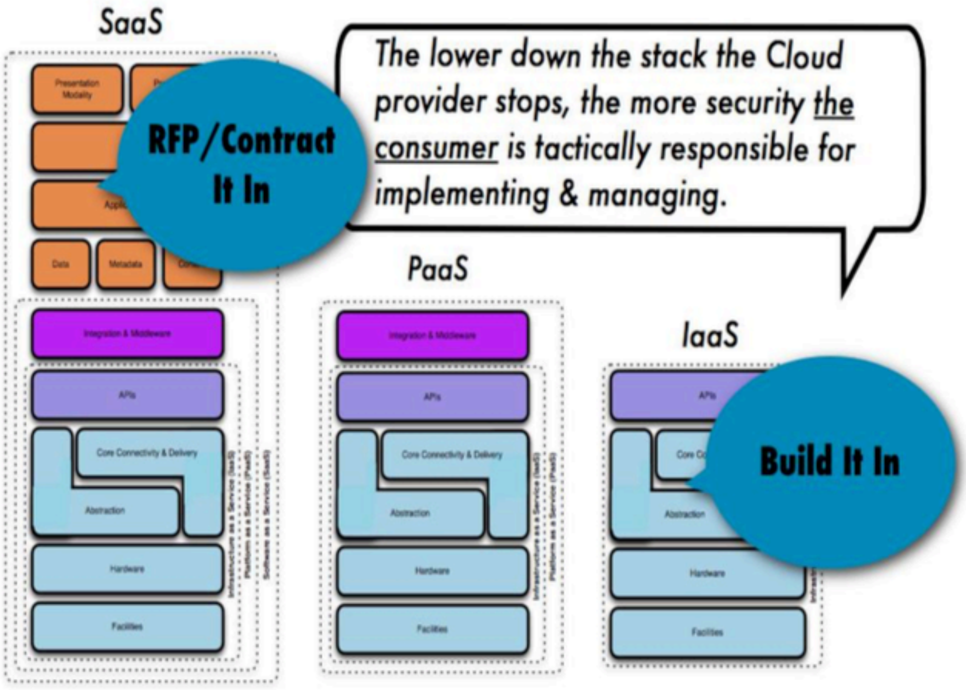
SaaS
RFP/Contract It In
The lower down the stack the Cloud provider stops, the more security the consumer is tactically responsible for implementing & managing.
PaaS
IaaS
Build It In
Presentation Modality
Data
Metadata
Integration & Middleware
APIs
Core Connectivity & Delivery
Abstraction
Hardware
Facilities

Figure 2-7 Service Models: Responsibilities [34]

### 2.6.4. ANYTHING AS A SERVICE

Anything/Everything as a service (XaaS) is a catch-all model which comes from the fact that the economies of scale, performance and elasticity of cloud are driving innovative ways of thinking for the provisioning and deployment of technology solutions which were previously reserved for on-premise appliances. Virtualisation and infrastructure convergence introduces the concept of commodity-based computing [79]; organisations are purchasing business-services and forgoing the flexibility of selecting hardware, software flavours and configuration options for the benefits of cost, consolidation and convenience.

XaaS can be used to describe any other service-based deployment which doesn't fall within the "big three" (IaaS, PaaS, SaaS) [80]. Common examples that I have encountered are included in **Table 2-2**:

| **XaaS Variant** | **Description** |
|---|---|
| Desktop as a Service | Delivery of virtualised desktop operating systems from the cloud. |
| Disaster Recovery as a Service | Outsourcing of all DR-related functionality to the public cloud. |
| Identity as a Service | Open Authorization (OAuth)/Security Assertion Markup Language (SAML)-based federation solutions along with cloud-based directory services. |

Table 2-2 XaaS Variants

| | |
|---|---|
| Storage as a Service | Cloud-based storage platforms. |
| Security as a Service | Delivery of security capabilities in the cloud: Network Function Virtualisation (NFV), cryptography-as-a-service, Web/Endpoint security as a service. |

**Table 2-2 (Cont.) XaaS Variants**

## 2.7. DEPLOYMENT MODELS

Today, four deployment models exist for cloud computing. It is my opinion that these are self-explanatory and can be defined in table format:

| **Deployment Model** | **Description** |
|---|---|
| Public | For use by the public. Located in datacentres accessible via public IP networking and DNS namespaces. The focus of this paper. |
| Private | Provisioning exclusively for a specific organisation. Datacentres can be located on-premise or off-premise. |
| Hybrid | The composition of two or more infrastructures (Public, Private, Community). |
| Community | For the use of a specific collection of consumers, generally with a shared interest. |

**Table 2-3 Cloud Deployment Models [32]**

## 2.8. CLOUD REFERENCE ARCHITECTURE AND TAXONOMY

Architecture is defined by the British Standards Institute (BSI) as the *"fundamental concepts or properties of a system in its environment embodied in its elements, relationships, and in the principles of its design and evolution"* [81]. It is not only the individual components of cloud we must understand but the relationships between these components. Architecture provides the organisation with a map, a view of their technology "on a page". It is imperative that as organisations begin to adopt third party services (CSPs) for the storage and processing of their information that they understand the data flows, systems and interfaces that will be managing these interactions.

### 2.8.1. CLOUD REFERENCE ARCHITECTURE FRAMEWORKS

The BSI definition of architecture [81] is abstract enough to cover all forms of systems architecture although the constituent parts of a cloud deployment differ significantly from legacy technology paradigms.

Much like an industry ambivalence on the meaning of cloud computing, several architectural approaches exist which cover the concepts, actors and capabilities of cloud. I believe that traceability is imperative when assessing the efficacy of an architectural model. For a cloud architecture to be effective, we must be able to trace business requirements through to technical controls. Several frameworks exist which satisfy this requirement when addressed together. These artefacts are covered in the following subsection(s):

#### 2.8.1.1. NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY

The NCCRA [41] is a conceptual model which is generic in nature. It was produced to provide US federal government with an effective tool for understanding the terminology, components and operations of cloud computing. The conceptual nature of the model is important; it does not provide any implementation guides and is vendor agnostic. As the guide states *"The NIST cloud computing reference architecture focuses on the requirements of "what" cloud services provide, not a "how to" design solution and implementation"* [41].

**Figure 2-8** presents an overview of the NIST Cloud Computing Reference Architecture. As you can see: actors, functional, service and deployment models are defined:

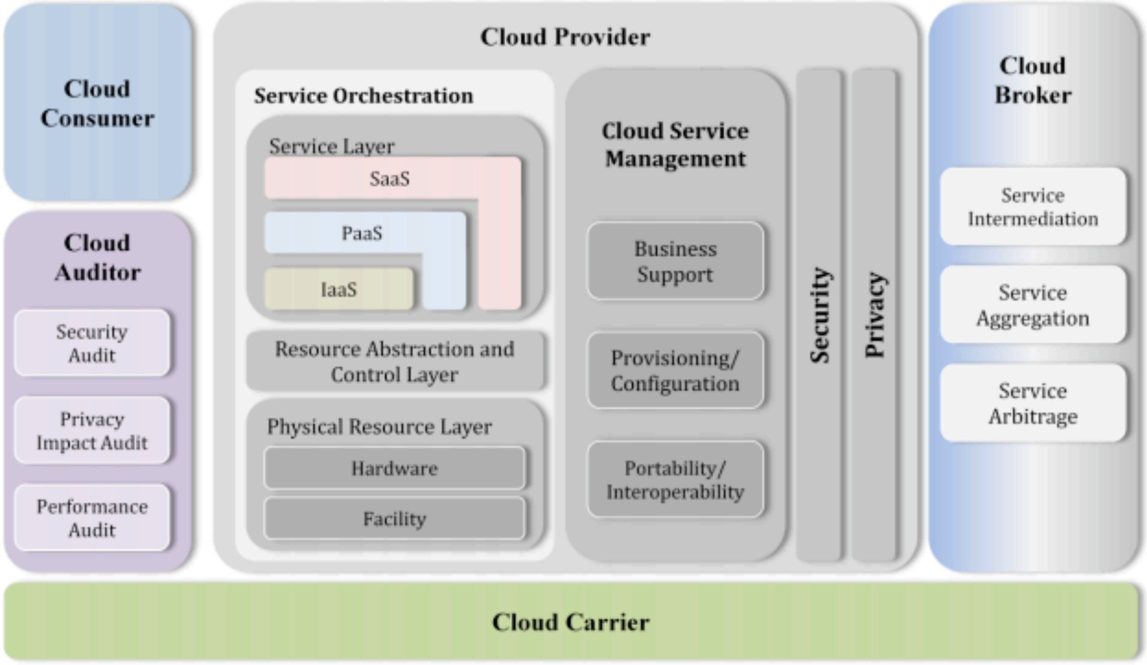

Figure 2-8 NIST Cloud Computing Reference Architecture [41]

The NCCRA defines actors and their relationships, it was not designed to cover security requirements. The NCCRA should be thought of a blueprint or a "conceptual architecture" in Sherwood Applied Business Security Architecture (SABSA) parlance [82] although the requirements for security are not defined. The NCCRA defines the constituent components of a cloud ecosystem.

Within "The NIST Definition of Cloud Computing" [32], the components of cloud are broken down as follows:

- The three types of service models defined in the NIST Reference Architecture: IaaS, PaaS, and SaaS;
- The four types of deployment models defined in the NIST RA: Public, Private, Hybrid, and Community; and
- The five Actors defined in the NIST RA: Provider, Consumer, Broker, Carrier and Auditor.

NIST published a Cloud Computing Security Reference Architecture (NCCSRA) to provide a *"…comprehensive formal model to serve as security overlay to the architecture described in NIST SP 500-292: NIST Cloud Computing Reference Architecture"* [83].

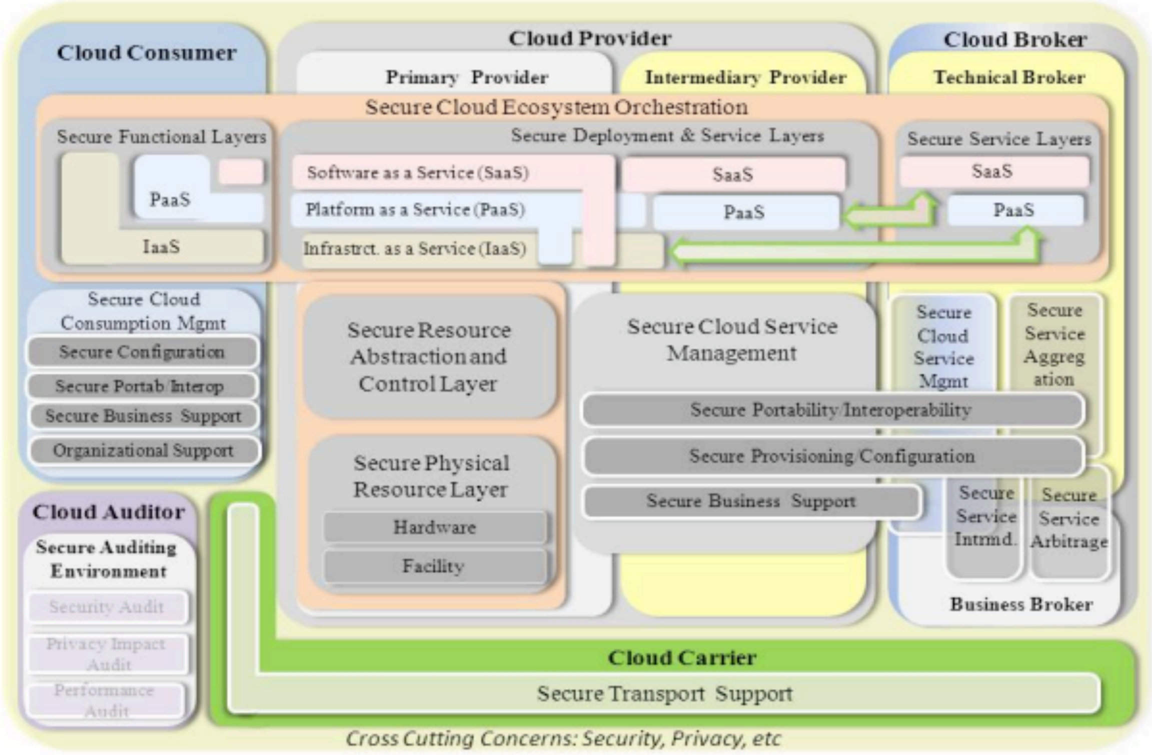
Cloud Consumer
Cloud Provider
Primary Provider
Intermediary Provider
Cloud Broker
Technical Broker
Secure Cloud Ecosystem Orchestration
Secure Functional Layers
PaaS
IaaS
Secure Deployment & Service Layers
Software as a Service (SaaS)
Platform as a Service (PaaS)
Infrastrct. as a Service (IaaS)
SaaS
PaaS
Secure Service Layers
SaaS
PaaS
Secure Cloud Consumption Mgmt
Secure Configuration
Secure Portab/Interop
Secure Business Support
Organizational Support
Secure Resource Abstraction and Control Layer
Secure Physical Resource Layer
Hardware
Facility
Secure Cloud Service Management
Secure Portability/Interoperability
Secure Provisioning/Configuration
Secure Business Support
Secure Cloud Service Mgmt
Secure Service Aggregation
Secure Service Intrmd.
Secure Service Arbitrage
Business Broker
Cloud Auditor
Secure Auditing Environment
Security Audit
Privacy Impact Audit
Performance Audit
Cloud Carrier
Secure Transport Support
Cross Cutting Concerns: Security, Privacy, etc

Figure 2-9 NIST Cloud Computing Security Reference Architecture [83]

The NCCSRA retains the actors, sever and deployment models but overlays a set of core security requirements. The paper does not prescribe technologies nor logical security capabilities. Much like the NCCRA, references are aligned to a conceptual architecture framework.

This document details a set of processes for applying a Cloud-adapted Risk Management Framework (CRMF). The CRMF defines a set of "*security components*" which are derived from the CSA Enterprise Architecture [84] (discussed in the next section) and through leveraging the NIST Risk Management Framework: 800-37 [85].

### 2.8.1.2. THE CLOUD SECURITY ALLIANCE

The CSA Enterprise Architecture is considered both a methodology and a set of tools to assist an organisation in their understanding of roles and responsibilities when deploying cloud services. The Enterprise Architecture defines a set of security capabilities and depending on the service and deployment models of the cloud application, it presents a set of control objectives for both the customer and the cloud service provider (CSP).

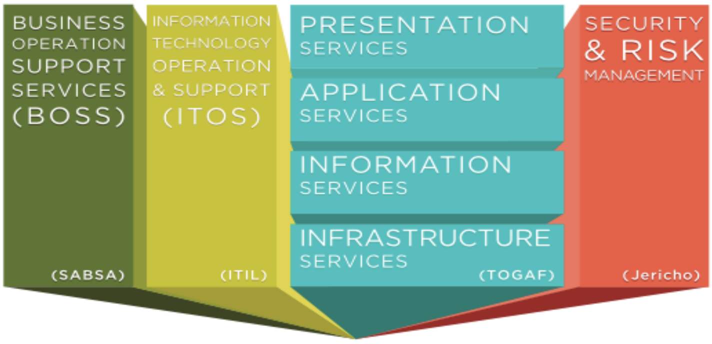
BUSINESS OPERATION SUPPORT SERVICES (BOSS)
(SABSA)
INFORMATION TECHNOLOGY OPERATION & SUPPORT (ITOS)
(ITIL)
PRESENTATION SERVICES
APPLICATION SERVICES
INFORMATION SERVICES
INFRASTRUCTURE SERVICES
(TOGAF)
SECURITY & RISK MANAGEMENT
(Jericho)

Figure 2-10 CSA Enterprise Architecture [84]

The CSA draws upon other established frameworks and architectural models to provide comprehensive coverage of the architectural requirements associated with cloud computing. The CSA Enterprise Architecture is broken down into "business domains" which are included below in **Table 2-4**:

| **Business Domain** | **Framework/Model** | **Explanation of Framework** |
|---|---|---|
| Business Operation Support Services (BOSS) | SABSA | Business-aligned architecture model to deliver security architectures at an enterprise and a solutions level. |
| Information Technology Operation and Support | ITIL | IT Service Management Framework used to align services with business requirements. |
| Presentation, Application, Information and Infrastructure Services | TOGAF | An enterprise architecture framework which follows a sequential and interactive lifecycle model for the deployment of solutions. |
| Security and Risk Management | Jericho Forum | Forum setup to address the issues associated with organisational deperimeterisation.<br>The forum closed in October 2013 citing that "deperimiterisation is now an established fact" [86]. A pertinent statement in our exploration of cloud pervasiveness. |

Table 2-4 CSA Enterprise Architecture Domains [84]

The CSA documentation provides multi-layered architectural guidance; the Enterprise Architecture allows a customer or a service provider to ascertain which actor is responsible for which requirements in a cloud deployment. Unlike NIST, the CSA provides a collection of security controls which are aligned across several industry and regulatory compliance frameworks through the CSA Cloud Controls Matrix (CCM) [87] (which we will return to in **Section 3.7**). The CSA model works because it follows an inclusive model, consulting multiple stakeholders.

#### 2.8.1.3. ITU-T – CLOUD COMPUTING REFERENCE ARCHITECTURE

The ITU-T Y.3502, Cloud Computing Architecture, August 2014 [88] provides a four-layer reference architecture for cloud computing. Having researched several architectures and frameworks for this thesis, it is my opinion that ITU-T.3502 provides a model most useful for an enterprise or solutions security architect. I draw this conclusion as the ITU-T.3502 model is multi-dimensional. It includes layers (or tiers) with which parallels can be drawn with n-tier architecture and it also documents a series of functional controls and capabilities which should be applied across functional business units within the enterprise. **Figure 2-11** details the layers and functions outlined in this section.

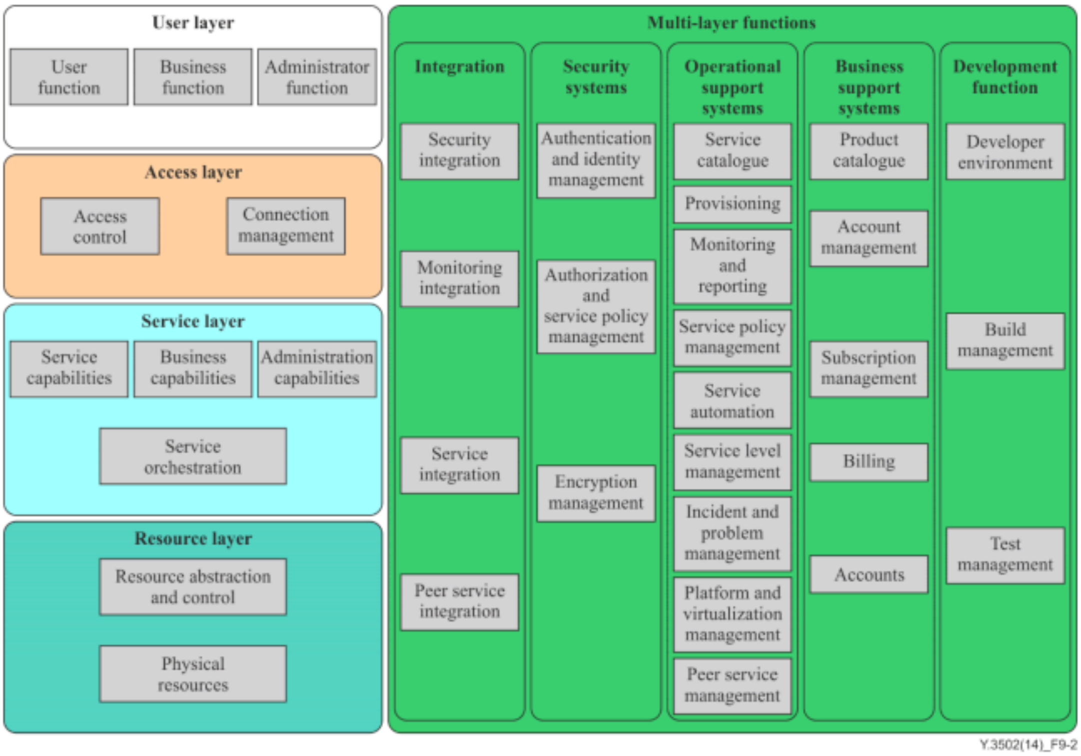

Figure 2-11 ITU-T Y.3502 Reference Architecture [88]

The purpose of this section is to identify the improvements our industry is making to provide cloud-first reference artefacts. It is my opinion that no framework could be adopted in isolation and provide comprehensive guidance for the adoption of public cloud. Having said this, all the above referenced frameworks provide a security specialist with technology-agnostic support and I encourage any architect looking at a public cloud migration to review each artefact.

### 2.8.2. TAXONOMY

A taxonomy is *"the study of the general principles of scientific classification"* [89]. Taxonomies are not new to information architecture or database design. SABSA [82] introduces the concept of taxonomies to define business attributes. The SABSA model is used within the CSA Enterprise Architecture to define business operation support services (**Table 2-4**).

Given that we are without definitive agreement on all forms of cloud nomenclature, a cloud computing taxonomy seems prudent. I wanted to understand how much material existed which around cloud taxonomies. Academic efforts have been made to define a "taxonomy of cloud computing services" (**Figure 2-12**) [90] although research is limited.

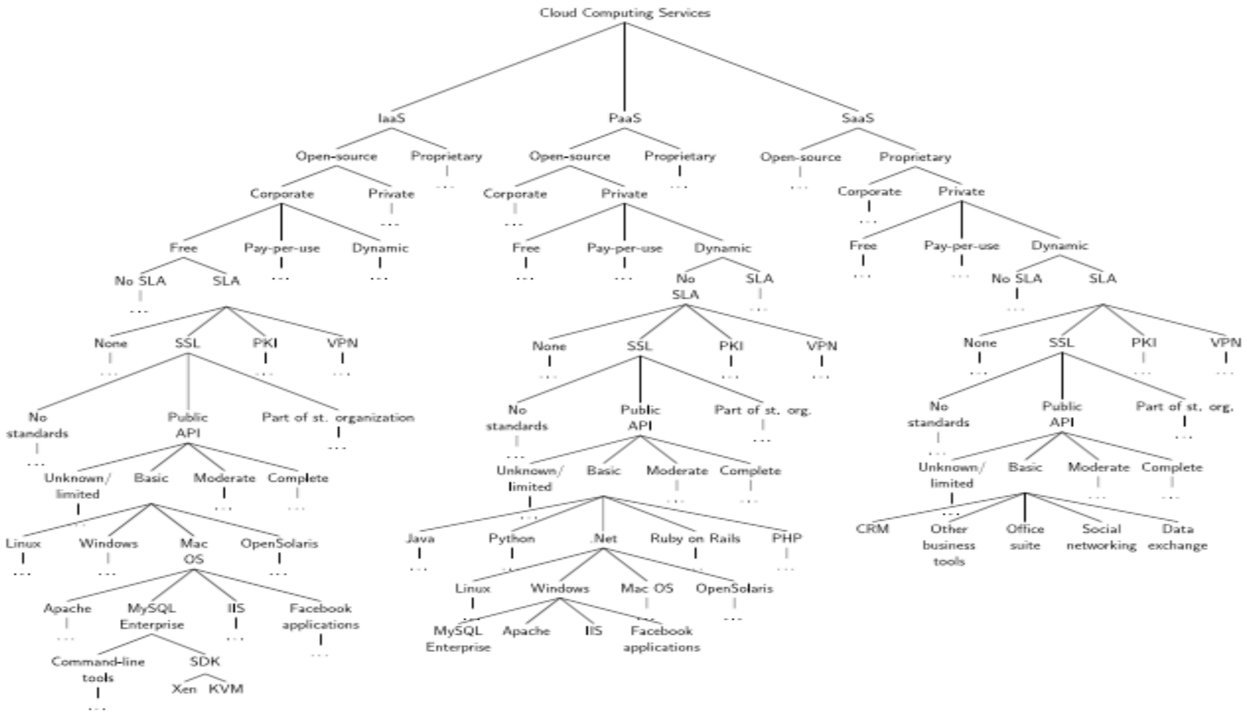

Figure 2-12 Taxonomy of Cloud Computing Services [90]

Microsoft [91] have produced a cloud taxonomy which delineates not only the required components of a cloud deployment but also the areas of responsibility between the consumer and the CSP in the case of IaaS, PaaS and SaaS. Responsibility and accountability are areas that the author will explore across **Chapter 3** and **Chapter 5**.

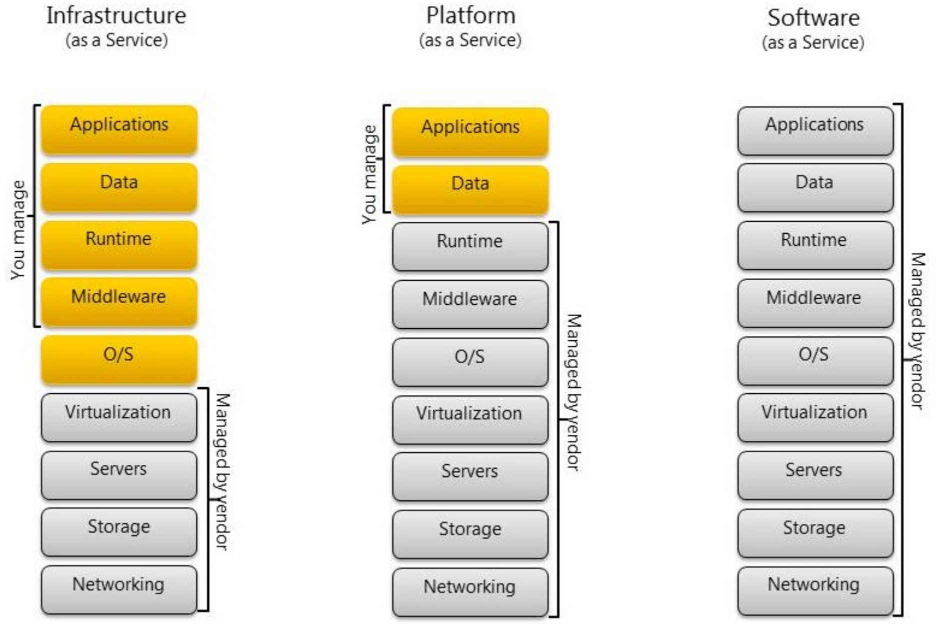

Figure 2-13 Microsoft Cloud Taxonomy [91]

I assert that this is due to the constantly changing pace of technology that extensive research in cloud taxonomies has not been undertaken. As “anything as a service” becomes more prevalent, these existing taxonomies will be considered incomplete and outdated.

## 2.9. CLOUD: NETWORKING AND DATACENTRE DEPENDENCIES

Cloud computing is currently leveraged to deliver shared service solutions, replacing on-premise applications and infrastructure, constructed to support the delivery of services within the data centre.

As workloads are migrated to the cloud, traditional network architectures are being increasingly less efficient and unnecessarily expensive. If our services are no longer in the datacentre, the requirement for local and wide area networks lessens.

Enterprise networks were designed with the data centre as the focal point. Users worked from head office and branch locations, they connected into applications and services at the datacentre over Virtual Private Networks (VPNs), leased lines and high-performance technologies such as Multi-Protocol Label Switching (MPLS) [92]. The architectural pattern adopted was commonly referred to as a "hub-and-spoke" design (**Figure 2-14**).

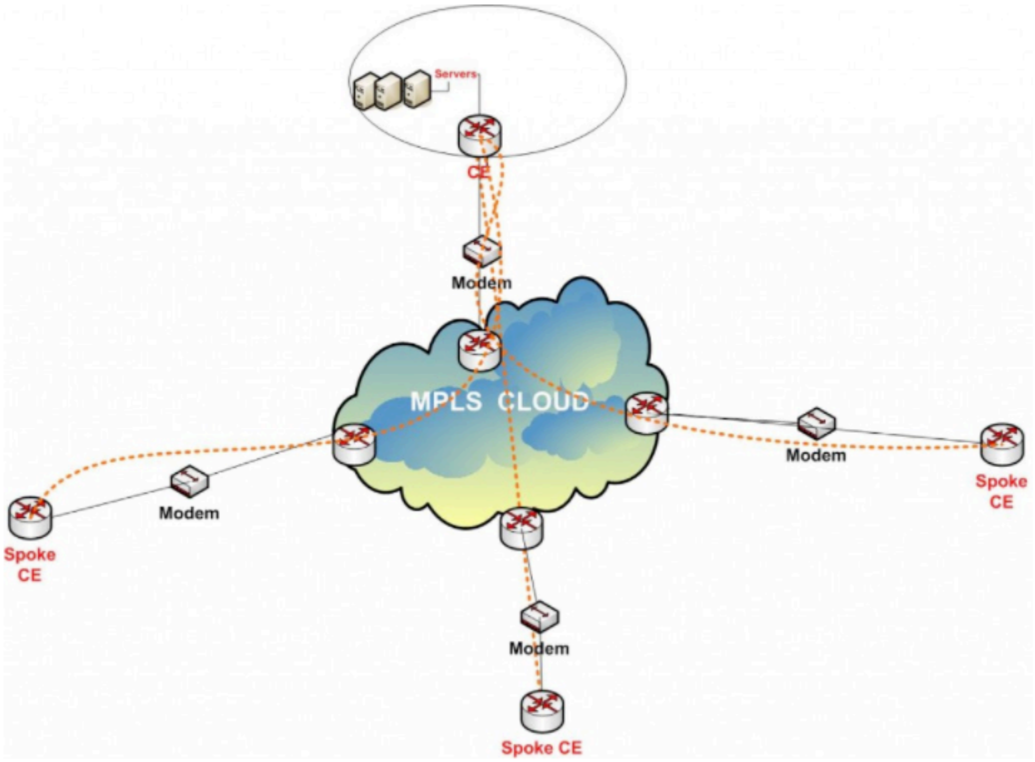

**Figure 2-14 Hub-and-Spoke Network Architecture [93]**

The hub-and-spoke model is expensive and often complex to manage although the efforts were justified as businesses required access to line-of-business applications in the datacentre.

Public cloud is changing the way organisations access resources. As applications are moving to the cloud, the requirement for inter-site communication over dedicated networking infrastructure is lessening [94]. Technologies such as Software Defined Wide Area Network (SD-WAN) are gaining popularity in the enterprise due to the surge in public cloud adoption. With SD-WAN, organisations are looking to leverage commodity Internet connections where possible, thus minimising expensive MPLS or leased circuits. SD-WAN technologies provide intelligent routing of traffic deciding which communications really need to be communicated over private lines and which can be routed "direct to Internet" (**Figure 2-15**).

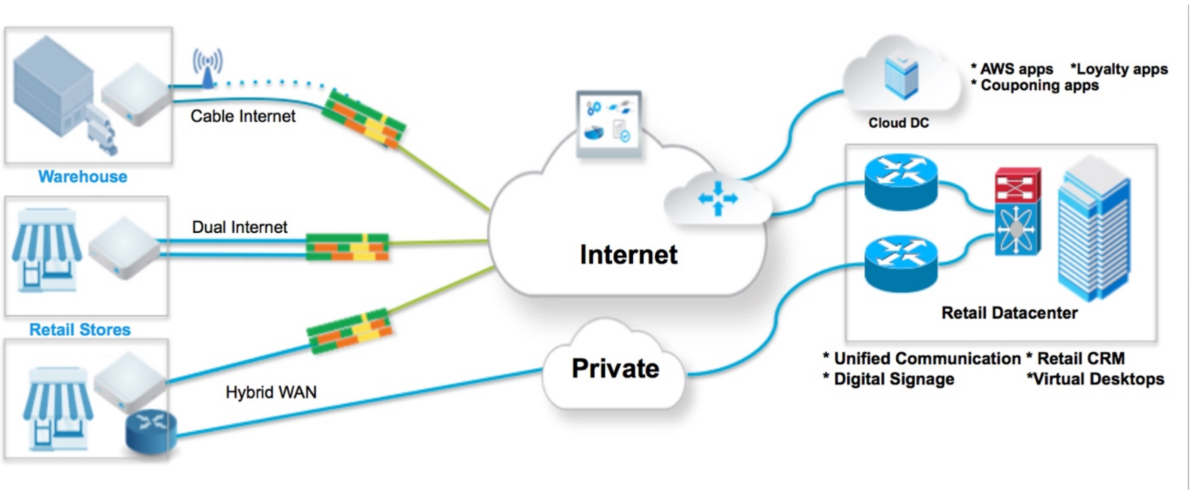

Figure 2-15 SD-WAN Overview [95]

## 2.10. CONCLUSION

The term cloud computing means many things to many different people. The definitions provided within our industry vary and are nuanced, although the attributes of service-orientated, elastic, ubiquitous (access) and scalable are resident within most characterisations of cloud.

Cloud computing is a complex ecosystem with many moving parts: actors, deployment models, reference architectures, service models and taxonomies.

Across my research, it became evident that cloud security standards have matured considerably over recent history. In the interests of project scope, a section regarding cloud standards was removed although it is worth acknowledging that many industry and academic organisations have invested time and effort in providing important artefacts to improve the reliability and security of cloud models [96]. Traditional thinking suggests that cloud adoption requires a concession in the organisation's ability to standardise and audit. A view that organisations must accept a "patchwork of cloud standards" [97]. Perhaps this position was valid five years ago, I do not believe it is defensible today. Organisations will need to amend their processes and train their people to audit cloud in new ways (via new standards) but this cannot be considered a risk/vulnerability of cloud and will not be discussed further in this paper.

History suggests that technology paradigms come-and-go. What is flavour of the month today is often displaced through a combination of technology advancement, Moore's Law[11] and vendor innovation although it is unlikely that cloud computing is a passing trend. It is predicted that the "we don't do cloud" policy rhetoric will soon be as outdated as an organisation adopting a "no Internet" policy [10].

The business benefits of cloud are well documented [51] although less focus is placed on the security benefits of cloud computing which as I have outlined in this chapter, are often as powerful a justification for cloud adoption as those associated with elasticity or cost saving. By deploying security capabilities in the cloud that are horizontally distributed across multiple locations, organisations significantly improve their ability to provide

[11] Moore's Law refers to exponential growth in processing power. Gordon Moore (1965) observed that transistors-per-inch on integrated circuits have doubled every year since their invention. The term is now colloquially used to describe a general exponential growth of computing power.

resilience and defend against denial of service attacks. The distributed nature of cloud computing also provides a vehicle for the dissemination of threat intelligence information which shortens the time from malware discovery to defence propagation.

If cloud is here to stay then, as security professionals, we need to ensure that our organisational processes and people are trained to embrace cloud services in a secure, repeatable fashion. The essential characteristics of cloud can, and are [34], debated regularly although elasticity, rapid provisioning and shared resources require new approaches to securing organisational data.

The chapter closes by repeating that cloud means many things to many people; with the explosion of the Internet of Things (IoT) and enterprise adoption of cloud services for line of business applications, public cloud adoption will continue to grow. With such diversity of use case and global coverage, my preliminary assessment is that attempting to define standards, threat models and taxonomies for cloud is significantly more onerous that established computing paradigms.

# 3. CLOUD RISKS

What is Risk? Risk is inherent in our daily lives. As human beings, we take both conscious and unconscious risks every day. A lot of research has been undertaken around the human appetite for risk and the contributing factors that make one risk palatable while others are considered "too risky". Gardner [98] draws on prominent theories for the explanation of risk decisions suggesting that, as human beings, we are driven by "the example rule" [98, pp. 18-19] – risks which we can immediately recall as local or relevant to us are treated as fundamentally "more likely". Gardner's assertion would therefore suggest a "catch 22" situation in our understanding and acceptance of public cloud: the more frequently we discuss the purported "risks of cloud", the more likely we are to deem cloud as "risky". As humans, we are predisposed to "confirmation bias"; the tendency to interpret information which supports our preconceptions.

How do we address risk in an information security context? More specifically, how do we address information risk in a cloud context? Organisations cannot make risk-based decisions without understanding the actors and relationships involved in their ecosystem: *"Understanding the relationships and interdependencies between the different cloud computing deployment and service models is critical to understanding the security risks involved in cloud computing*" [83].

This chapter will look at the risks which present themselves with cloud computing. The objective of the chapter is to define what we mean by risk and identify those risks which are present specifically because of the architecture of cloud computing. A lot has been written around the risks of cloud computing with Wisegate [99] suggesting over 50% of organisations believe the risks are just too great for certain types of information. Given our discussion regarding the sharp increase in public-cloud adoption for sensitive information (**Section 2.6.1**), there is clearly an inconclusive position. Industry research like that of Wisegate seems to contradict the quantitative adoption figures provided by cloud vendors [28, 26].

Does cloud *introduce* new forms of risk which didn't exist in previous computing ecosystems? It is important that we understand how many of these are unique to cloud and a result of the intrinsic nature of cloud architecture.

Before it is possible to explore cloud risks, it is critical that the concept of risk is understood. Too often the terms "risk", "threat" and "vulnerability" are used interchangeably thus making pragmatic risk analysis impossible and management decisions opaque.

This chapter will explore information risk management frameworks and assess their suitability for cloud computing. We will define what we mean by "information risk" and clarify what threats and vulnerabilities mean and where they fit into a risk equation. In this chapter, I will present a classification scheme for cloud vulnerabilities and threat events to delineate between issues which are a direct result of public cloud and those which are applicable to both cloud and non-cloud environments.

## 3.1. WHAT IS RISK?

It is important to understand risk from multiple perspectives; Dinu's [100], definition of risk is from an economic lens: *"Risk is defined as the probability or threat of damage, injury, liability, loss, or any other negative occurrence that is caused by external or internal vulnerabilities, and that may be avoided through pre-emptive action".* Dinu asserts that risk may be avoided through planning and application of safeguards in advance of an event.

Pritchard's [101] philosophical definition is generic enough for a layperson to understand the concept: *"…a potential unwanted event, where its riskiness is measured in terms of the probabilistic likelihood of it occurring, such that the higher the probability in question, the more risky the event".* Whilst probability of an unwanted event is an imperative consideration, I assert that impact is also a mandatory factor in any risk equation. The US

Department of Homeland Security (DoHS) defines risk as *"the potential for an unwanted outcome resulting from an incident, event, or occurrence, as determined by its likelihood and the associated consequences"* [102].

The DHS further elaborates on risk through their "Risk Lexicon" which is a comprehensive set of terms and meanings around the practice of risk management (**Figure 3-1**). Importantly, the Risk Lexicon identifies the psychological aspects of risk and whilst the DoHS Lexicon is not focused on information or cyber risk, the components are consistently used within risk management methodologies associated with our industry.

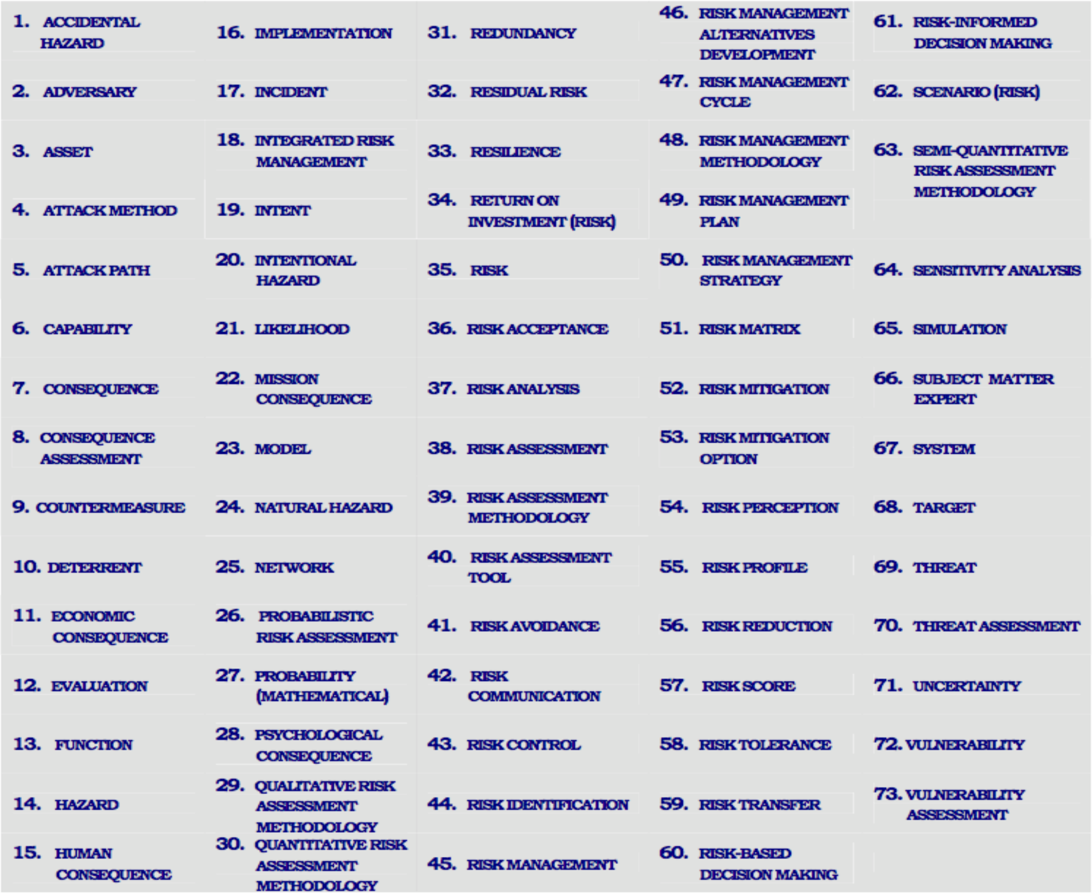

1. ACCIDENTAL HAZARD
2. ADVERSARY
3. ASSET
4. ATTACK METHOD
5. ATTACK PATH
6. CAPABILITY
7. CONSEQUENCE
8. CONSEQUENCE ASSESSMENT
9. COUNTERMEASURE
10. DETERRENT
11. ECONOMIC CONSEQUENCE
12. EVALUATION
13. FUNCTION
14. HAZARD
15. HUMAN CONSEQUENCE
16. IMPLEMENTATION
17. INCIDENT
18. INTEGRATED RISK MANAGEMENT
19. INTENT
20. INTENTIONAL HAZARD
21. LIKELIHOOD
22. MISSION CONSEQUENCE
23. MODEL
24. NATURAL HAZARD
25. NETWORK
26. PROBABILISTIC RISK ASSESSMENT
27. PROBABILITY (MATHEMATICAL)
28. PSYCHOLOGICAL CONSEQUENCE
29. QUALITATIVE RISK ASSESSMENT METHODOLOGY
30. QUANTITATIVE RISK ASSESSMENT METHODOLOGY
31. REDUNDANCY
32. RESIDUAL RISK
33. RESILIENCE
34. RETURN ON INVESTMENT (RISK)
35. RISK
36. RISK ACCEPTANCE
37. RISK ANALYSIS
38. RISK ASSESSMENT
39. RISK ASSESSMENT METHODOLOGY
40. RISK ASSESSMENT TOOL
41. RISK AVOIDANCE
42. RISK COMMUNICATION
43. RISK CONTROL
44. RISK IDENTIFICATION
45. RISK MANAGEMENT
46. RISK MANAGEMENT ALTERNATIVES DEVELOPMENT
47. RISK MANAGEMENT CYCLE
48. RISK MANAGEMENT METHODOLOGY
49. RISK MANAGEMENT PLAN
50. RISK MANAGEMENT STRATEGY
51. RISK MATRIX
52. RISK MITIGATION
53. RISK MITIGATION OPTION
54. RISK PERCEPTION
55. RISK PROFILE
56. RISK REDUCTION
57. RISK SCORE
58. RISK TOLERANCE
59. RISK TRANSFER
60. RISK-BASED DECISION MAKING
61. RISK-INFORMED DECISION MAKING
62. SCENARIO (RISK)
63. SEMI-QUANTITATIVE RISK ASSESSMENT METHODOLOGY
64. SENSITIVITY ANALYSIS
65. SIMULATION
66. SUBJECT MATTER EXPERT
67. SYSTEM
68. TARGET
69. THREAT
70. THREAT ASSESSMENT
71. UNCERTAINTY
72. VULNERABILITY
73. VULNERABILITY ASSESSMENT

Figure 3-1 DHS Risk Lexicon [103]

Risk nomenclature varies across industry and academic bodies, I was surprised to see that all the DoHS Risk Lexicon was appropriate for a conversation regarding information risk. Of the 73 attributes, only the elements pertaining to a "hazard" could be contentious in a cyber context; hazards are commonly associated with circumstances *"…perceived to be capable of causing harm or costs to human society"* [104]. As explained in **Section 2.3.1,** cloud systems are increasingly being leveraged to deliver transportation and healthcare services, I assert that altering the confidentiality, integrity or availability of information within these systems could cause harm to human society.

**Figure 3-2** diagrams a commonly accepted risk equation at its most foundational level. For a risk to occur, we need to apply qualitative or quantitative analysis measures to ascertain how likely the risk is to occur and what is the impact to the organisation if this happens. Impact is often associated with financial loss although impacts can be reputational, operational, legal and/or health and safety related. In an information security content, impact is generally assessed through a formalised process of data classification and business impact assessment (BIA); failure to classify information and understand its importance to the enterprise results in an inability to apply commensurate people, process and technology controls to information. This is a fundamental vulnerability and given the need for stakeholder engagement outside of IT, one which is (in my extensive experience) overlooked in organisations.

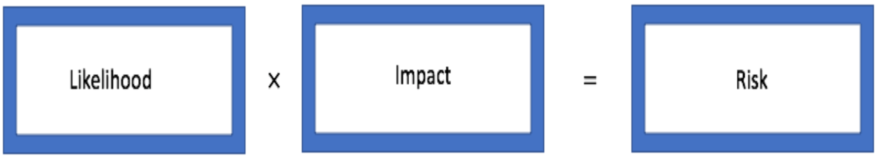

Figure 3-2 Generic Risk Equation

Other definitions exist which see risk as potentially positive. The International Organisation for Standardisation (ISO) takes a simplistic definition being *"the effect of uncertainty on objectives"* [105]. This definition is interesting as it implies that risk isn't always a negative. The Information Security Forum subscribes to the ISO definition as part of their Information Risk Assessment Methodology (IRAM2) (covered in **Section 3.5.1**) and IRAM2 explicitly states *"this definition of risk implies that risk is not necessarily negative; in fact, it recognises some risks could result in organisations exceeding their objectives"* [106]. IRAM2 [106, p. 3] goes on to define "information risk" specifically and this does always have a negative consequence: *"Information risk is the risk of loss to an organisation resulting from the compromise of certain attributes of its information assets, namely confidentiality, integrity or availability*". Having reviewed available literature on the subject, I assert the following:

1. In a general context, risk can have both positive and negative connotations.
2. When discussion information risk, the exposure relates to the compromise of confidentiality, integrity and/or availability of information as is always negative.

The ISF's definition of information risk most appropriately defines the challenges we will be exploring in this paper**.** IRAM2's updated definition used to define information risk brings with it a more granular risk equation (**Figure 3-3**).

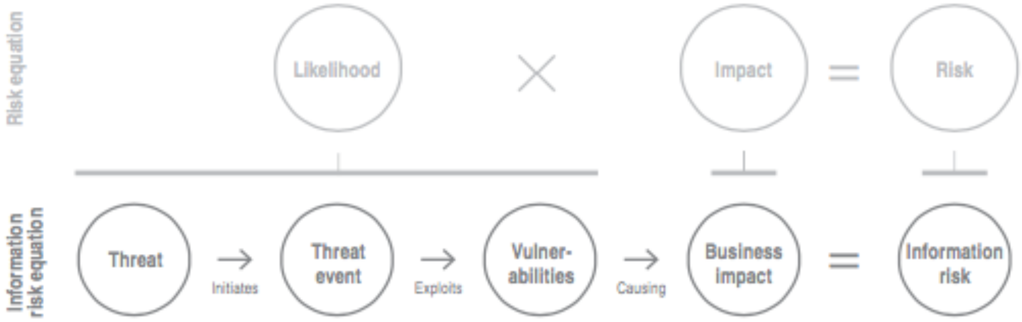

Figure 3-3 IRAM2 Information Risk Equation [106]

Multiple industry/academic organisations share the ISF's view of information risk. The NIST Guide for Conducting Risk Assessments [107] outlines a generic risk model which adopts an almost identical set of steps in assessing risk (**Figure 3-4**):

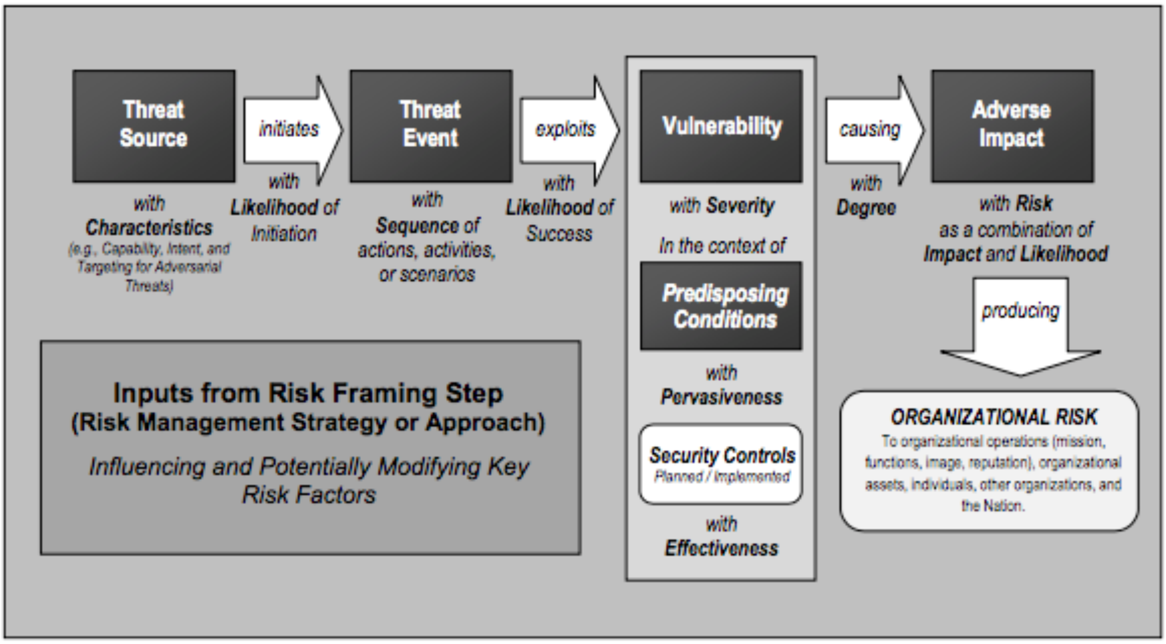
Threat Source
initiates
with Likelihood of Initiation
with Characteristics (e.g., Capability, Intent, and Targeting for Adversarial Threats)
Threat Event
with Sequence of actions, activities, or scenarios
exploits
with Likelihood of Success
Vulnerability
with Severity
In the context of
Predisposing Conditions
with Pervasiveness
Security Controls Planned / Implemented
with Effectiveness
causing
with Degree
Adverse Impact
with Risk as a combination of Impact and Likelihood
producing
ORGANIZATIONAL RISK
To organizational operations (mission, functions, image, reputation), organizational assets, individuals, other organizations, and the Nation.
Inputs from Risk Framing Step (Risk Management Strategy or Approach)
Influencing and Potentially Modifying Key Risk Factors

Figure 3-4 NIST Generic Risk Model [107]

For the purposes of my analysis, either the IRAM or NIST models could be used. I have decided to continue with the ISF IRAM2 model as it provides the necessary granularity to assess control strength and includes threat event and actor capabilities which are required for our assessments of resource isolation in **Chapter 4**.

For completeness, I have included reference to the Information Systems Audit and Control Association (ISACA) Cloud Risk Assessment Framework (CRAF) [108]. The CRAF was the only "cloud-specific" risk framework I discovered across my research. The CRAF looks to amalgamate information from several distinguished industry resources [109, 110, 35] although as a framework which can be practically applied for the assessment of threats and vulnerabilities pertaining to cloud, I found it to be too theoretical. This artefact should be reviewed by anyone looking at ways to measure and classify cloud risk (defining impact) although in my opinion assessing likelihood via the CRAF is problematic and prone to error.

## 3.2. RISK ASSESSMENT

*"... For operational plans development, the combination of threats, vulnerabilities, and impacts must be evaluated in order to identify important trends and decide where effort should be applied to eliminate or reduce threat capabilities; eliminate or reduce vulnerabilities; and assess, coordinate, and deconflict all cyberspace operations..."* [107].

**The National Strategy for Cyberspace Operations**
**Office of the Chairman, Joint Chiefs of Staff, United States, Department of Defence**

Risk cannot always be avoided, nor necessarily is risk-avoidance the desired result. Avoiding risk entirely will often render a business without the opportunity to leverage the benefits that an opportunity may bring. Gartner [111] notes that this is particularly pertinent for cloud: *"Failing to analyse cloud risk will result in missed opportunity and/or unacceptable risk to the business. Risk management is a mature discipline that can determine how much cloud risk is acceptable"*. ENISA [35] takes a similar view stating *"Risk should always be understood in relation to overall business opportunity and appetite for risk – sometimes risk is compensated by opportunity"*.

It is imperative that organisations take balanced risk decisions and that risk is managed; this does not necessary mean "avoided". What organisations need is pragmatic risk management. Fundamentally, risk can be handled in four ways which are covered in **Table 3-1**:

| **Risk State** | **Explanation** |
|---|---|
| Avoided | Risk is identified and measures are taken to avoid the risk entirely. If organisations attempt to avoid all risks, business productivity is likely to be impacted. |
| Mitigated | Measures are taken to lessen the impact or likelihood of a risk. Controls can be applied which lessen the impact or likelihood of an event. |
| Transferred | The risk is transferred to a third party. This may be another person or department in the enterprise or through the purchase of insurance[12] |
| Accepted | The business takes the decision to accept the risk. |

**Table 3-1 Risk Status [112, p. 42]**

Risk assessment is the process by which an organisation decides upon the risk action to take in a specific situation. Assessing risk provides context and relevance to an organisation. Organisation A may decide that a risk needs to be avoided at any cost, Organisation B might decide the risk can be accepted or perhaps mitigated through the inclusion of controls. Risk assessments fall into two key categories:

### 3.2.1. QUANTITATIVE RISK ASSESSMENT

Quantitative risk assessment comes into play when we can map a dollar amount to a specific risk [113].

Through quantitative risk assessment, a numerical quantity or amount can be expressed and assigned to a risk. In quantitative analysis, statistics are gathered and used to determine the probability of a threat event occurring.

Quantitative risk assessment is well-suited to environments and industries with rich statistical data over extended periods. The home and life insurance industries operate via quantitative means to assess likelihood of burglary or life expectancy of their customers. Trend data exists for these events and reaches back for centuries; the same cannot be said for cyber security.

Quantitative analysis is practical when we are looking to obtain a monetary value figure. It is best-suited to qualify impact (over likelihood). Organisations should follow a process of data classification and Business Impact Assessment (BIA); during these activities, a financial value is often placed on data records. A more coarse-grained approach is to understand penalties from regulators for a data breach or non-compliance with a

[12] Cyber Insurance is growing in popularity. Organisations are taking out policies to cover losses in the event of a breach of information. Whilst appealing to customers, the ability to apply quantitative likelihood equations is onerous for insurers and impacts are impossible to estimate; this is resulting in punitive premiums and a slow uptake by organisations [237]

standard. If an organisation's file server has 1000 records of customer data and the cumulative cost per record of a breach is $158 [114] then value can be assigned as a worst case financial impact. Quantitative assessment leaves a smaller margin for error although it is not always practical.

A good example where quantitative risk assessment can be used in a cyber context would be to ascertain the likelihood of a cyber-attack emanating from a country; this is measurable information which can be extracted from a Security Information & Event Management (SIEM) system. Unfortunately, organisations can only measure what they are managing. The collection of all sensor information, for all technology systems is expensive and often cumbersome. History of an event occurring should not be a unilateral factor in a likelihood calculation which is why qualitative analysis is almost always used to support quantitative figures.

### 3.2.2. QUALITATIVE RISK ASSESSMENT

Most information and cyber risk assessments rely on information from the qualitative category. A qualitative risks assessment requires a subjective evaluation and should be performed by a subject matter expert (SME).

It is often impossible to analyse quantitative data to assess the likelihood of an organisation experiencing a compromise of their data. In an ideal world, the security professional would provide a percentage likelihood of a data breach although cyber security presents a volume of variables which cannot be suitably measured as absolutes. For example:

- Attribution of threat actor location is onerous due to anonymising services and botnets.
- Security control strength and application consistency varies across organisations.
- Forensic breach analysis information is not publicly provided.

Qualitative risk analysis relies on expertise and previous experience to ascertain the likelihood and impact of data breach. It is important to acknowledge that qualitative analysis should still follow a repeatable framework for risk management.

I will continue with our data breach example in **Section 3.2.1** to exhibit the need for qualitative analysis:

If we assume a level of volatility in the number of data records on a file server, the ability to apply a financial figure is impossible. It is also impractical to apply a "worst-case" rating to all risks. Threats to an organisation must be prioritised and prioritisation is achieved through understanding how likely a risk is to occur. The application of controls to mitigate the impact of risk are almost always considered; this process also increases the subjectivity of our analysis and moves the onus to qualitative analysis. Industry models such as NIST 800-53 [13] and ISF IRAM2 rely on qualitative analysis. **Figure 3-5** details ENISA's qualitative risk analysis model which takes a conventional risk metrics, impact and likelihood, and provides qualitative ratings for each category. Another common scoring system is to use "low, medium and high".

[13] 800-53 uses the definition of "semi-quantitative" for risk calculation although expertise and industry experience are still required to produce meaningful data. This is therefore a qualitative approach.

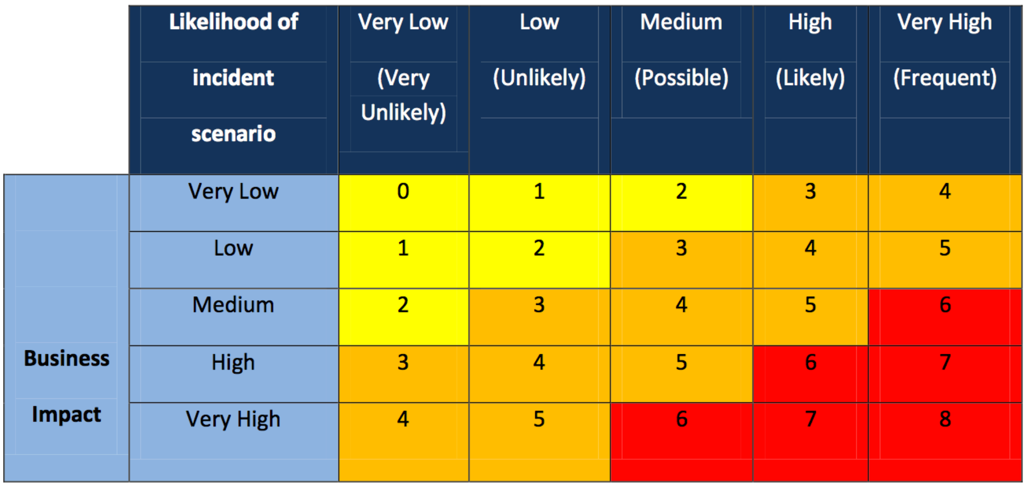

|  | Likelihood of incident scenario | Very Low (Very Unlikely) | Low (Unlikely) | Medium (Possible) | High (Likely) | Very High (Frequent) |
|---|---|---|---|---|---|---|
| Business Impact | Very Low | 0 | 1 | 2 | 3 | 4 |
|  | Low | 1 | 2 | 3 | 4 | 5 |
|  | Medium | 2 | 3 | 4 | 5 | 6 |
|  | High | 3 | 4 | 5 | 6 | 7 |
|  | Very High | 4 | 5 | 6 | 7 | 8 |

Figure 3-5 ENISA Risk Ratings [35]

The time has been taken to call out risk assessment categories as in later chapters, an exploration of cloud-specific threats and vulnerabilities will take place. It is the author's hypothesis that our industry focuses too heavily on inherent impact (explained in **Section 3.7** ) and avoids the qualitative assessments needed to make balanced risk decisions. If mature organisations in all industry verticals are migrating to platforms such as Office 365 and Salesforce.com, we cannot say that the likelihood of a data breach is any higher than an on-premise solution. These organisations invariably have established risk management and information security functions which systematically assess all technology projects.

Hypervisors and virtualisation (core to public cloud, see **Chapter 4**) are relatively new technologies and introduce vulnerabilities which were previously not applicable in a world of physical appliances. It is important however that we understand the likelihood of a vulnerability being exploited. I will continue to support my assertion through the application of the IRAM2 model in **Sections 4.8, 4.9**. I will apply a qualitative risk analysis across the actors, threats and vulnerabilities which apply to public cloud.

The security benefits outlined in **Section 2.3.2** cover compelling (security) improvements through public cloud adoption and they certainly go a long way to mitigate the inherent impact of a vulnerability being exploited.

## 3.3. HOW DOES THIS APPLY TO CLOUD?

Qualitative versus quantitative is only one of the many considerations for risk assessment in information security. Different industry bodies categorise the types of risk association with information risk in several ways. Another consideration is the scarcity of "cloud-only" risk assessment methodologies. As I will subsequently outline, perhaps this is related to a lack of genuine "cloud-only" risks.

Gartner [111] defines a series of cloud-specific risks (**Table 3-2**). Gartner's categorisation is thorough and decomposes cloud risk into five areas: agility, compliance, supplier, availability and security. It is my opinion that the Gartner classifications introduce confusion and complexity. Risk, in the context of this paper, is the loss to an organisation resulting from the compromise of the confidentiality, integrity and/or availability of data.

Gartner's separation of "security" into a specific category suggests that information security should be considered in isolation, it is my opinion that failure to manage risk in the four other categories could impact the confidentiality, integrity and/or availability of information.

| **Risk Category** | **Description** |
|---|---|
| Agility | The ability of the cloud service to provide the features, functions, forms and levels of service that are necessary to meet business requirements, both immediately and over time. |
| Compliance | Legal requirements may be mandated by regulators or through a contract. This may take the form of very specific requirements or restrictions on how specific forms of data may be used in the public cloud, and it may also take the more ambiguous form of a standing requirement to be able to demonstrate an appropriate level of attention to risk assessment and control. |
| Supplier | The public cloud market is highly competitive and capital-intensive, which makes many cloud service providers (CSPs) financially fragile. CSP business failures have occurred on short notice |
| Availability | Availability of service and data such as unrecoverable data loss or service disruptions. Contingency planning is not always practical for every form of cloud service. |
| Security | Broad category that legitimately means different things to different organizations, but generally this addresses the limiting of access to data, and the protection of data from hostile acts that would impact confidentiality, integrity or availability. |

Table 3-2 Gartner Cloud Risk Categories [111]

Gartner's view serves to evidence the subjectivity of the risks associated with cloud along with the importance of context and applied risk management. I believe that equally compelling argument could be proposed which argues a diametrically opposite position for each of their five categories (**Table 3-3**). This highlights the subjectivity of opinion regarding public cloud.

| **Benefit Category** | **Description** |
|---|---|
| Agility | Cloud computing reduces time-to-market and allows for service provision in multiple datacentres. Capabilities such as Amazon Elastic IP improve failover and migration scenarios whilst VM cloning allows for almost instantaneous standing up of instances in multiple geographies. |

Table 3-3 Security Benefits Categorisation

| | |
|---|---|
| Compliance | Amazon Web Services **[20]** provide cloud environments which are pre-validated against regulatory compliance frameworks. Deployment of solutions into specific AWS regions ensures that data can be retained within a specific geography **[115]**. |
| Supplier | Broad selection of SaaS, PaaS and IaaS suppliers drives commercial competitiveness and technology innovation. |
| Availability | Services available in multiple geographies, with in-built failover and redundancy. AWS provides Availability Zones and Regions for this purpose. "Regions" are leveraged to manage latency and availability zones provide redundancy in region through connections to multiple ISPs and power grids [115] |
| Security | Refer to **Section 2.3.2**. A compelling case is presented that suggests public cloud improves an organisation's security posture. |

**Table 3-3 (Cont.) Security Benefits Categorisation**

ENISA is a centre of expertise for cyber security in Europe [116]. ENISA have dedicated time to produce many advisories covering cloud computing and risk management. ENISA documents a top ten set of risks for cloud [35, p. 9]. They further break cloud risks into a model which closely aligns to people, process and technology; a structure common to information security when discussing risks and controls. Their more detailed analysis covers 23 "Risks" which are covered in **Table 3-4**.

| **Risk Category** | **Description** |
|---|---|
| Policy and Organisational | Vendor Lock-In \| Governance\| Compliance \| Reputation \| Service Termination \| Acquisition \| Supply Chain Failure |
| Technical | Resource Exhaustion \| Isolation Failure \| Malicious Insider \| Interface Compromise \| Data Interception \| Data Leakage \| Insecure Data Deletion \| Denial of Service (DDoS) – Distributed/Economic \| Loss of Encryption Keys \| Malicious Probes \| Compromised Service Engine \| Hardening Conflicts |
| Legal | Subpoena and E-Discovery \| Changes of Jurisdiction \| Data Protection Risks \| Licensing Risks |

**Table 3-4 ENISA Risk Categories [35]**

Cloud computing can multiply the impact or likelihood of an existing risk but as the ISF defines, information risk is concerned with the compromise of the confidentiality, integrity and/or availability of data. Cloud arguably introduces new threat actors, threat events and vulnerabilities. If an enterprise can understand these three factors, they can take a balanced approach to risk management. As NIST, ENISA and ISF all define, for a risk to occur, we need a threat and a vulnerability. A threat without a vulnerability isn't a risk. Equally, a vulnerability without a threat actor isn't a risk.

Wilhelm [117] provides a succinct definition regarding the differences between a risk and a threat: *"In its simplest terms, a threat is something that can do damage to a system (such as malware). The risk describes the likelihood and impact of the threat…".*

Taking our agreed definition of risk, we can define that for all risks of cloud, the risks are to the compromise of confidentiality, integrity and/or availability of data. The impact can be financial, reputational, legal or regulatory, operational or health and safety. I will continue to discuss how impact is altered in a public cloud model.

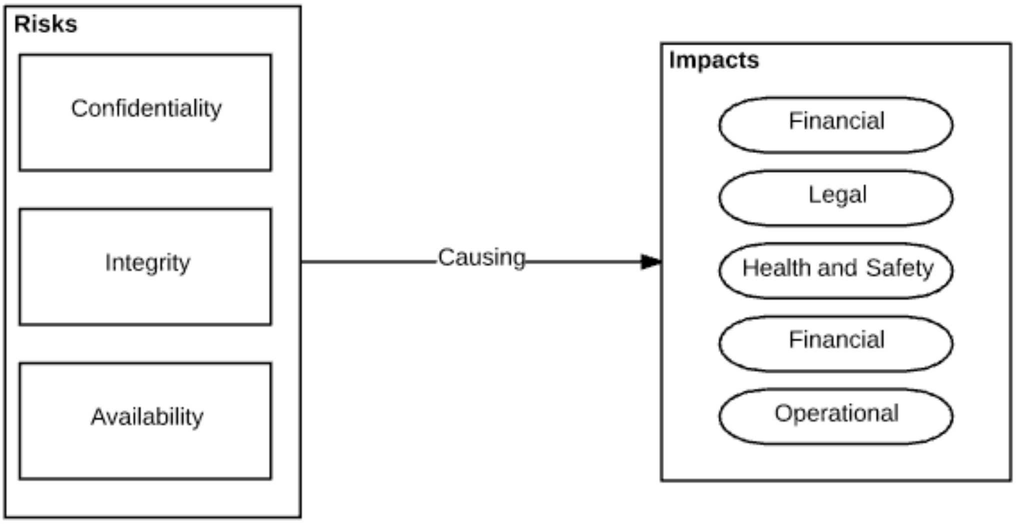

Figure 3-6 Cloud Information Risk Categories

The risks in **Figure 3-6** would be risks for an organisation whether applications and infrastructure were deployed on-premise or in the public cloud. It is important to differentiate between what a risk is and the impact and/or likelihood of a risk occurring. The recent (February 28th, 2017) outage with Amazon's Simple Storage Service (S3) infrastructure serves well to illustrate my point:

Through routine debugging activity, an incorrectly entered command caused catastrophic availability implications for Amazon S3 in the "US-EAST-1" Region on 28th February 2017 [118]. This unintentional IT mistake (debugging of a billing system) meant that customers who relied on S3 resources in that region were taken offline for several hours. Amazon's thorough incident report explaining that errors in people and process are simply unavoidable [119] and it is important that service restoration occurs as quickly as possible.

If we objectively consider the components of this risk equation, we see that there is nothing unique about the threat actor, the threat event or the vulnerability associated with the S3 outage. An accidental (**Figure 3-9**) threat event (misconfiguration), exploited a vulnerability (user error) to cause business impact although the repercussions were widespread causing outages across the Internet [118]. Will such a high-profile outage at Amazon cause organisations to validate existing confirmation bias in relation to public cloud security? Gigerenzer [120] (**Section 3.3.1)** discusses the human predisposition to disproportionately weight high-impact, low-likelihood attacks over those which more frequently occur. Cloud outages at mature CSPs happen very rarely and services are recovered quickly.

We can draw on another technological parallel when discussing risks, threats and vulnerabilities which are applicable for cloud: Online fraud/cybercrime is now the most common crime in the United Kingdom with almost one-in-ten people falling victim [121, 122]. "Online/electronic crime" no longer requires delineation from "crime". Criminals are using technology to carry crimes which were previously performed in person;

technology being the mechanism for delivery. Card skimming used to be an activity carrying a high likelihood of detection as "skimmers" needed to be added to ATMs, our online world has created skimming 2.0 where an attacker can achieve the same outcome without the need for physical hardware [123].

In **Section 4.3**, I will discuss the attributes of modern computing although it is appropriate here to highlight that many of the vulnerabilities and threats outlined for cloud, were documented at a time when virtualisation and "cloud-like" technologies had not reached the prevalence they have in 2017. This embryonic phase meant that organisations who were early-adopters carried the risks of trial-and-error associated with any nascent solution. As we will return to in **Chapter 5**, public cloud platforms are now capable of offering comparable security technologies to those housed in a customer's data centre.

NIST [107, p. 10] introduce a concept that I assert is directly applicable to public cloud computing – a predisposing condition: "*A predisposing condition is a condition that exists within an organization, a mission or business process, enterprise architecture, information system, or environment of operation, which affects (i.e., increases or decreases) the likelihood that threat events, once initiated, result in adverse impacts to organizational operations and assets, individuals, other organizations, or the Nation*". I discovered the term "predisposing condition" after I defined "exacerbated by cloud" although NIST's definition validates my opinion.

Multitenancy is a trade-off; organisations benefit from the elasticity, cost and performance benefits of a shared service although the impact of a breach of CIA in one tenant could impact others (tenants). This paper will attempt to assess the likelihood of this occurring.

If cloud does not bring with it unique risks, why is there so much academic and industry information citing the "risks of cloud"? I assert that this could be a result of a fear of the unknown. Cloud is a new paradigm and brings with it new ways of working. The cost and efficiency savings associated with cloud naturally bring with them an opportunity for organisations to downsize staff numbers; cloud facilitates vendor consolidation, the removal of on-premise appliances and a reduction in Wide Area Network (WAN) traffic. These factors ultimately improve efficiencies and allow organisations to reduce staffing levels. In the same way turkeys do not vote for Christmas, those who maintain appliance-based, on-premise architectures are unlikely to espouse the benefits of low-maintenance, public cloud platform.

We will take the consistent components of the NIST and ISF models for our research and critically review the threats and vulnerabilities associated with cloud computing. Cloud certainly introduces inherent and specific vulnerabilities although there are far fewer which are directly a result of public cloud adoption than appears with a cursory glance. The impact of these "new" vulnerabilities are however significant which further supports the need to better understand likelihood. It is also my opinion, which I will support in **Chapter 4**, that the likelihood of vulnerability exploitation, by a selection of prevalent threat actors, is significantly less probable for public cloud vulnerabilities than more traditional attack vectors that we see in existence today.

To frame the conversation in **Chapter 4** and better qualify my above assertion, **Figure 3-7** depicts the components involved in understanding a cloud risk assessment framework. These components have been gathered from both 800-30 [107] and IRAM2 [106] The diagram identifies areas of "Cloud Focus" which I feel identifies the components of the risk equation which materially change because of (public) cloud: Threat actors remain broadly as they were before public cloud adoption. Actors and their associated events are either adversarial or accidental in nature. Public cloud has not introduced a new category of threat events; in an adversarial context, an actor needs to apply similar offensive methods to those she would have used on-premise. A similar situation exists for accidental events; fundamentally, user error is user error and the misconfiguration of a server could occur irrespective of the location of the instance. The categories of risk remain as they were pre-public cloud. So, what changes to the risk model has public cloud introduced? Public cloud exacerbates the impact of several "common" vulnerabilities which are exploited by threat events.

For my study, I will break cloud threats into three areas (**Table 3-5)**:

| **Category** | **Description** |
|---|---|
| Unique to Cloud | Threats which do not exist in previous computing models – I aim to prove that these are scarce but warrant explicit attention. |
| Exacerbated by Cloud | Threats made inherently and inextricably more severe through cloud: Limited and acknowledged in this work. |
| General | Threats which may reside within the environment but cloud residency is not a consideration. I will aim to prove that this covers most vulnerabilities prevalent in computing today. |

Table 3-5 Threat Classification

In the interests of retaining the focus of my argument, I have documented an **Appendix A: ENISA Cloud Vulnerability Assessment** where I will address each "cloud vulnerability" raised (by ENISA) and analyse if it is exacerbated by public cloud.

While I believe that public cloud does not introduce a wave of new vulnerabilities, the impact of resource sharing between different organisations means that there is a "by proxy" or secondary impact to consider which was never applicable in an on-premise deployment. It is therefore critically important that we can evidence the following factors:

- The likelihood of exploiting the architecture of resource sharing is small enough to have the associated risks "accepted" by the organisation.

- The application of security controls suitably "mitigate" the risks to a level palatable for the organisation.

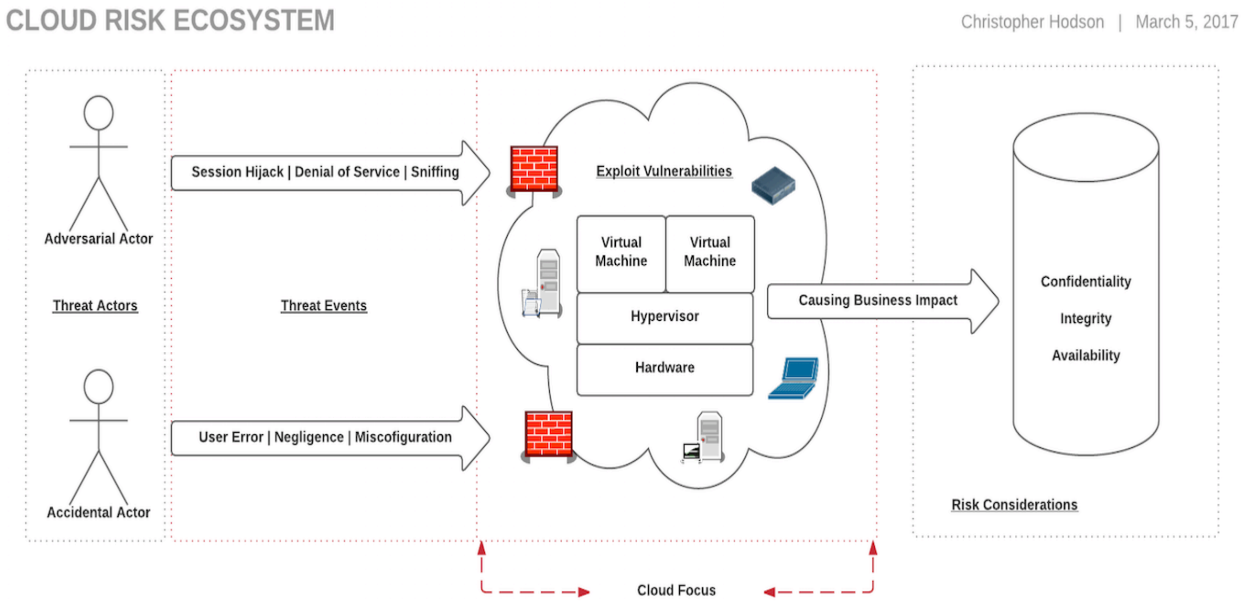

Figure 3-7 Cloud Risk Ecosystem

### 3.3.1. FEAR OF CLOUD

Quantitative and qualitative risk assessment methodologies provide organisations with set of holistic methods to assess and quantify risk in their own environments. What cannot be overlooked in any assessment of risk are emotion, fear and self-preservation. Essentially, the innate characteristics we take from our ancestors.

Across the months that followed the atrocities of September 11th 2011 [124], there was a significant drop in the number of passengers travelling on aeroplanes. The very visceral impact of the events that day caused people to feel safer in other forms of transport.

Gigerenzer asserts that people tend to fear what he calls "dread risks": low-probability, high-consequence but with a primitive, overt impact. In 2004, his study entitled "*Dread Risk, September 11, and Fatal Traffic Accidents*" [120] pulled figures from the months and years before 9/11 and those immediately afterwards. Gigerenzer proves beyond reasonable doubt that the events of 9/11 caused more people to travel across the United States (US) via automobile rather than take a flight. This increase in road travel resulted in an increased number of cars on the road and consequentially road traffic accidents which resulted in fatalities. Whilst Gigerenzer proved through statistics that flying was considerably safer than getting in a car, people feel safer in cars. This myopic approach to risk management has parallels in the world of cloud; I acknowledge that the inherent impact of a cyber-attack is significantly increased if multiple tenants share the same infrastructure although the likelihood of such an attack via means exclusive to public cloud is drastically lower than, say, credential theft or a watering hole attack [125]. Based on my research and risk processes followed in **Chapter 4**, it appears that an aversion to public cloud is not based on measured risk or pragmatic assessment of likelihood.

## 3.4. CLOUD THREAT ACTORS

It is important that we consider the different threat actors involved in a public cloud environment. Shostack [126] explains that cloud computing does introduce new threats actors:

| **Cloud Threat Actor** | **Description** |
|---|---|
| CSP Insider | Staff at the third-party - The employees of the cloud service provider |
| Cloud Tenant Users | A user or administrator capable of exploiting a vulnerability in the cloud ecosystem to compromise another tenant. This user could be an administrator of the cloud platform or another customer of the service. |

**Table 3-6 Cloud Threat Actors [126]**

Shostack identifies new actors although malicious and accidental insiders and administrators have existed as established threat actors for as long as information risk management has been a discipline. The CSP insider is simply an insider: someone with a greater level of privilege and physical access than a traditional external, malicious individual or group. Whilst not downplaying the significance of insider threat, the insider being CSP based does little to change their motivations or capability. I would agree that the impact to accidental or nefarious activities could be exacerbated although it is important to remember that security controls should be deployed commensurate with the sensitivity of the information being stored or transmitted in the public-cloud environment.

Shostack suggests public cloud introduces new actors because "when you move your data or operations to someone else's cloud, you add a trust boundary" [126]. This is true from a legacy threat modelling perspective but the class of actor is "malicious insider" irrespective of location. Shostack's assertion is also predicated on an ability to apply network-based security. As enterprise adoption of public cloud continues to grow [26, 25] and sophisticated cyber-attacks continue to achieve success, the traditional "network boundary" or perimeter becomes opaque and of limited efficacy. Forrester going as far as to say that defining trusted interfaces is now impossible [127]. If trust is impossible to achieve and industry analysts are recommending a "Zero Trust" model [127] then I assert that we are better placed applying a consistent set of security controls based on information sensitivity irrespective of location. This is increasingly pertinent as traffic patterns are progressively moving away from the data centre and towards cloud services [25].

The security controls required for an accidental or malicious actor should not change based on their location or residency. In the same way that a remote working connecting via a Virtual Private Network (VPN) should equally be afforded the same security controls and consistent operating experience. I wanted to understand if security controls were more robust for an on-premise accidental threat actor than that of a CSP insider? More specifically, the process controls associated with vetting of staff. My rationale being that the more thorough the background checks, the better our ability to assess likelihood of initiation (LoI) (please refer to **Section 4.7.1** for a detailed explanation of LoI). A natural conclusion to draw would be that controlling the vetting process would provide better visibility and control over staff; however, again using Amazon as our test case, the scale and completeness of their vetting service is exemplarily and so the AWS operational implementation of least privilege. AWS perform full vetting and background checks on all staff in-line with their multitude of regulatory compliance attestation. Staff have no logical access to customer instances and all control plane (management of cloud) access is strictly limited and monitored through bastion hosts, principal of least privilege and zoned data centres). [128]

Having researched the subject, I share a similar view regarding Shostack's second category: Cloud tenant users/admins. This category of actor would not be present in an entirely on-premise implementation. Tenant users could be accidental or malicious in nature. Cloud tenant abuse only becomes an issue if avoidable vulnerabilities are present in the cloud ecosystem. I will present these across **Chapter 4** and will include the appropriate mitigations and compensating controls. With regards to customer instance management (example: IaaS resource allocation) management interfaces for cloud services should appropriately logically secured. A tenant from one organisation should be without a route to impact the posture of another tenant.

The ISF IRAM2 breaks threats, and consequential actors, into three categories: Malicious, Accidental and Environmental (**Table 3-7**):

| **Threat Actor Category** | **Description** |
|---|---|
| Adversarial | Malicious actor. Deliberate and nefarious actions against the enterprise. |
| Accidental | Related to unintentional action by the threat actor. |
| Environmental | Threats which are initiated through factors and elements outside of the enterprise. Man-made and natural hazards fall into this category. |

Table 3-7 Threat Actor Categories

NIST supports the above categories and adds "structural" which is defined as *"failures of equipment, environmental controls, or software due to aging, resource depletion, or other circumstances which exceed expected operating parameters"* [107]. Environmental and structural threat categories are categories of threats which are essentially a combination of actor and threat event. Environmental threat events are technically speaking of an accidental categorisation although they are not initiated through a physical actor. Structural and Environmental are two areas which are certainly improved by cloud computing. The author was unable to find information to the contrary. Keeping with our AWS scope, the structural and environmental benefits of public cloud are clearly visible. As I covered in **Table 3-3**, Availability Zones and Regions provide an organisation with tenant residency in multiple public datacentres without the cost and operational overhead of managing private locations. The redundant architecture of AWS makes it inherently more capable of sustaining an outage than an organisation relying on a private datacentre architecture.

The pre-existing categorisations of adversarial, accidental, environmental and structural fit the cloud model and require no amendments. Cloud does introduce new actors into existing categories, namely Cloud Insider and Co-Host User although these are instances of existing categories. If a service was hosted by a third party in their data centre, the categorisation would remain identical. It is not public cloud that introduces the category of actor although I acknowledge that exploitation of a vulnerability initiated by one of these actors would be exacerbated by public cloud due to the multitenant architecture. Multitenancy will be covered in detail in **Chapter 4**.

## 3.5. CLOUD THREAT EVENTS

For the purposes of risk analysis and control selection, it is important that cloud threats events are decomposed into a series of appropriate categories. A threat event is the action or lack thereof of a threat actor.

The ISF and NIST both explicitly define a series of adversarial and accidental threat event categories and it is my opinion that these are appropriate in understanding cloud risk.

### 3.5.1. ISF

The three ISF IRAM2 risk categories are decomposed in a series of “threat event types” which align to threat actor categories outlined in “3.3.1: Adversarial, Accidental and Environmental”.

Adversarial threat events are detailed in **Figure 3-8.** All threat event categories are considered applicable for a public-cloud eco-system.

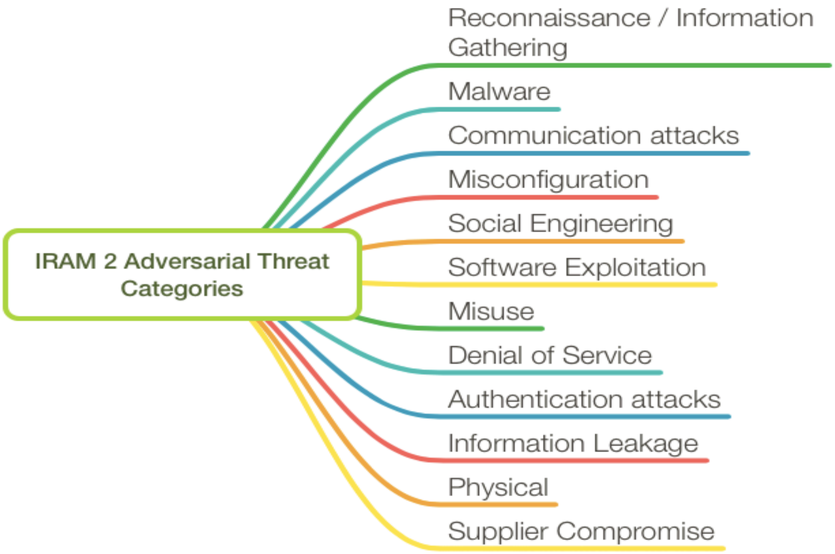

Figure 3-8 IRAM2 Adversarial Threat Events

None of the threat events outlined in **Figure 3-8** are exclusively cloud considerations. It could be argued that reconnaissance and information gathering are made easier as public cloud services are required to be externally resolvable; we will however identify controls in **Section 4.8.1** which mitigate the threats associated with reconnaissance. In an AWS context, it is ultimately the explicit action of the customer to enable inbound access to an instance. It is also important to point out that any application instance with a requirement to service public connection requests requires Internet Protocol (IP) and Domain Name System (DNS) information to published.

There is a permutation of the "communication attacks" / "authentication attacks" threat which is exacerbated by cloud and identified as a specific 'Account or Service Traffic Hijacking' cloud threat by the CSA as part of the "Notorious Nine: Cloud Computing Top Threats in 2013" [129]. Essentially, through compromising account credentials used to access AWS, an attacker would subsequently have potential access to further machines under an administrator's control. This threat isn't to be overlooked although mature and freely available controls exist (multifactor authentication) which mitigate the threat. Converged infrastructure models (**Section 4.3**) mean that the threat exists irrespective of a public or private (cloud) deployment model.

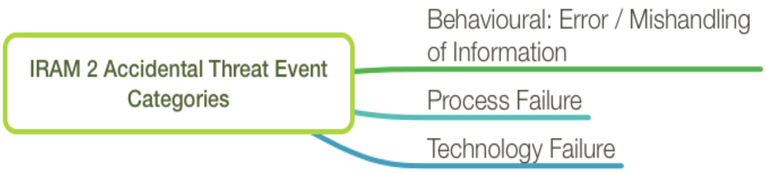
IRAM 2 Accidental Threat Event Categories
Behavioural: Error / Mishandling of Information
Process Failure
Technology Failure

**Figure 3-9 IRAM2 Accidental Threat Event Categories**

Accidental threats can occur in all deployment models. The threat events called out by IRAM2 are certainly not to be ignored but they should be treated as they would be in any technology deployment. A good architect and designer is thinking about technology and process failure with all solutions. Mishandling of information is a threat that needs to be mitigated through people, process and / or technology but none of these areas are unique to cloud. In the case of technology failure, the highly available and redundant architecture of public cloud would almost certainly improve an organisation's defences against the threat of technology failure.

### 3.5.2. CLOUD SECURITY ALLIANCE

The CSA's Cloud Computing Top Threats [130] identifies twelve threats and appropriate control guidance (via CCM) for adoption of cloud services. This paper broadly supports the more generalised (although not cloud-focused) views of NIST and the ISF. The advantage of the Top Threats paper is that it provides a decomposition of threats using the STRIDE threat model:

| **Threat** | **Description** | **STRIDE Classification** |
|---|---|---|
| Data Breach | Incident where organisation sensitive to the enterprise is breached. Confidentiality, Integrity and / or Availability of said data is compromised. | Information Disclosure |

Table 3-8 CSA Top 12 Cloud Threats **[130]**

|  |  |  |
|---|---|---|
| Weak Identity | Failure to use Multi-Factor Authentication (MFA), weak authentication | Spoofing Identity \| Tampering with Data \| Repudiation \| Information Disclosure \| Denial of Service \| Elevation of Privilege |
| Insecure APIs | APIs used to manage and interact with cloud platform – vulnerabilities in construction, weak / no encryption, authentication issues | Tampering with Data \| Repudiation \| Information Disclosure \| Elevation of Privilege |
| System and Application Vulnerabilities | Exploitable vulnerabilities in programs and Operating Systems | Spoofing Identity \| Tampering with Data \| Repudiation \| Information Disclosure \| Denial of Service \| Elevation of Privilege |
| Account Hijacking | Phishing and general fraud. Man in the Middle (MiTM) | Spoofing Identity \| Tampering with Data \| Repudiation \| Information Disclosure \| Denial of Service \| Elevation of Privilege |
| Malicious Insiders | Insider threat vulnerabilities | Spoofing Identity \| Tampering with Data \| Information Disclosure |
| Advanced Persistent Threats | Multi-phase, targeted attacks looking to establish a foothold in an environment with the goal of extracting sensitive information. | Information Disclosure \| Elevation of Privilege |
| Data Loss | The prospect of losing data | Repudiation \| Denial of Service |
| Insufficient Due Diligence | General issue for organisations but exacerbated by the structure and availability of public cloud. | Spoofing Identity \| Tampering with Data \| Repudiation \| Information Disclosure \| Denial of Service \| Elevation of Privilege |
| Abuse / Nefarious use of cloud service | Poorly secured cloud deployments, free trials and shadow IT signups. | Denial of Service |
| Denial of service | Prevention of a service being available to a customer through resource exhaustion. | Denial of Service |
| Shared Technology Issues | Use of infrastructure components in public cloud which have not been designed to securely support resource isolation and multitenancy. | Information Disclosure \| Elevation of Privilege |

Table 3-8 (Cont.) CSA Top 12 Cloud Threats [130]

This list is included for completeness. There are no new threats outlined in the CSA list that have not be gathered through other sources for this study. Where this report adds value is the focus on the types of threat aligned to an established threat model.

### 3.5.3. NIST

NIST [85] suggests that threat events *"can be expressed in highly general terms (e.g., phishing, distributed denial-of-service), in more descriptive terms using tactics, techniques, and procedures, or in highly specific terms (e.g., the names of specific information systems, technologies, organizations, roles, or locations)."*

NIST decomposes threat events through adversarial and non-adversarial categories. In **Figure 3-10**, I have detailed the adversarial threat events as outlined by NIST. NIST defines eight categories of threat event with subcategories for each category.

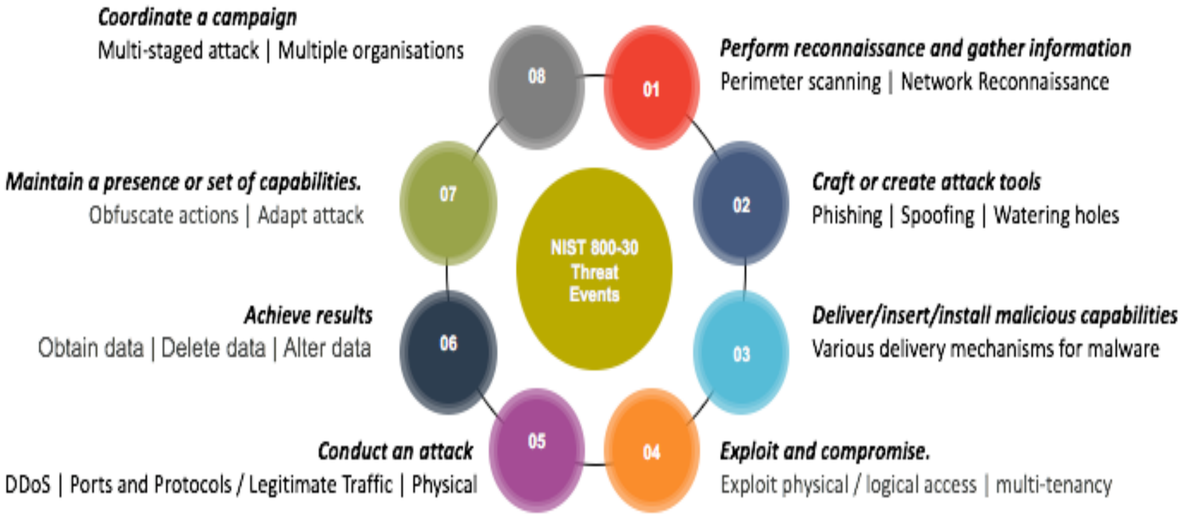

NIST 800-30 Threat Events
01
Perform reconnaissance and gather information
Perimeter scanning | Network Reconnaissance
02
Craft or create attack tools
Phishing | Spoofing | Watering holes
03
Deliver/insert/install malicious capabilities
Various delivery mechanisms for malware
04
Exploit and compromise.
Exploit physical / logical access | multi-tenancy
05
Conduct an attack
DDoS | Ports and Protocols / Legitimate Traffic | Physical
06
Achieve results
Obtain data | Delete data | Alter data
07
Maintain a presence or set of capabilities.
Obfuscate actions | Adapt attack
08
Coordinate a campaign
Multi-staged attack | Multiple organisations

**Figure 3-10 NIST 800-30 Adversarial Threat Events [107]**

Most interestingly is the explicit reference to cloud computing within their "Exploit and Compromise" phase. This is the first threat event we have identified which is a "cloud security threat" (**Figure 3-11**) and consequentially requires special attention to better understand the likelihood of this threat being realised in an enterprise production environment. As explained in **Chapter 4**, there are many forms of multitenancy and the threats associated with each are different; so are the controls available to mitigate the risks in each instance. It is also important to consider the capabilities of the threat actor; these vary significantly depending on the end goal of the attacker.

Non-adversarial threats are a combination of events initiated by an accidental human actor and those considered "*environmental*" by the ISF. What is important is the delineation between adversarial and non-adversarial. I have identified that public cloud architecture improves the customer's ability to mitigate or remove structural and environment threats. I believe I have put forward a compelling case to say that whilst accidental threats will always exist, public cloud does not introduce specific accidental threats. There will always be accidental events in any computing environment.

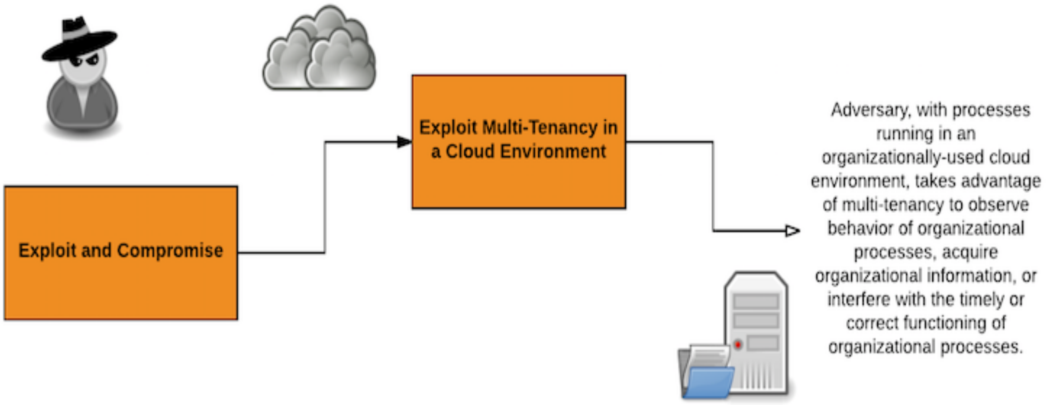
Exploit and Compromise
Exploit Multi-Tenancy in a Cloud Environment
Adversary, with processes running in an organizationally-used cloud environment, takes advantage of multi-tenancy to observe behavior of organizational processes, acquire organizational information, or interfere with the timely or correct functioning of organizational processes.

Figure 3-11 NIST Multitenant Threat

NIST state that *"Organizations can eliminate certain threat events from further consideration if no adversary with the necessary capability has been identified"* [107]. This is applicable when environmental threat events are not applicable for that region. This is an area where quantitative risk assessment can be applied. Assessing likelihood of a hurricane or flood at a datacentre is made possible through review of meteorological data readily available online.

### 3.5.4. ENISA

The ENISA model does not explicitly call out threat events; rather it focuses on vulnerabilities; without a thorough understanding of threat actors and the associated events which may exploit a vulnerability, it is an impossible task to make qualified risk decisions. ENISA do not focus on threat events, although their study of vulnerabilities [35] is thorough and can be treated as comprehensive vulnerabilities catalogue for our research.

### 3.5.5. LEGAL CONSIDERATIONS IN THE CLOUD

In **Section 3.3**, we define "Legal and Regulatory" as a category of impact warranting attention. Neither IRAM2 nor the NIST Risk Framework identify legal threat events when defining their information risk equations although threat events associated with public cloud computing can have a significant legal impact.

ENISA [35] do include legal risks as part of their cloud guidance for information security. These are documented in **Table 3-9**.

In the context of information, risk is the loss resulting from the compromise of certain attributes of its information assets (confidentiality, Integrity, availability). Risk can only be assessed once the threats are understood, vulnerabilities identified and strength of controls assessed. The legal "risks" that are frequently discussed are risks to the confidentiality, integrity or availability of information – it is the impact which has a legal association.

| **Risk** | **Explanation** | **Vulnerabilities** | **Risk Rating** |
|---|---|---|---|
| Subpoena | *In the event of the confiscation of physical hardware because of subpoena by law-enforcement agencies or civil suits, the centralisation of storage as well as shared tenancy of physical hardware means many more clients are at risk of the disclosure of their data to unwanted parties.*<br>*At the same time, it may become impossible for the agency of a single nation to confiscate 'a cloud' given pending advances around long distance hypervisor migration.* | Lack of resource isolation<br>Storage of data in multiple jurisdictions and lack of transparency<br>Lack of information on jurisdictions | High |
| Changes of jurisdiction | *Customer data may be held in multiple jurisdictions, some of which may be high risk. If data centres are in high-risk countries, e.g., those. lacking the rule of law and having an unpredictable legal framework and enforcement, autocratic police states, states that do not respect international agreements, etc, sites could be raided by local authorities and data or systems subject to enforced disclosure or seizure. Note that we are not implying here that all subpoena law-enforcement measures are unacceptable, merely that some may be so and that some legitimate seizures of hardware may affect more customers than the targets of a law-enforcement action depending on how the data is stored.* | Storage of data in multiple jurisdictions and lack of transparency<br>Lack of information on jurisdictions | High |

Table 3-9 ENISA Legal Risks [35]

| | | | |
|---|---|---|---|
| Data Protection | *It can be difficult for the cloud customer to effectively check the data processing that the cloud provider carries out, and thus be sure that the data is handled in a lawful way. It must be clear that the cloud customer will be the main person responsible for the processing of personal data, even when such processing is carried out by the cloud provider in its role of external processor. Failure to comply with data protection law may lead to administrative, civil and criminal sanctions, which vary from country to country, for the data controller.* | Storage of data in multiple jurisdictions and lack of transparency about this<br>Lack of information on jurisdictions | High |
| Licensing Risks | *Licensing conditions, such as per-seat agreements, and online licensing checks may become unworkable in a cloud environment. For example, if software is charged on a per instance basis every time a new machine is instantiated then the cloud customer's licensing costs may increase exponentially even though they are using the same number of machine instances for the same duration. In the case of PaaS and IaaS, there is the possibility for creating original work in the cloud.* | Lack of completeness and transparency in terms of use. | Medium |

Table 3-9 (Cont.) ENISA Legal Risks [35]

We should pay close attention to the vulnerabilities outlined in **Table 3-9**; excluding "resource isolation", the remaining vulnerabilities are inherent within people and process. A lack of transparency in a terms of use or misunderstanding of the laws and regulations of jurisdictions are not insurmountable challenges to resolve, nor are they technical impediments to public cloud adoption.

#### 3.5.5.1. GENERAL DATA PROTECTION REGULATION (GDPR)

The General Data Protection Regulation (GDPR) (Regulation (EU) 2016/679) is intended to protect the data protection rights of citizens within the European Union (EU). GDPR comes into force in May 2018 and unlike previous EU-wide data protection artefacts, it is regulation, not a directive. This means that member states must implement the terms verbatim. Previous legislation allowed countries to interpret clauses and created disparity in the application of the previous EU wide mandate: EU Directive 95/45 EC – Data Protection Directive.

Unlike directives before it, GDPR carries with it heavy penalties for a reach of personally identifiable information (PII). These fines can amount to 4% of an organisation's global turnover [131].

Cloud computing is often cited as being an impediment to GDPR compliance [132, 133] because of some of the GDPR clauses which place a greater responsibility on organisations to understand their data flows and location of sensitive information:

**The Right to Erasure**: Article 17 [131] states that a subject has the right to request that all PII a company has relating to an individual is deleted.

**Data Portability**: Article 18 [131] states that an individual has the right to transfer their data from one processor to another without impediment from the data controller.

**Breach Reporting**: A consideration for the cloud customer, not the data subject; a breach of PII must be reporting to the local Supervisory Authority (SA) (UK – this would be the Information Commissioner) within 72 hours of a breach being discovered. This is covered in Article 33 [131].

There is certainly a need for a comprehensive understanding of roles and responsibilities between the customer and CSP although, irrespective of GDPR, holistic and mature information security management requires an organisation to understand and document their data assets and understand flows of information. As an example, ISO 27002 [109] has clear guidance regarding supplier relationships and the need to understand the types and volume of data residing in third party locations. TOGAF Version 9 [134] makes explicit reference to "Data Migration Diagrams" which are used to provide a "visual representation of the spread of sources/targets and serve as a tool for data auditing and establishing traceability" [135]. What has changed with GDPR is the impact of non-compliance. 4% of global revenue is a figure which makes board representatives listen. There are "cloud considerations" in a GDPR context including those for the CSP. CSPs are now responsible for the data that they process and must take ownership for the protection of information under their control [136]. I assert that this shift to a shared responsibility model will encourage CSPs to evidence strong information security and data privacy controls.

GDPR does bring with is considerations regarding data sovereignty but requirements surrounding where an organisation stores their data pre-date the regulation. The use of public cloud services certainly requires an organisation to invest the time in due diligence to understand precisely where cloud service providers will store their information. Mature CSPs are acutely aware of the legal ramifications of data retention in multiple jurisdictions. Amazon [29, pp. 9-15] can evidence a highly-available, fault tolerance architecture which provides the customer with assurances their data will be retained in regions they select and with technical security controls commensurate to the classification of the information being stored / processed.

Much like the conversation around risk categories, legal and compliance could be argued as a vulnerability inherent in cloud computing or a benefit. If organisations neglect their responsibilities as data owners and blindly hand their data to any CSP, they increase the likelihood of failing to comply with legal and compliance requirements. Conversely, public cloud services can provide pre-attested environments for regulations such as PCI-DSS and ISO 27001 which improve an enterprise's compliance posture [66].

Having spent 18 years in the technology field, I assert that regulatory compliance frameworks are often misunderstood in our industry, they should be treated as a minimal acceptable level of security, not a gold standard. Progressive thinking aims to reduce risk first and comply with regulations second. Clearly, one should achieve the other although as an example, famously Target were PCI-DSS compliant at the time of their infamous 2013 data breach [137].

## 3.6. CLOUD VULNERABILITIES

For a risk to manifest itself, a vulnerability needs to be exploited. The ISF defines a vulnerability as: *"…a weakness in people, process or technology in an environment, which could be exploited by one or more threats."* [106] One of the primary objectives of this paper was to assess the risks of cloud computing and to identify which were the result of people, process and technology. Upon reflection, it is in fact vulnerabilities which need to be defined as applicable to people, process or technology. I can support this assertion through the research performed in this thesis: having reviewed preeminent industry papers surrounding information risk, I conclude that on-premise and public cloud threat actors and their associated events are comparable. In **Section 3.5** I documented, compared the categories of threat event provided by the ISF, NIST and the CSA. None of these events could be credibly positioned as only applicable to public cloud computing.

Risk, in an information risk context, is concerned with the confidentiality, integrity and availability of information as it always was pre-cloud. Cloud-based, multitenant architectures introduce novel vulnerabilities and potentially exacerbate the impact of a successful compromise of information. The CSA supports my assertion claiming that cloud has created new security vulnerabilities as well as amplifying existing ones [130]. If we can identify and remove cloud-centric vulnerabilities, we remove the associated risks.

There is no commonly agreed definition of the mandatory components to a risk equation. ENISA's risks, identified in **Section 3.3** are comprehensively documented and compartmentalised into "Technical", "Policy & Organisational" and "Legal", this includes a breakdown of associated vulnerabilities. ENISA calls out 31 vulnerabilities which are "specific to cloud" and 22 vulnerabilities which are present in a cloud ecosystem although also resident in other computing models.

I wanted to understand these 53 vulnerabilities in more detail. If 31 vulnerabilities were uniquely attributable to cloud computing then (cloud) adoption sounds littered with opportunities for exploitation and ultimately data loss. Upon my evaluation of ENISA's vulnerabilities, I challenge that many of the "cloud-specific vulnerabilities" are in fact vulnerabilities which present themselves in any modern computing architecture. To support my assertion, I would like to introduce **Appendix A: ENISA Cloud Vulnerability Assessment** where I address each of the ENISA vulnerabilities in turn and provide commentary regarding each risk's context and appropriate mitigations.

Having already introduced the concept of a "predisposing condition", I take this concept and apply a categorisation scheme based thereon. This is applied across the findings in **Appendix A: ENISA Cloud Vulnerability Assessment**. This classification is applied to vulnerabilities and these are detailed for completeness in **Table 3-10** (below):

| **Category** | **Description** |
|---|---|
| Unique to Cloud | Vulnerabilities which do not exist in previous computing models – these are scarce but warrant explicit attention. |
| Exacerbated by Cloud | Vulnerabilities made inherently and inextricably more severe through cloud: Limited and acknowledged in this work. |
| General | Vulnerabilities which may reside within the environment but cloud residency is not a consideration. This category accounts for most vulnerabilities prevalent in computing today. |

Table 3-10 Cloud Vulnerability Classification

Undoubtedly, due to the multitenant architecture of public cloud computing, there will be vulnerabilities which are not present in an environment solely provisioned for a single organisation although through my analysis I conclude that there are significantly fewer unique vulnerabilities which are truly introduced by cloud computing. At a macro level, I assert that two truly exclusive vulnerabilities are cloud-based:

| Cloud Vulnerability | Explanation |
|---|---|
| Resource Sharing | Shared services model of cloud introduces vulnerabilities concerned with availability of service and data leakage. Inadequate resource isolation could result in side channel attacks and tenants breaking outside the confines of their environment and extracting information from other tenants. |
| Data Residency | Cloud introduces fresh requirements to understand the impact of data being potentially domiciled in locations outside of the company datacentre and, on occasion, in geographical locations with various data privacy laws and regulations. Whilst only applicable in public and hybrid cloud scenarios, this truly is a unique cloud vulnerability. |

Table 3-11 Cloud-Specific Vulnerabilities

I cover the legal risks associated with cloud computing in **Section 3.5.5,** I will continue to explore resource sharing in detail across **Chapter 4**.

It is important to ask the following questions when assessing the severity and applicability of a vulnerability:

- Is this vulnerability exclusively cloud-based?
- Is there a threat event which can exploit this vulnerability?
- Are their controls available which mitigate the impact of the vulnerability being exploited?

**Appendix A: ENISA Cloud Vulnerability Assessment** breaks the 53 Cloud Security Vulnerabilities outlined by ENISA and categorises them cloud-initiated, cloud-exacerbated or simply "general" vulnerabilities present in all environments. For the purposes of analysis, I have ignored deployment and service model considerations. This was an intentional measure; I wanted to analyse if vulnerabilities exist in the inherent construction of cloud before asserting if public/private or SaaS, IaaS, PaaS introduces a greater level of risk. The appendix includes details and justifications for my vulnerability classifications although an overview is included below for completeness:

| Vulnerability Category | Instances |
|---|---|
| Unique to Cloud | 4 (two legal, two resource sharing) |
| Exacerbated by Cloud | 13 |
| General | 14 |

Table 3-12 ENISA Cloud Vulnerabilities

I acknowledge that the ENISA findings are thorough. Any organisation embarking on project dealing with sensitive information would be well-served review the paper. What my findings highlight is that over half of the vulnerabilities would be present in any contemporary technology environment. I explore this further in **Section 4.3**.

The volume of vulnerabilities "exacerbated by cloud" was expected and intrinsically linked to resource sharing. As I explained in **Section 3.3**, the sharing of resources can exacerbate the impact should a threat actor exploit a vulnerability. Several vulnerabilities are also defined as being exacerbated by cloud as they require process changes in organisations to deal with cloud. A useful example being "User Provisioning Services"; most (if not all) mature CSPs offer single sign on capabilities. Options for SAML integration are included, as is Active Directory (AD) authentication. This is therefore not a technology problem. Organisations may however have to alter joiner, mover, leaver (JML) processes to deal with provisioning in the cloud.

The most interesting and insightful information I took away from this analysis was the number of exacerbated vulnerabilities that are a result of a need for process change as opposed to any technical vulnerability. Based on these findings, I can assert that organisations would wise to focus on operational processes change when dealing with public cloud adoption.

I feel that additional commentary is necessary regarding several vulnerabilities that ENISA have defined as being "Non-Cloud". I would strongly challenge this categorisation for four of their 22 vulnerabilities. ENISA consider "security awareness", "vetting of processes (and staff)", "unclear roles and responsibilities" and "liability uncertainty" as all being "non-cloud". I qualify my challenge through the research conducted as part of this paper. The CSA [34] identify that when organisations move to cloud, the responsibilities for securing the environment shift from exclusively "the customer" to a shared model (customer and CSP). I have repeatedly referenced Amazon's [138] similarly supportive position. This shift in responsibility brings with it a need to redefine responsibilities and ensure that CSP staff are vetted and capable of performing tasks previously carried out on-premise.

## 3.7. MITIGATING AND MINIMISING RISK

As covered in **Section 3.1**, risk is the result of a threat actor initiating a threat event, to exploit a vulnerability thus causing adverse impact. To mitigate or avoid risk, organisations look to security controls and safeguards to reduce likelihood, impact or both.

Controls and safeguards exist everywhere, in the physical and digital world. For example: seatbelts are mandatory in the United Kingdom when driving an automobile. Safety warnings are placed on household materials containing dangerous substances and food items contain use by dates. All these measures are there to minimise either the impact (seat belts) or likelihood (warning labels) of a threat event occurring.

In the world of information security, risks to the compromise of confidentiality, integrity or availability of data require the application of security controls to lessen both the impact and / or the likelihood of a data breach. To ensure traceability of requirements, risk management generally concerns looking at risk from two perspectives which align to the ultimate impact of the risk [106, p. 23].

**Inherent Impact**: Impact or likelihood of a threat event occurring without the application of controls.

**Residual Impact**: The remaining exposure to an organisation after the application of controls.

In an information security context, removing all risks to the confidentiality, integrity and availability of data isn't possible. All interactions with technology inherently require the acceptance of a degree of risk; no network is infallible and in some scenarios, the costs of defence against highly-sophisticated, motivated and financially-backed actors are just too high. This is one of the reasons why conscious risk management is such a valuable discipline. Accepting risk is a perfectly viable solution, assuming it is a qualified assessment.

Industry guidance exists for the application of safeguards to mitigate risk. The CSA provides their "Cloud Controls Matrix" (CCM) [87] which is produced to give organisations pragmatic guidance in the application of information security controls for cloud models. The CCM provides organisations with a fundamental set of guiding principles for the application of security controls for cloud, across all major deployment and service models. The CSA have provided a threat-centric cloud assessment [130] which applies the STRIDE [139] threat model and control traceability. The CCM significantly aides an organisation in demystifying regulatory compliance headaches brought about by a patchwork of cloud standards as it aligns across many regulatory frameworks including (COBIT, HIPPA, ISO 27001, PCI-DSS v3).

The ISF IRAM2 [106] provides a comprehensive set of controls to mitigate and remove the threats associated with the comprise of CIA. As with the CSA CCM, the skills of a security SME are still required to digest conceptual requirements into meaningful, technical security controls. Tangentially, this situation highlights the demand for security specialists at different layers of a security architecture framework. The SABSA Reference Model [82] covers this in detail. The IRAM2 model provides a qualitative risk assessment framework for assessing the likelihood and impact of a CIA compromise. They (ISF) use the concept of "control strength" to assess how thorough / comprehensive a control is applied to mitigate a threat event exploiting a vulnerability. Controls should be applied to mitigate and remove vulnerabilities; the application of controls for threats that do not exist wastes money and increase operational expenditure.

The CCM is a critically important document to ensure holistic coverage of focus areas but the efficacy of a solution is paramount. A box can be ticked to confirm that anti-malware solutions exist but are they effective and operationally maintained? With the above in mind, we need to ensure that a combination of guiding principles and qualitative control assessment can be adopted for our study. NIST [83] leverage the CCM as part of the NCCSRA.

It is important that we consider several control fundamentals appropriate for our study. These relate to the efficacy of a control. Where compliance frameworks have previously fallen-down in their 'tick-box' nature – simply having a control is of limited value if it is not implemented correctly and actively maintained. Controls can take many forms across people, process and technology. Controls can be automated or manual. Controls are also implemented at various phases of a cyber-attack / data exfiltration – preventative, detective and responsive. Irrespective of control type, control strength [106] is an important attribute we will use in our analysis:

***Effective***: Reliable evidence that the control should substantially mitigate the risks identified.

***Partially***: Some evidence that the control should substantially mitigate the risks identified.

***Ineffective*:** No available evidence about control effectiveness or evidence control is insufficient

Assessing inherent impact is of limited value in isolation. It does however provide business risk stakeholders with visibility of the impact (cost, time, legal implications) of a "do nothing" approach. "Risk Acceptance" is a perfectly reasonable course of action in some situations but only once impact and likelihood are fully understood by person(s) responsible for the information in question. Residual impact should always be presented which details the controls and mitigations available to remove or limit the impact and or likelihood of a data compromise.

For the purposes of penetration testing and application code review, it is often beneficial to understand inherent (unmitigated) vulnerabilities as root-cause remediation (prior to productionisation) is always preferable to the application of retrospective safeguards although for a risk assessment, the impact and likelihood can often be mitigated to a level palatable to the organisation.

There is a natural assumption to make that all controls are technical. IT and cloud are both very technical areas of specialisation although as I have explored in **Section 3.6** many areas of weakness emanate from people or process. Often these vulnerabilities can only be migrated with people or process controls.

## 3.8. CROSS-INDUSTRY CLOUD ADOPTION

Industry regulation is required in all sectors to provide consumer confidence and protect the interests (and sometimes lives) of employees, partners and customers. Some industries are more heavily regulated than others; especially in areas related to IT and data security. If established industry regulation prohibits the use of public cloud, this would suggest to me that cloud architecture is inappropriate for certain forms of data.

One of the objectives of the paper relates to the secure implementation of public cloud for organisations:

***To understand if public cloud can be implemented securely for the enterprise.***

To support my risk analysis in **Chapter 4**, I wanted to find answers to the following research questions:

- ***Is public cloud adoption prevalent in specific industry verticals?***
- ***Is public cloud only suitable for certain sensitivity of data?***
- ***Do mainstream compliance requirements preclude the use of public cloud?***
- ***Are global businesses really adopting public cloud to replace their on-premise infrastructure?***

Industry reports [37] suggest that over 50% of IT teams are using services running in AWS and adoption of the Microsoft Azure platform is also on the rise [37, 28]. I wanted to find out if industries were bucking this trend and avoiding cloud. A survey conducted by The Economist [27] suggests cloud computing is being adopted in all major industry verticals (**Figure 3-12**).

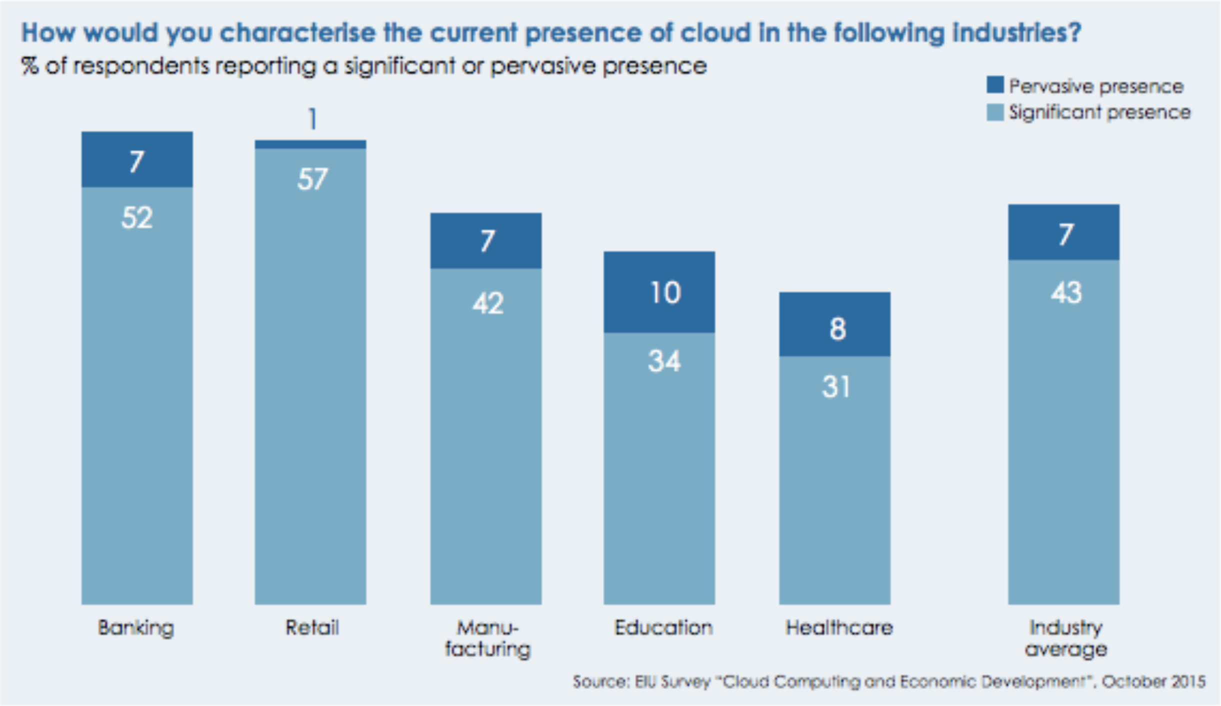

Source: EIU Survey "Cloud Computing and Economic Development", October 2015

Figure 3-12 Cloud Adoption Across Industries [27]

Another key area of consideration is regulation. If regulatory compliance precludes the use of public cloud, it cannot be considered a viable solution for certain forms of sensitive information. To assess if regulation prevents a company from adopting cloud, I have researched applicable regulation associated with banking. My rationale being that banking is a heavily regulated industry and data breaches have historically targeted payment data [137]. The Financial Conduct Authority (FCA) regulates financial services firms and the markets in the UK. The

FCA have released specific guidance regarding the use of cloud services, in which they confirm that is "*no fundamental reason why cloud services (including public cloud services) cannot be implemented, with appropriate consideration, in a manner that complies with our rules*" [140].

As this thesis is time-bound, it is not possible to analyse an item of regulation in each sector, instead I have selected a specific item which would have the broadest applicability in as many industries as possible. The Payment Card Institute Data Security Standard (PCI-DSS) [141] was created in 2004 and is an information security standard applicable to all organisations who process store or transmit credit card information. In 2017, this accounts for the overwhelming majority of companies in almost all industries. Much like the position provided by the FCA, PCI-DSS does not preclude the use of public cloud to store or process payment card information. Due to the increased adoption of cloud services, the PCI Security Standards Council have released specific guidance [142] to ensure that organisations understand how to implement payment services in the cloud. **Figure 3-13** details an example shared responsibilities model for compliance with PCI-DSS across deployment models.

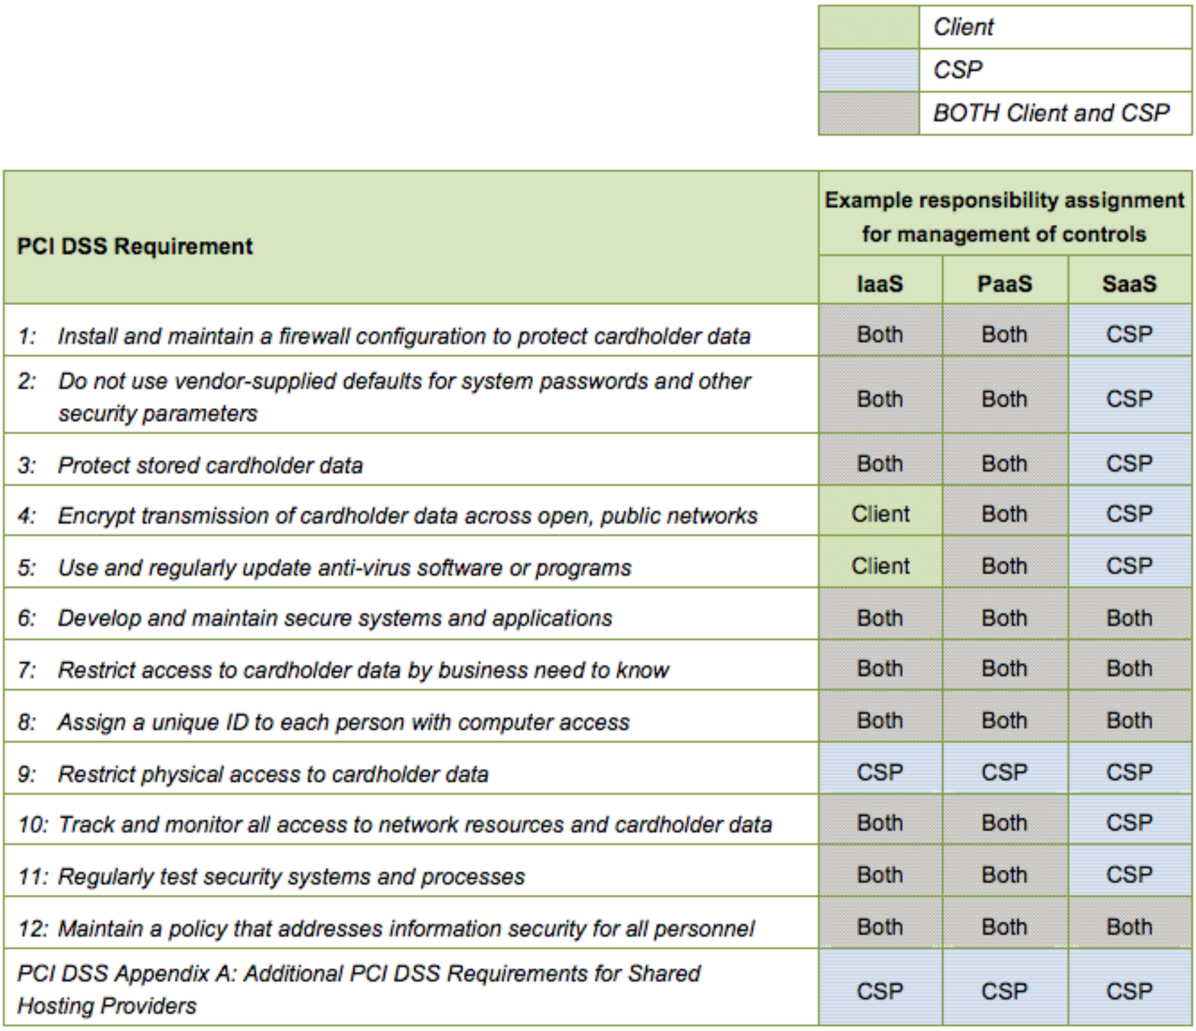

| | |
|---|---|
| 🟩 | *Client* |
| 🟦 | *CSP* |
| ▨ | *BOTH Client and CSP* |

| **PCI DSS Requirement** | **Example responsibility assignment for management of controls** | | |
|---|---|---|---|
| | **IaaS** | **PaaS** | **SaaS** |
| *1: Install and maintain a firewall configuration to protect cardholder data* | Both | Both | CSP |
| *2: Do not use vendor-supplied defaults for system passwords and other security parameters* | Both | Both | CSP |
| *3: Protect stored cardholder data* | Both | Both | CSP |
| *4: Encrypt transmission of cardholder data across open, public networks* | Client | Both | CSP |
| *5: Use and regularly update anti-virus software or programs* | Client | Both | CSP |
| *6: Develop and maintain secure systems and applications* | Both | Both | Both |
| *7: Restrict access to cardholder data by business need to know* | Both | Both | Both |
| *8: Assign a unique ID to each person with computer access* | Both | Both | Both |
| *9: Restrict physical access to cardholder data* | CSP | CSP | CSP |
| *10: Track and monitor all access to network resources and cardholder data* | Both | Both | CSP |
| *11: Regularly test security systems and processes* | Both | Both | CSP |
| *12: Maintain a policy that addresses information security for all personnel* | Both | Both | Both |
| *PCI DSS Appendix A: Additional PCI DSS Requirements for Shared Hosting Providers* | CSP | CSP | CSP |

Figure 3-13 PCI-DSS Cloud Responsibilities Mapping [142]

What my research in this section has identified is that location should not be the defining factor when assessing information risk. Both the FCA and the PCI Council support the use of public cloud so long as an understanding of roles and responsibilities exists and that appropriate controls can be applied commensurate with the classification of information being stored or processed.

AWS have a comprehensive set of case studies [143] identifying customer adoption of public cloud across industry verticals and for production systems processing sensitive information. The rhetoric suggesting that public cloud in inappropriate for certain forms of data is inaccurate.

## 3.9. SHADOW IT: THE CATCH-ALL VULNERABILITY

*Among the most significant security risks associated with cloud computing is the tendency to bypass information technology (IT) departments and information officers.* [130]

***Cloud Security Alliance – 2016***

My personal definition: Shadow IT is concerned with the unauthorised use of applications and services procured and deployed without the approval and / or visibility of the IT department. Information technology (IT) departments have been battling against the unauthorised use of software for decades. The convenience and service-based nature of cloud exacerbates the issue of shadow IT and creates off-premise data stores but cloud computing did not introduce shadow IT.

The challenges that Shadow IT introduce with cloud are significant. I was surprised that more focus was not given in ENISA's vulnerability categorisation for cloud computing [35]. Whilst Shadow IT is not uniquely introduced by cloud, it does require amendments to business process and the application of security controls. For these reasons, we will address the challenges that public cloud Shadow IT bring in **Chapter 5**.

## 3.10. CONCLUSION

I wrote in the introduction to this thesis that cloud is a new paradigm, my research across this chapter has made me re-evaluate this statement. Whilst cloud introduces new opportunities for organisations to leverage service-based IT services, it is not conceptually different to business relationships of years gone by.

Organisations have relied on third-parties for centuries to provide services with the inherent consideration of storing and processing information. With cloud computing, we are now storing digital information in new locations but it is unfair to assert that prior to cloud, organisations were insular, secretive enterprises with zero exposure to the outside world.

One of public cloud's greatest benefits is also one of its most talked-about vulnerabilities: "shadow IT". We have identified in this chapter that cloud computing requires a fresh approach to process and people training. Organisations need to evaluate their procurement and supply-chain processes. The example of a marketing team procuring application resource on AWS and proceeding to upload gigabytes of sensitive information is a by-design characteristic of public cloud: on-demand access.

Neophobic views have no place in the technology world; things move too fast. There are security considerations in the adoption of public cloud but due diligence is a requirement of any business project. The security controls, capabilities and safeguards from established vendors in public cloud allow organisations to ensure a commensurate level of protection in the cloud. I have evidenced (**Table 3-3**) that public cloud adoption can make organisations more secure as they can benefit from carrier-grade, best-of-breed logical and physical controls which are cost-prohibitive for some. Environments are built to comply with legal and regulatory requirements from inception [29]. Public cloud service providers have security as their core business – they therefore take this very seriously.

Having compared preeminent information risk management frameworks [106, 107], it is evident that an absence of "cloud-focus" can exacerbate a natural aversion to outsourcing of control for data. Whilst significant steps have been taken to address the shortfall in cloud standards documentation [96], there still an opportunity for contextualised risk assessment frameworks with a cloud focus.  As we discussed in **Chapter2**, public cloud introduces new actors, architectural considerations and interpretations of “internal” and “external”.

My studies have highlighted that public cloud computing introduces a predisposing condition which could exacerbate the financial, legal or reputational impact of the compromise of CIA for corporate data.  Public cloud computing does not introduce new threat events nor are the vulnerabilities inherent in public cloud reserved only for public cloud environments.  The vulnerabilities which are compounded by public cloud adoption relate to inadequate resource isolation; a subject which will be thoroughly decomposed and critically-analysed in the following chapter.

# 4. MULTITENANCY AND RESOURCE ISOLATION: THREATS AND VULNERABILITIES

This chapter will form the practical analysis component of my thesis. Through my research thus far, I have identified requirements for organisations to amend processes and educate users into the methods for procuring, supporting and using public cloud. Failure to follow these processes, such as understanding the legal implications of housing data in different international jurisdictions, can have significant impacts on an organisation. Whilst these process changes are acutely important, they do not present vulnerabilities inherent in the construction of cloud. The issues present themselves as a repercussion of (technology) change.

It is important that risk decisions consider not only threat actors, threat events and vulnerabilities but equally important are the controls, mitigations and compensations which may lower risk to a tolerable level in an enterprise. The motivations of threat actors are also incredibly important; information security resources are finite and it is important that organisations are investing time and money mitigating vulnerabilities that will be exploited, via a prioritised approach.

In this chapter, we will look at the vulnerabilities and threat events which exist for multitenancy. A risk assessment, using IRAM2, will be applied for each threat event associated with resource isolation. We will conclude by comparing the likelihood of resource isolation / multitenancy threats with those of other, traditional security vulnerabilities which have been exploited to compromise the CIA of information.

As this chapter introduces several technical concepts and complex risk management equations, I considered it prudent to provide some explicit signposting of my objectives and the methods I plan to use to achieve these. I assert that the concept of resource sharing, and by association, an inability to achieve appropriate resource isolation, is the primary vulnerability introduced by public cloud architecture. I therefore challenged myself to answer several pertinent questions associated with public cloud:

1) ***What is multitenancy and more specifically: Is this coarse-grained definition appropriate for all public cloud implementation?***

2) ***Is multitenancy the exclusive reserve of public cloud?***

3) ***Are the vulnerabilities associated with resource sharing exploitable with a reasonable degree of likelihood by range of threat actors with varying levels of skill and persistence?***

4) ***Are the vulnerabilities of multitenancy appropriately contextualised? Do other attack paths exist which are more likely exploitable by all / any of the threat actors used in this study?***

Through answering these research questions, I can satisfy several objectives outlined in **Section 1.2**. Questions 1 & 2 will be answered through personal research into definitions of multitenancy from several sources and research into contemporary datacentre architecture. Questions 3 & 4 will be answered through a thorough research exploration of threat events which exploit multitenancy vulnerabilities followed by the practical application of the ISF IRAM2 model to assess likelihood of initiation.

## 4.1. WHAT IS MULTITENANCY?

Multitenancy between organisations, at its core, is what separates cloud computing from paradigms of the past. The most significant security vulnerabilities associated with cloud computing are the sharing of resources and the housing of data in remote locations. Gartner [144] goes as far as to say:

*"Enterprise IT organisations' top objection to cloud computing is that by allowing outside data centres to handle their business data they are potentially exposing their data to competitors or other intruders."*

Definitions of multitenancy across academia and the technology industry are far less varied than for defining cloud computing. Brown et al [145] play on an everyday analogy: *"Multitenancy is similar in nature to multiple families in the same condominium. Generally speaking each has their own space, however there is a risk that one family may have access to another families space or information"*. The Open Web Application Security Project (OWASP) [110] provides the most appropriate definition for this study*: "Multi-tenancy in cloud means sharing of resources and services to run software instances serving multiple consumers and client organizations (tenants). It means physical resources (such as computing, networking, storage) and services are shared, also the administrative functionality and support may also be shared. One of the big driver for providers is to reduce cost by sharing and reusing resources among tenants."* OWASP's definition explicitly references the sharing of not only technology but also administrative and support functions; key considerations and areas of potential vulnerability.

In the case of multitenancy, there are people, process and technology vulnerabilities which form the basis of an aversion to cloud adoption [35]. Multitenancy is a concept as opposed to a specific technological paradigm. Virtualisation and the sharing of resources are intrinsically-linked to multitenancy. It is my opinion, backed up by the research in this paper, that virtualisation and resource sharing are the technical building blocks which deliver multitenancy for public cloud. I also disagree with NIST's exclusion of multitenancy from its "Essential Characteristics" [32] as without multitenancy, the economies of scale and flexibility are removed for the CSP making cloud, generally, unprofitable.

The remainder of this section will focus on the technical implementation of multitenancy: Its instantiation through virtualisation and converged infrastructure shared by tenants from different organisations. I explained in **Section 3.1** that multitenancy can exacerbate the impact of information compromise as resources are being shared; I wanted to find out how likely such an event is given the ever-growing use of public cloud and maturity of cloud security controls [115].

## 4.2. WHY FOCUS ON MULTITENANCY?

As explored in Chapter 3, cloud introduces several vulnerabilities and it exacerbates the impact of certain pre-existing vulnerabilities (see **Appendix A**). Of the technical vulnerabilities inherently presented in cloud, a lack of resource isolation carries with it the most serious impact if exploited. ENISA classifies the risk[14] of "isolation failure" (**Figure 4-1**) as "High" and describes this class of risk to include *"the failure of mechanisms separating storage, memory, routing, and even reputation between different tenants of the shared infrastructure (e.g., so-called guest-hopping attacks, SQL injection attacks exposing multiple customers' data stored in the same table, and side channel attacks")* [35].

[14] The term "risk" is taken from ENISA. It is my opinion that that isolation failure is a high-level aggregation of process and technical vulnerabilities. The risks remain those of an impact to the confidentiality, integrity and / or availability of organizational data assets.

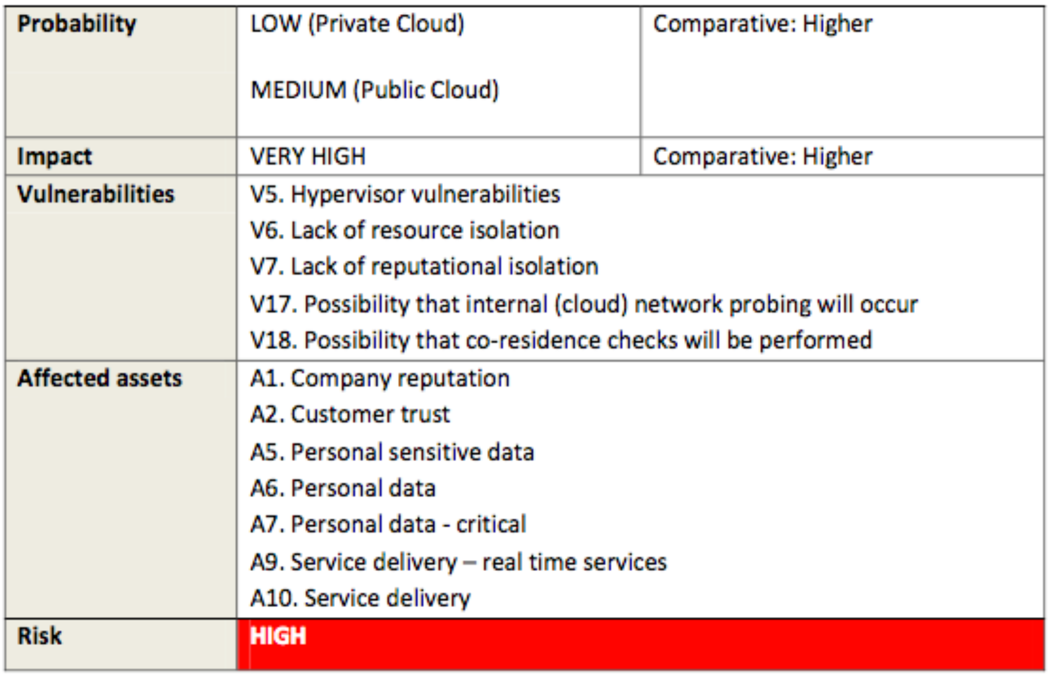

| | | |
|---|---|---|
| **Probability** | LOW (Private Cloud)<br>MEDIUM (Public Cloud) | Comparative: Higher |
| **Impact** | VERY HIGH | Comparative: Higher |
| **Vulnerabilities** | V5. Hypervisor vulnerabilities<br>V6. Lack of resource isolation<br>V7. Lack of reputational isolation<br>V17. Possibility that internal (cloud) network probing will occur<br>V18. Possibility that co-residence checks will be performed | |
| **Affected assets** | A1. Company reputation<br>A2. Customer trust<br>A5. Personal sensitive data<br>A6. Personal data<br>A7. Personal data - critical<br>A9. Service delivery – real time services<br>A10. Service delivery | |
| **Risk** | **HIGH** | |

Figure 4-1 ENISA Isolation Risk [35]

The impact of a compromise to resource tenancy is undeniable. Should this occur, other tenants may obtain access to sensitive information or cause operational and availability impact to production systems; however, ENISA's definition of "High" for probability / likelihood requires further analysis. This would depend on the attributes of the threat actor and the exploitability of any vulnerability.

In **Chapter 2**, I assessed the service and deployment models for cloud computing. The threat events that exploit isolation vulnerabilities have a differing risk profile depending on the combination of service and deployment model. ENISA (**Figure 4-1**) explicitly reference probability across deployment models: Private and Public Cloud. Accepting insider threat as a credible vector in most enterprise environments, the paper would be better served identifying service models (in a probability context).

## 4.3. CONTEMPORARY DATA CENTRE STRATEGY AND THE ROLE OF MULTITENANCY

In the introduction to this chapter, I posed a question surrounding the exclusivity of multitenant architectures:

"***Is multitenancy the exclusive reserve of public cloud***?"

To answer this question, I have researched several industry definitions and will include these findings in this section.

As acknowledged by the International Data Corporation (IDC), the role of Information Technology (IT) is dramatically changing within the modern business: *"In the past decade, information technology (IT) evolved from an enabler of back-office business processes to the very foundation of a modern business. In the increasingly digital and mobile world, the datacentre is often the first and most frequent point of contact with customers"* [146].

Terms like "converged infrastructure" and "software defined networking" (SDN) are prevalent in the contemporary technology ecosystem. These solutions rely on virtualisation and horizontal scaling to achieve the flexibility and elasticity needs of the digital organisation. If an organisation retains physical hardware to support a traditional n-tier architecture [15] (one physical appliance for each layer), they will be spending significantly more than their competitors. These cost and efficiency benefits are seeing virtualisation as foundational in almost all IT deployments.

Whether on-premise or in the public cloud, the technical architecture of our platforms today is very similar. Virtualised, converged infrastructure solutions are shortening the time to provision servers and to deploy applications. The IDC survey results in **Figure 4-2** outline the operational benefits of EMC VBlock [147] architecture and evidence significant advantages to moving away from physical servers on a per application basis.

Convergence approaches are bringing significant savings not only in time-to-market but also in operational efficiencies. Openstack [148] (**Figure 4-3**) is an open-source technology platform which allows an enterprise to control pools of network, storage and compute via Openstack API or a management dashboard. Openstack technologies can be applied to public and private cloud. Amazon AWS adopts a similar single, unified management pane approach to VM management. Irrespective of the service and deployment model of infrastructure and applications, orchestration via centralised management is saving time, improving consistency (and therefore security) and lowering operational expenditure.

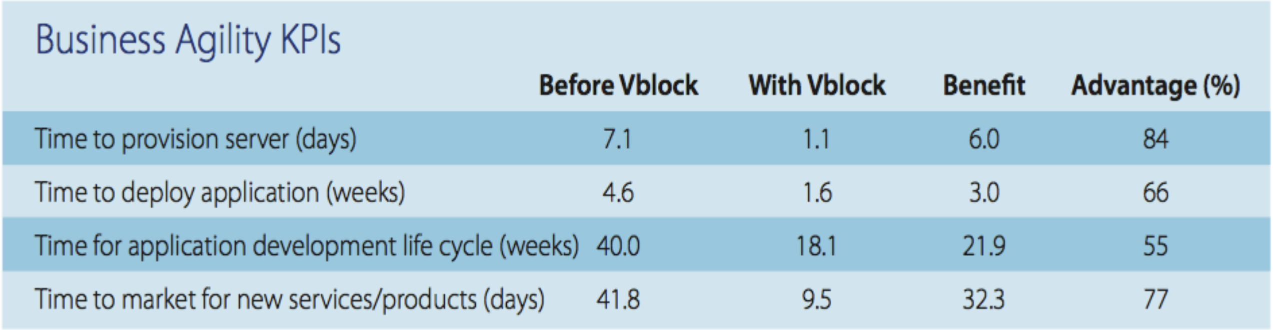

Business Agility KPIs

| | Before Vblock | With Vblock | Benefit | Advantage (%) |
|---|---|---|---|---|
| Time to provision server (days) | 7.1 | 1.1 | 6.0 | 84 |
| Time to deploy application (weeks) | 4.6 | 1.6 | 3.0 | 66 |
| Time for application development life cycle (weeks) | 40.0 | 18.1 | 21.9 | 55 |
| Time to market for new services/products (days) | 41.8 | 9.5 | 32.3 | 77 |

**Figure 4-2 Business Agility KPIs: Converged, Virtualised Infrastructure [146]**

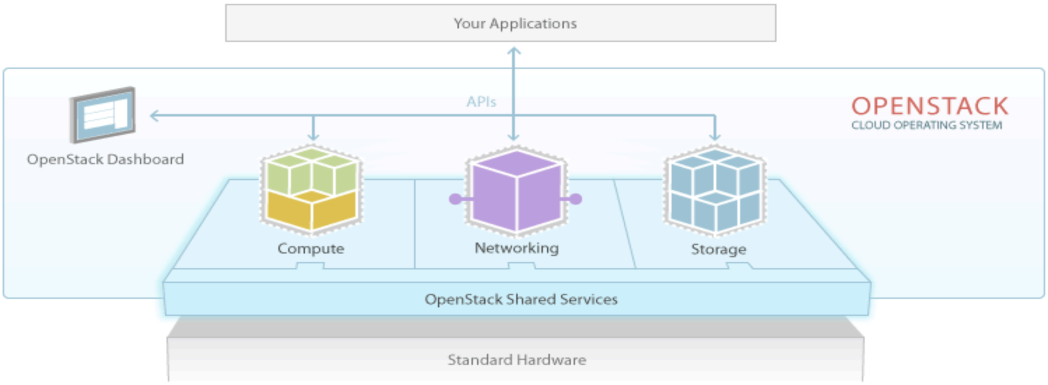

**Figure 4-3 Openstack Cloud Operating System [148]**

We must refer to some of our definitions from **Chapter 2**; are we moving to a world where all on-premise infrastructure is fast becoming "private cloud"? In Introduction, we defined cloud as: *"…a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."* [41]. This isn't an optional requirement, the business benefits brought about my virtualised, converged infrastructure and networks are core components of an organisation's journey to becoming "digitally native". If all infrastructure is becoming converged and virtualised – where and how do we delineate cloud from "non-cloud"? The only difference now is the fact that in some converged, virtualised environments, the tenants are distinct organisations as opposed to departments within the same company.

MacVittie [149] supports my theory stating that: *"there are "customer" level distinctions to be made internal to an organization, particularly a large one.". In the same work, the author explains that: "each of these "entities" can – and often does – have its own budgets and thus dedicated resources"* [149]. This is important because whilst multiple applications within an organisation could be deployed on the same physical hardware, they may well have different requirements for the confidentiality, integrity and availability of information. In a public versus private discussion we are often replacing the names of our tenants from "customer" to "department" or "function" in public and private multitenancy (respectively) (**Figure 4-4**).

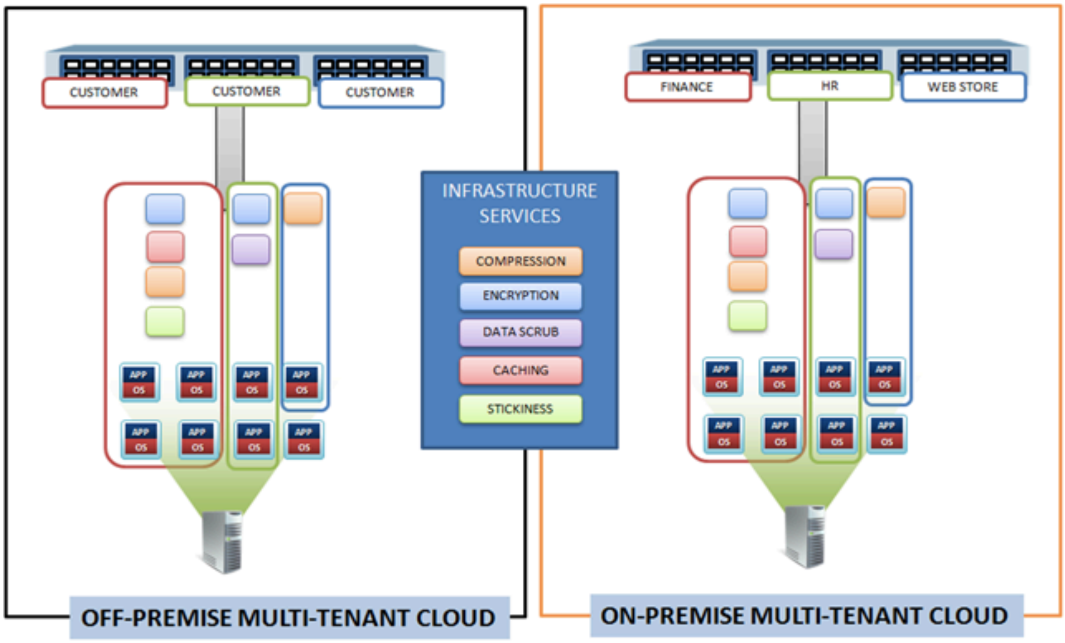

Figure 4-4 Public / Private Multitenancy Shared Services [149]

The threat events and vulnerabilities for private cloud are identical to those presented in public cloud. I make this assertion having thoroughly reviewed seminal work in this space which provides clear definitions of cloud computing along with the associated threat events and vulnerabilities [87, 35, 83].

Another important consideration when attempting to delineate public and private cloud is the prevalence in 2017 of enterprise federation solutions. Identity federation and single sign-on (SSO) solutions provide the user with an unobtrusive experience when authenticating to application and infrastructure services. **Figure 4-5** provides a logical view of a common SSO architecture in which users authenticate once but are provided access to multiple systems which could be on-premise or in public cloud. This "frictionless authentication" aids user experience but further abstracts the location of information from the user.

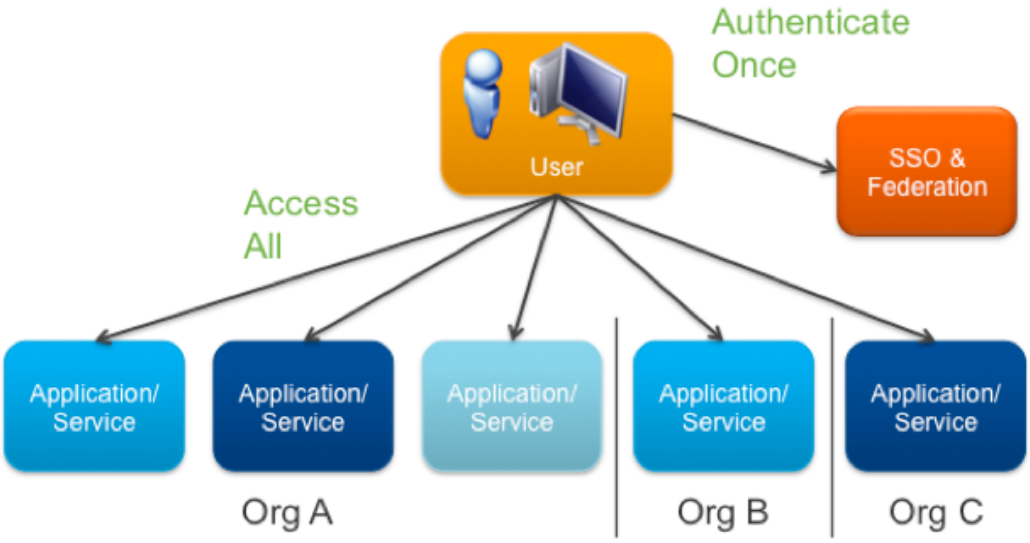

Figure 4-5 Enterprise Federation and SSO [150]

## 4.4. MULTITENANCY: DECOMPOSITION ACROSS SERVICE MODELS

If I am to understand the threats and vulnerabilities associated with multitenancy, it is critically important to understand which multitenancy can be applied across a technology stack.

Multitenancy can be employed at different levels throughout the technology stack and the threats/vulnerabilities need to be considered on a per service model basis. **Figure 4-6** covers the decomposition provided by Kabbedijk et al [151]. This model draws many parallels from the Open Systems Interconnection (OSI) Network Model [152]; as we move up layers from hardware to application instance, the application of multitenancy becomes more granular, less open to portability and extensibility although also more focused on an application instance or the security needs of a use case.

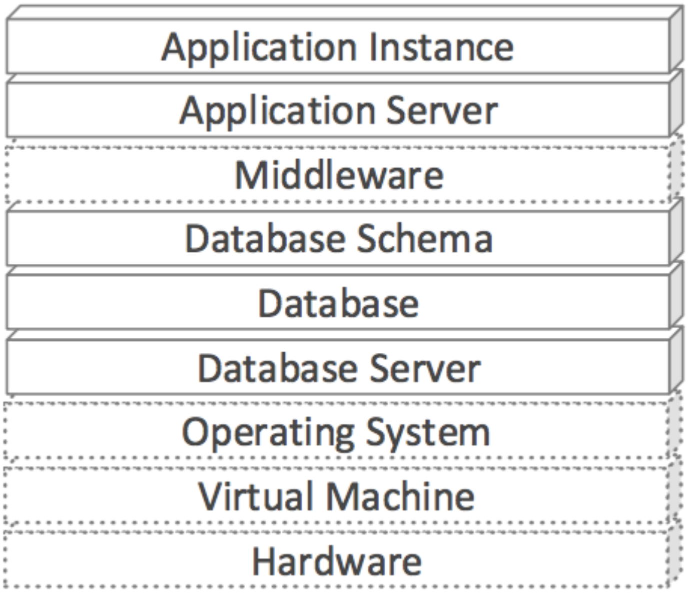

Figure 4-6 Software Stack: Multi-Tenancy Application [151]

The Gartner approach fully acknowledges that the inherent benefits of multitenancy vary dependant on the model adopted. As we move from left to right in **Figure 4-7**, the options for elasticity and operational cost saving increase but so do the risks associated with a breach in the confidentiality, integrity or availability of information.

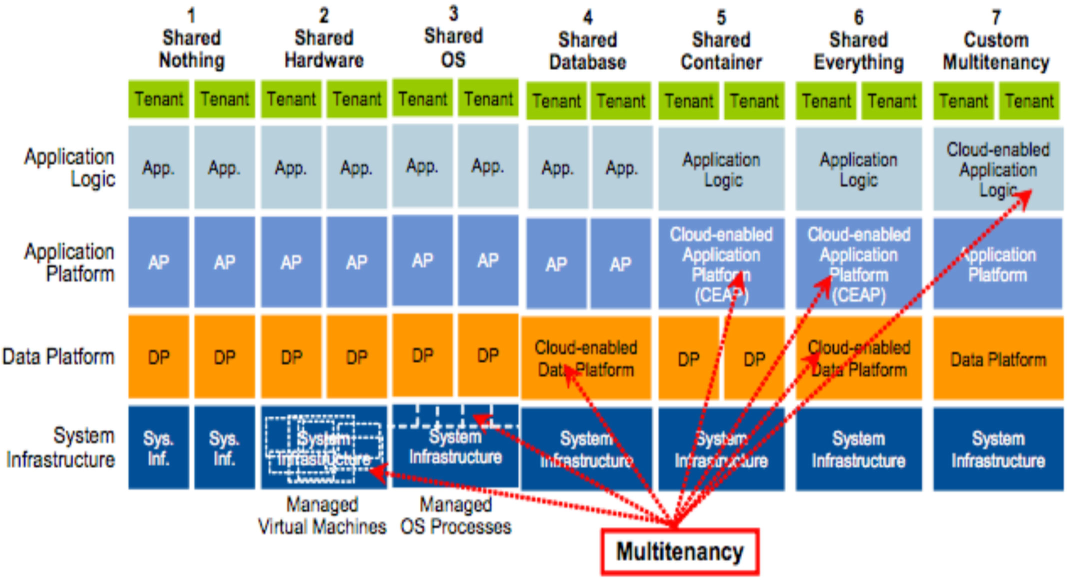

Figure 4-7 Gartner Model of Multitenancy [144]

**Table 4-1** provides an overview of the models outlined by Gartner (**Figure 4-7**). This research is over six years old although the models remain appropriate for consideration. Gartner (correctly) predicted that two multitenancy models will have a profound impact on the future of cloud computing: Shared-Hardware Multitenancy and Shared-Everything Multitenancy. These two models fit broadly into the classifications in **Section 2.6** to describe IaaS and SaaS respectively.

| **Model** | **Description** | **Benefits** | **Considerations** |
|---|---|---|---|
| Shared Nothing | No multitenancy. CSP provisions hardware and software for customer. | Minimised impact and likelihood of a data breach \|True isolation. | Expensive: procurement and maintenance of dedicated infrastructure \| Customer still restricted to standard version of software. |
| Shared Hardware | Shared hardware leveraging a pool of virtualised machines through hypervisor technology. | Dynamic sharing of resources \| Low cost entry to cloud | Potential hypervisor vulnerabilities \| Operational overheads remain |
| Shared OS | Tenant application instances allocated in one OS instance | Flexibility in resource allocation compared to hardware model. | Exploited OS vulnerability impacts all tenants of the platform |

Table 4-1 Gartner Multitenancy Models [144]

| | | | |
|---|---|---|---|
| Shared Database | Tenants share database technology but retain individual application instances. | Extensible: Can be combined with other multitenancy models \| Cost savings on dedicated database management system | Limited flexibility / elasticity |
| Shared Container | One application container instance (application server), separate logical / physical database | Cost savings of presentation and application layer components. | Container must be developed to support resource sharing and scaling \| exploited application layer container could compromise backend database through injection attacks or credential compromise |
| Shared Everything | Application instance and database platform shared across tenants | Significant cost benefits \| Rapid ramp-up / ramp-down of tenants \| | Requirement for a cloud-first development model. All components build specifically for the cloud. |

**Table 4-1 (Cont.) Gartner Multitenancy Models [144]**

The Cloud Security Alliance decomposes multi-tenancy into its constituent components across IaaS, PaaS and SaaS in their "Security Guidance for Critical Areas of Focus in Cloud Computing" [34]. The CSA model corroborates the framework presented by Gartner. **Figure 4-8** details the breakdown of multi-tenancy for each service model. As we move from IaaS through to PaaS and SaaS, the scope of multitenancy grows. We must explore if an intrinsic association of increased risk exists in parallel to this growth of multitenancy scope. In an IaaS deployment, the scope of sharing is contained to the physical security of the datacentre(s), system hardware, abstraction (usually consisting of virtualisation and hypervisor), connectivity and Application Programming Interfaces (APIs). As we move into PaaS, integration and middleware components also become shared across multiple tenants. Finally, in a SaaS model, multitenancy encompasses the entire cloud eco-system. In a SaaS deployment, Application and Presentation (OSI) layer components are also shared. Frequently, data are shared in file stores and databases with logical security controls to protect the confidentiality, integrity and availability of information. An interesting finding from the author is that also the CSA CCM V3.0.1 contains over 130 control specifications with only two[15] which are not universally applicable across IaaS, PaaS and SaaS. This indicates that whilst responsibility for the management and application of controls differs depending on service model, the controls themselves remain standard.

[15] IAM08: Controls relating to the storage and access of identities for an application – not applicable to IaaS | IVS02: Secure storage of virtual machine images – not applicable to SaaS or PaaS

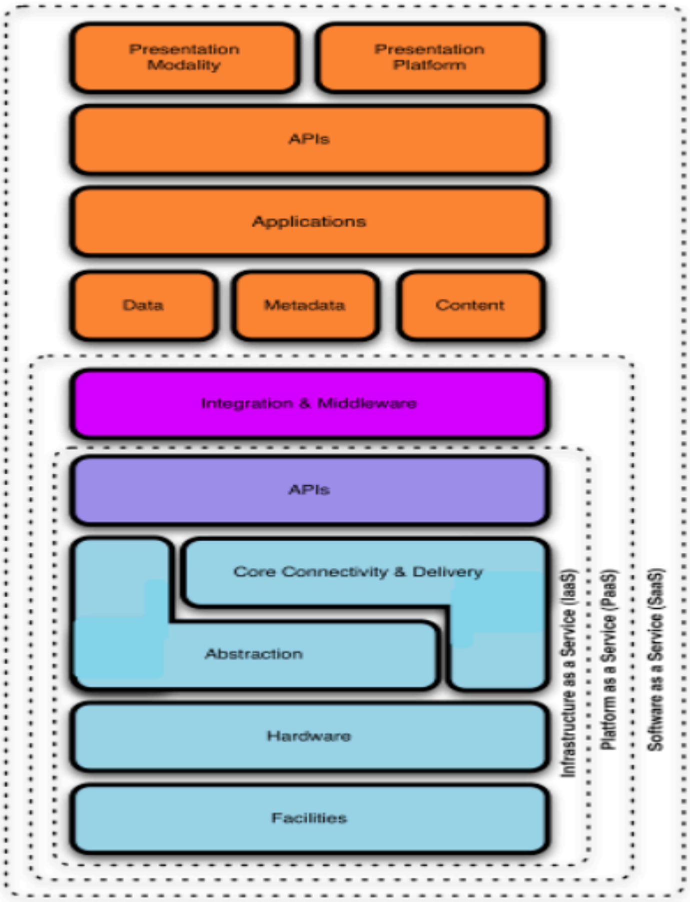

Figure 4-8: CSA Security Guidance for Critical Areas of Focus in Cloud Computing: Multi-Tenancy [34]

Multitenancy, at its core, deals with the sharing of resources across organisations but service model is not the only security consideration when assessing where and how multitenancy can be securely applied. In most cases, technology solutions must be cost effective. In a commercial scenario, the projected income or revenue must exceed the costs associated with building and supporting a solution. Risk management follows a similar line in the consideration of controls and safeguards; if the cost of a control outweighs the expected annualised loss expectancy, the control would be considered a poor investment. This equation assumes that impact can be measured with a financial lens; often it cannot.

**Figure 4-9** [153] depicts three different approaches to the application of database security in an example Software as a Service application. As we see, the approaches outlined offer different levels of logical (and potentially physical) security. Depending on the sensitivity of the information being stored/processed, the risk appetite of the organisation and the budget available, the controls being considered could vary significantly.

Risk is contextual and depending on the classification of information, the controls in each case could be considered suitable for the information residing within the databases. It Is also important to remember that no blanket rule exists in relation to the security posture of a CSP. Actors, threat events and vulnerabilities exist in all cloud service models (IaaS, PaaS and SaaS). It is important to remember that organisations cannot outsource accountability; organisations must perform due diligence on CSPs regardless of adopted service model.

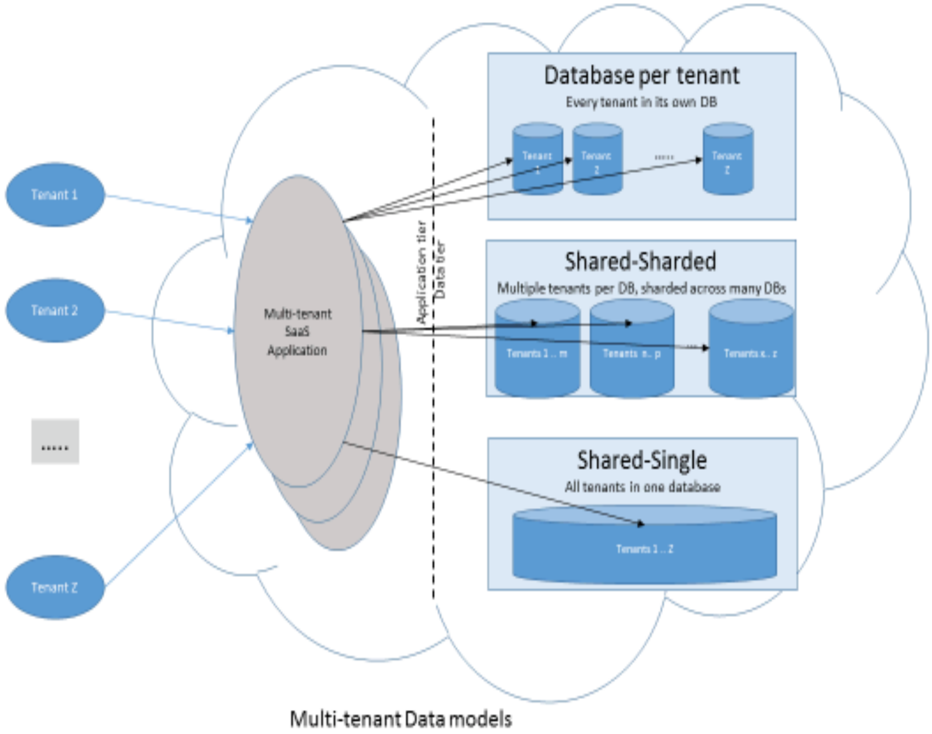

Figure 4-9 Applying Multi-Tenancy in SaaS [153]

With so many variables and permutations in the application of multitenancy, the categorisations of ENISA [35] and NIST [107] are too course grained to make a balanced risk decision regarding the threat events and vulnerabilities associated with multi-tenancy without a thorough understanding of the environment being assessed.

The CSA explicitly reference how responsibilities migrate from CSP to customer as we shift through the stack from SaaS through to IaaS (**Figure 2-7**). Regardless of where responsibility lies, accountability cannot and should not be outsourced. It is the customer who is ultimately responsible for ensuring that the appropriate people, process and technology controls exist commensurate with the information being stored and processed. In an outsourced scenario, how these controls are evidenced will change.

## 4.5. VULNERABILITIES ASSOCIATED WITH MULTITENANCY

Based on the research performed in **Section 4.4**, I conclude that public cloud computing introduces vulnerabilities irrespective of service model. Where academic and industry papers cover the "risks" of multitenancy, they generally focus on the technical vulnerabilities related to the hypervisor [154, 18].

Vulnerabilities within hypervisors are on an upward trend [155] and this shows no sign of slowing [156]. As customers demand increased functionality, hypervisor providers are expanding their code bases to cater for this increase in capability. As the code footprint for a hypervisor increases, so do the avenues for exploitation [155]. It is true to say that hypervisors have significantly less Lines of Code (LoC) than a conventional operating system (OS) but vulnerabilities are discovered frequently. Researchers and hackers have found a way to monetise vulnerability discovery through bug bounty programs and security conferences which financially remunerate researchers who find vulnerabilities in hardware and software [157].

I have already identified resource sharing / lack of isolation as an avenue for exploitation across public cloud. Whilst the vulnerability exists, the impact and likelihood of compromising the confidentiality, integrity or availability of data through the exploitation of resource isolation varies and should be measured using a framework such as IRAM2 [106] or NIST 800-30.

I was surprised that vulnerabilities associated with forensic analysis and data destruction are not (by ENISA's classification) included in a risk associated with resource isolation. "*Insecure or Incomplete Data Deletion*" is explicitly referenced [35, p. 10] with a vulnerability of "*Sensitive Media Sanitisation*". I agree that sanitisation is the process which needs to be followed but it is the fact that data is shared between tenants which is ultimately the reason why ineffective data deletion occurs in a public cloud environment.

The ENISA model does not identify forensics as a top cloud risk. After researching the subject, I was unsure if forensic analysis of an environment is made more difficult through the adoption of a multitenant, public cloud. It is fair to say that having multiple organisations residing on physical hardware will make taking a disk offline almost impossible, public cloud services introduce the options for real-time copying of disk images which could improve operational response times and malware infection remediation [35, p. 58]. It is commonly accepted that public cloud computing does make legally admissible forensic disk evidence harder to gather [158].

ENISA [35] further cites the loss or compromise of logs as a forensics risk. Mature cloud service platforms [159, 160] provide native log streaming capabilities to ensure that operational and security logs can be exported, batch or real time, into a Security Information and Events Management (SIEM) system, Big Data or Business Intelligence solution; this is another situation where inherent impact is being considered rather than performing a qualitative risk equation.

As this thesis is a time-bound study, it was impractical to include every vulnerability study which was identified as "in-scope" for public cloud although several areas of academic research are noteworthy. Brown et al. [145] research the risks specifically associated with multitenancy. Their paper is structured in such a way as to identify not only inherent risk but also appropriate countermeasures. A method for assessing the strength, breath or maturity of controls is not discussed nor are vulnerabilities presented which have not been covered in our review. The Journal of Applied Sciences [161] presents a thorough approach to cloud threat modelling which assesses the efficacy of several prominent threat model methodologies. The paper takes a similar approach to this work in their decomposition of threats and vulnerabilities although this author feels that many of these (threats and vulnerabilities) would fall into the "exacerbated by cloud" category as opposed to being the intrinsic result of a cloud deployment. This is however an excellent resource and should be reviewed by anyone looking at adopting a qualitative risk assessment for a multitenant environment.

The Open Web Application Security Project (OWASP) [110] presents 7 risks which they suggest are directly related to multitenancy in a cloud ecosystem. If we are to apply consistent terminology, these risks are vulnerabilities and will be considered thus for our research purposes. Of the 7 risks, six are identified as appropriate considerations for IaaS, PaaS and SaaS, this supports my view that inherent multitenancy vulnerabilities exist irrespective of service model. The impact and likelihood of exploiting a vulnerability depends on service model and the application of security controls. In **Figure 4-10** we combine ENISA cloud vulnerabilities with OWASP's risks of multitenancy to highlight the volume of vulnerabilities which are intrinsically-linked to resource isolation.

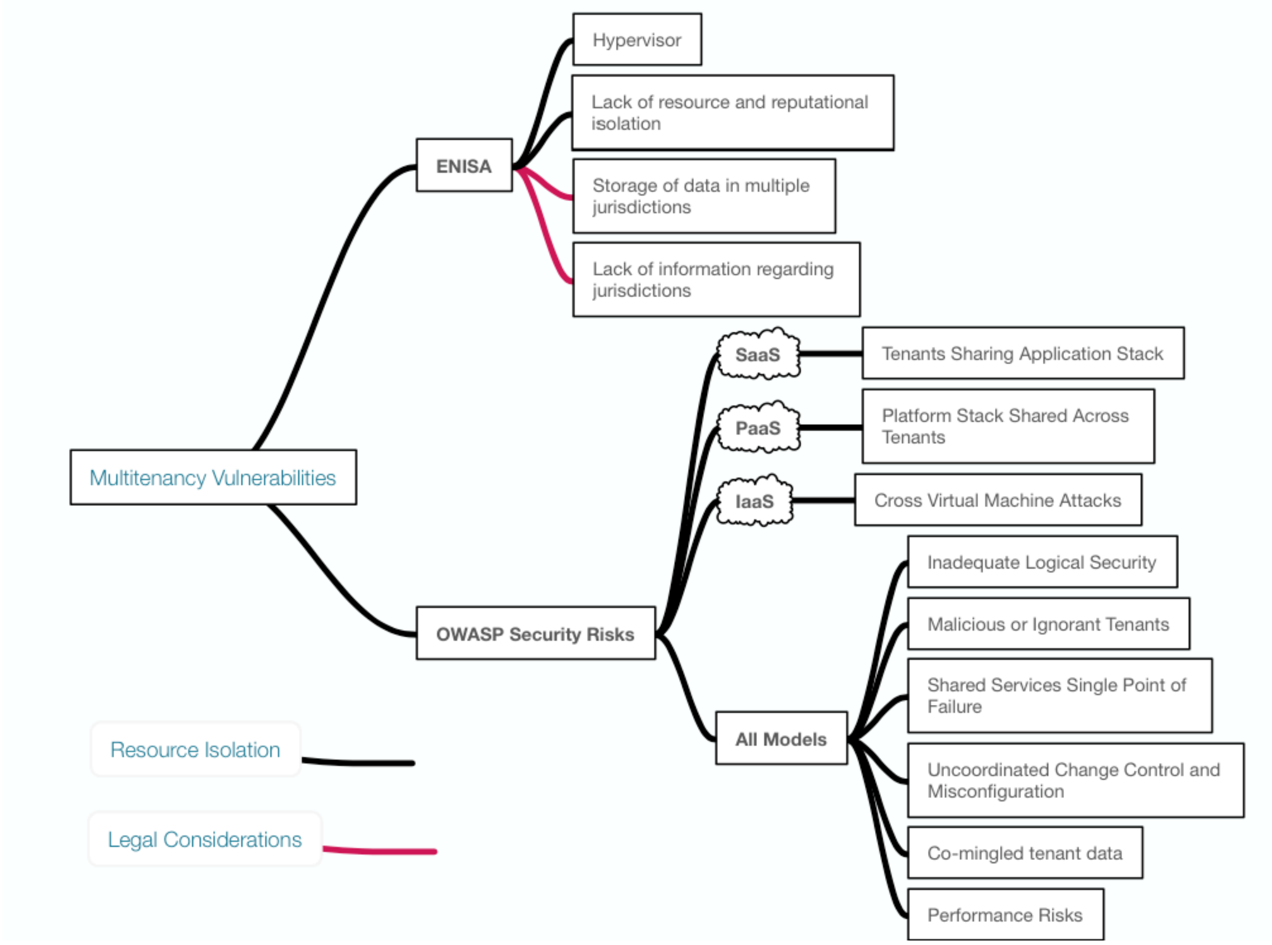
Multitenancy Vulnerabilities
ENISA
Hypervisor
Lack of resource and reputational isolation
Storage of data in multiple jurisdictions
Lack of information regarding jurisdictions
OWASP Security Risks
SaaS
Tenants Sharing Application Stack
PaaS
Platform Stack Shared Across Tenants
IaaS
Cross Virtual Machine Attacks
All Models
Inadequate Logical Security
Malicious or Ignorant Tenants
Shared Services Single Point of Failure
Uncoordinated Change Control and Misconfiguration
Co-mingled tenant data
Performance Risks
Resource Isolation
Legal Considerations

**Figure 4-10 Multitenancy Vulnerabilities**

When combined, ENISA's vulnerabilities and OWASP's risks appropriately and comprehensively detail the possible vulnerabilities associated with resource sharing and therefore multitenancy. As covered in **Chapter 3**, controls and countermeasure are used to lessen or remove a vulnerability or to deter / prevent a threat event from occurring. OWASP provides seven countermeasures which it deems appropriate for multitenancy [110]. These are solid architectural principles although the skills of an SME would be required to translate these high-level policy statements into an actionable set of controls; without such a stakeholder, the abstracted terms in a policy statement can be easily misconstrued and we introduce a similar situation which occurs with qualitative risk analysis – that being the subjectivity of terms like "secure" or "inadequate". "Architecting for Multi-multitenancy" is a sound architectural principle although by way of guidance for mitigating the risks associated with cloud computing, it is not prescriptive enough to add material value across a risk assessment process.

These countermeasures (controls) I outline here are a combination of people, process and technology. A triad common to information security for generations. It is important to remember that vulnerabilities should only be a consideration for an organisation if a threat actor can initiate a threat event which can exploit the vulnerability and the organisation is unable to apply security controls to lessen the impact and / or likelihood to a palatable level to accept the residual risk.

In the following sections, we will further explore the vulnerabilities which are applicable to public cloud. As evidenced through my research in **Appendix A: ENISA Cloud Vulnerability Assessment**, most public cloud vulnerabilities are related to people and or process. This assertion is applicable whether the vulnerability is originating in cloud, exacerbated by cloud or a general vulnerability. Public cloud and the concept of resource sharing can exacerbate the impact of vulnerability exploitation; compromising the CIA of information in your tenant, could have ramifications for other tenants in that environment. These vulnerabilities exist with on-premise solutions although public cloud does present a potentially higher impact We will continue in this chapter to highlight the likelihood of cloud-originating events allowing an organisation to make a balanced risk decision on the use of public cloud.

Gartner (Procter et al.) provides validation of my assertion through their analysis of the risks and benefits of public cloud [111]. They state that by 2019, cloud security audits will be based on a model which delineates risk (in our model "unique to cloud") and implementation detail. This compartmentalisation being necessary as currently organisations are failing to separate native and inherent vulnerabilities of public cloud and those introduced through poor operational or implementation detail.

## 4.6. HYPERVISORS AND VIRTUALISATION: THE FOUNDATION OF MULTITENANCY

Multitenancy forms the basis of public cloud computing and is almost always based on a foundation of virtualisation [151]. Multitenancy and virtualisation are often, incorrectly, used interchangeably. Multitenancy is the sharing of resources, at many layers, across a hardware and software stack. Virtualisation is the predominant means of delivering multitenancy.

In **Section 2.5**, essential characteristics of cloud computing were presented. NIST [32] presents five characteristics which do not include virtualisation and multitenancy. These omissions are interesting and do not support the views of the CSA [34]. As we will explore in the following sections, it is evident that the technological vulnerabilities exclusively reserved for public cloud computing all focus on the combination of virtualisation and hypervisor technology.

In **Section 4.3**, I reached the conclusion that Virtual Machines (VMs) are not the exclusive reserve of cloud computing; VMs are commonplace in all computing environments where cost efficiencies, scalability and ease of administration are of paramount concern. Virtualisation (and therefore virtual machines) are generally implemented using hypervisor technology. Hypervisors provides a level of abstraction between physical hardware and an operating system and provides the scalability to support multiple tenants on the same hardware. Hypervisors mediate and support the sharing of physical resources between VMs.

To use a biological example, the hypervisor is the brain of the virtualisation eco-system; it controls the allocation of resources for all virtual machines. The hypervisor acts as a mediation layer for all (virtualised) OS instances on that physical machine. All network traffic should route ingress / egress of the hypervisor, although exceptions to this rule do exist (Zhang [18]). As the brain of the system, the hypervisor needs to retain a privilege level capable of performing sensitive operations on behalf of guest operating systems. An exploration of protection mechanisms is warranted to explain this concept further.

Protection Rings are a mechanism to protect data and functionality from faults (fault tolerance) and malicious behaviour (computer security) [162]. X86 processor architecture [163] allows for four privilege levels (0-3). Zero being the most privileged and three the least privileged (**Figure 4-11**). Only two levels are used in a non-virtualised environment: Ring 0 is reserved for operating system functionality and user applications are run in Ring 3. System calls and user-model calls being carried out in their respective zones. As hardware manufacturers have added virtualisation support (Intel VT [164], AMD-V [165]) to processors, a new protection ring has been added to some code instruction sets. Elisan et al [166] call this layer a "ring-1" and it allows the hypervisor to monitor the guest operating systems running in ring 0 but not true 'ring 0' (in that they're using virtualised hardware as opposed to physical hardware). This layer is known as a 'thin hypervisor' [167] It should be noted that for this 'ring -1' to exist, the processor (of the physical machine) must support virtualisation extensions (Intel VT, AMD-V). Thin hypervisors form the basis for rogue hypervisor attacks which will be discussed in the following subsections. Microsoft Hyper-V [168] is commonly referred to as a type 2 hypervisor although this is not accurate. Upon installation of the Hyper-V server role (within Windows), a process is followed similar in construction, to that of a rogue hypervisor attack which places the Host OS into a VM and installs the Hypervisor below the OS (**Figure 4-11**). In a Xen Hypervisor architecture (primary focus of this study), the Host OS / Dom0 resides in Ring 1. In all situations, user applications remain in ring 3.

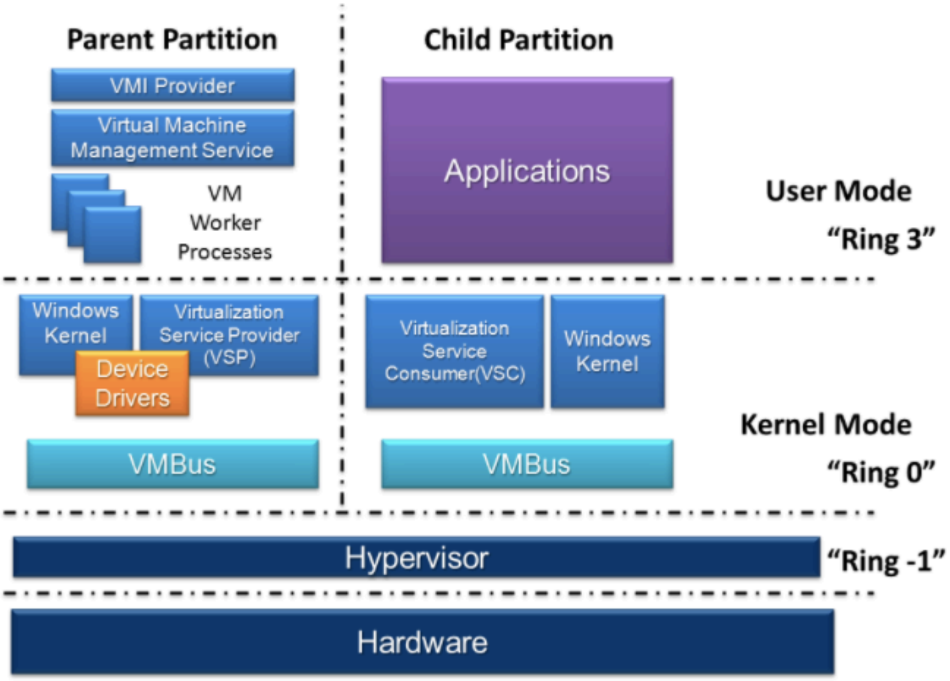

Figure 4-11 Protection Ring Architecture with Processor Virtualisation [169]

**Figure 4-12** [170] depicts the components and interactions of a type 1 hypervisor. The dotted line delineates a 'trusted ecosystem' or rings 0 and 1. Six interactions are identified as significant and an understanding of these communications is critically important if we are to identify and contextualise threat events associated with multitenancy. Virtualisation introduces new forms of interaction between the operating system and physical hardware. This exchange of information and resource management must be clearly explained if the threats and vulnerabilities to cloud computing are to be assessed.

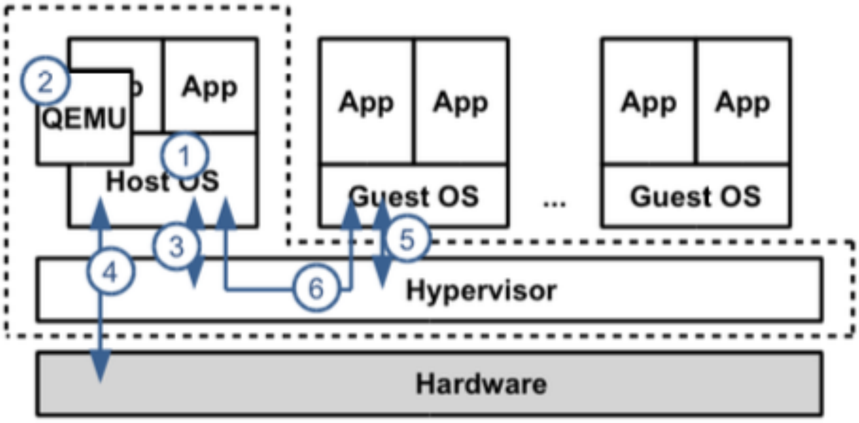

Figure 4-12 Hypervisor Interactions [170]

| Number | Name | Description |
|---|---|---|
| 1 | Host OS | "Dom0" in Xen vernacular. A privileged domain responsible for managing unprivileged domains [22]. Dom0 runs the XEN Management Toolset. |
| 2 | Emulator | QEMU (Quick Emulator) [171] performs hardware virtualisation. QEMU emulates multiple components (Processor, Memory, Network, Input / Output (I/O)) |
| 3 | Hypercalls | Communication between Dom0 and Hypervisor through Hypercalls – examples being system admin using Dom0 for guest OS configuration. |
| 4 | Device Drivers | The Host OS may include device drivers for direct communication with machine hardware. |
| 5 | Guest OS <> Hypervisor | Guest OS communicates directly with Hypervisor via hypercalls when it has an explicit need to request a service. |
| 6 | Guest OS <> Host OS / Emulator | Communication via VM Exits. |

Table 4-2 Hypervisor to VM Communication

Steps 5 and 6 (above) are the primary vectors for the VM Escape (5) and VM-to-VM attacks (6) we will discuss in the following sections. A VM Exit is an event which occurs when a VM's code is interrupted and the hypervisor code begins to execute to handle some event (e.g., emulate memory access, deliver a virtual timer interrupt). A hypercall is like a system call and is used by the guest VM to request explicit services from the hypervisor [170]. Hypercalls are used in situations only where paravirtualisation[16] is leveraged [155] .

There are frequently times with a currently running VM needs to pass control to the hypervisor for the running of system privileges. When this is necessary, a VM Exit is invoked. VM Exists occur for a plethora of reasons. On a 64bit x86 processor running virtualisation extensions, there are 56 reasons for a VM Exit and these form the biggest threat to the security of a virtualised environment [170]. VM Exits are called when the guest OS attempts to access memory at a physical address, execute an I/O instruction or read / write to specific registers.

Hypervisor functionality can be broken down into 11 functional areas [155]:

[16] Paravirtualisation is method of providing I/O virtualisation in XEN. It reduces both cost and complexity although the guest OS must be modified to host a "front end driver". This driver syncs with a back end driver which resides on the hypervisor and has direct access to hardware.

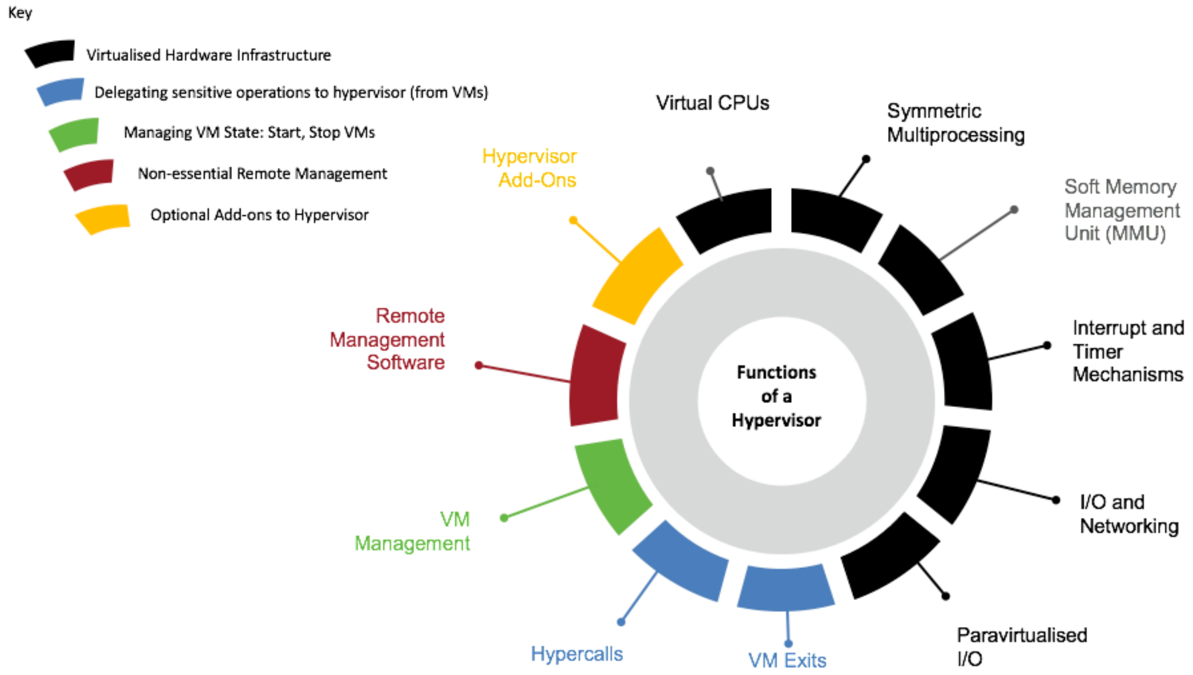
Key
Virtualised Hardware Infrastructure
Delegating sensitive operations to hypervisor (from VMs)
Managing VM State: Start, Stop VMs
Non-essential Remote Management
Optional Add-ons to Hypervisor
Functions of a Hypervisor
Virtual CPUs
Symmetric Multiprocessing
Soft Memory Management Unit (MMU)
Interrupt and Timer Mechanisms
I/O and Networking
Paravirtualised I/O
VM Exits
Hypercalls
VM Management
Remote Management Software
Hypervisor Add-Ons

Figure 4-13 Hypervisor Functionality

There are two commonly accepted forms of hypervisor:

**Type 1 Hypervisor:**

An abstraction layer sitting directly on hardware and mediating access to virtual machine images. A Type 1 hypervisor is commonly referred to as a "bare metal" hypervisor as it resides directly between hardware and virtual machines. Amazon Web Services (AWS) [20] is based on a customised Xen software stack. AWS *"has the largest share of compute capacity in use by paying customers — many times the aggregate size of all other providers in the market"* [21]; for this reason, we will be using the XEN architecture (**Figure 4-14**) in our discussion of Type 1 hypervisor vulnerabilities.

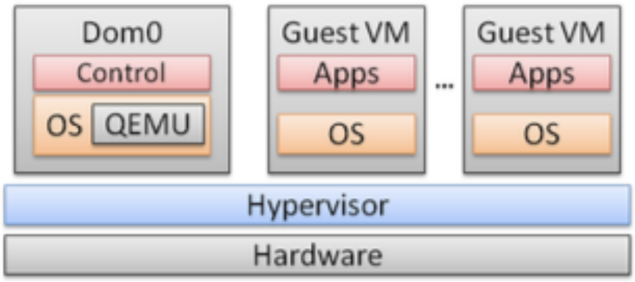

Figure 4-14 Xen Architecture [155]

In a Xen architecture, Ring 1 functionality is provided through the "Dom0" which takes the role of Host OS. Dom0 facilitates the management functions. Dom0 is constructed of a Linux-based kernel. Having Dom0 carry out VM management functions calls the guest VMs to run without VM customisations (HVM - fully virtualised mode [155]). Dom0 also exposes emulated hardware devices through a single instance of QEMU [171] which allows for near native guest VM system performance although this emulation layer has been known to present vulnerabilities to the virtualisation ecosystem [172]

Type 1 hypervisors are used across popular public IaaS technology solutions offered by Amazon [20] and Rackspace [173].

**Type 2 Hypervisor:**

A precis of type 2 hypervisors has been included for completeness. Type 2 hypervisors consist of a software capability installed within an OS instance thus requiring an additional layer and subsequently attack surface. Type 2 hypervisors are generally suited to workstation emulation and testing; they are not generally used for production environments [19]; a requirement of a type 2 hypervisor is the installation of a host OS to house the hypervisor. This approach introduces both performance and security considerations for the enterprise. Interestingly, across my research I discovered that most reported vulnerabilities associated with hypervisors and resource isolation related to type 2 hypervisor environments [174, 175].

VMWare Workstation [176] and Oracle Virtualbox [177] are examples of type 2 hypervisors.

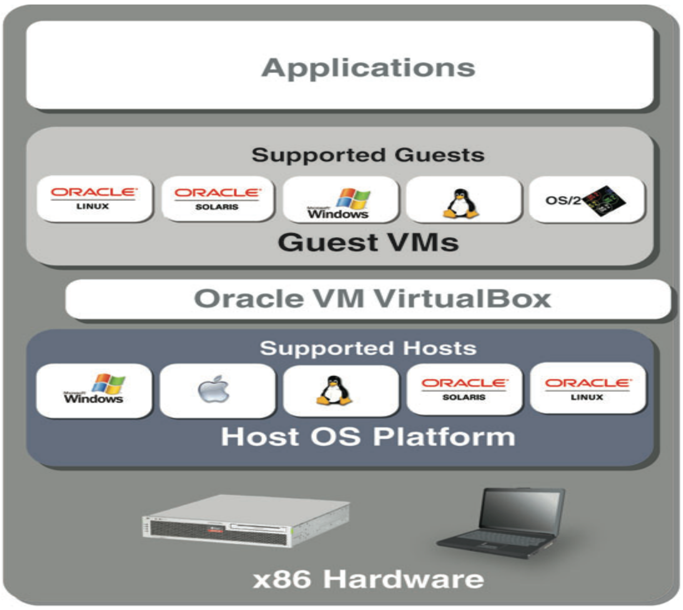

Figure 4-15 Oracle VirtualBox Architecture [178]

## 4.7. THREAT ACTORS

In my introduction to this chapter, I raised two questions relating to potential vulnerabilities of multitenancy:

***"Are the vulnerabilities associated with resource sharing exploitable with a reasonable degree of likelihood by range of threat actors with varying levels of skill and persistence?"***

***"Are the vulnerabilities of multitenancy appropriately contextualised and do other attack paths exist which are more likely exploitable by all / any of the threat actors used in this study?"***

To answer these questions, I need to analyse the threat actors which are appropriate in a discussion regarding the exploitation of multitenancy. In **Section 3.5**, explicit reference was made to three threat event / actor categories:

- Adversarial
- Accidental
- Environmental

The (multitenancy / resource isolation) vulnerabilities and threat events outlined in this paper focus significantly on adversarial threat events. Our study of vulnerabilities associated with resource isolation will take the view of an attacker for the following reasons:

1. There are adversarial threat events which are (at least ostensibly) “cloud-enabled”.
2. Most academic and industry literature reviewed by the author [154, 179, 18] focus on adversarial events.
3. Adversarial events allow us to compare other malicious threat events to contrast the seriousness of resource isolation attacks compared with other contemporary threat events.

### 4.7.1. MALICIOUS ACTOR CATEGORIES AND ATTRIBUTES

In the same way that “exploiting multitenancy” is too general a term in assessing impact and likelihood of a threat event, “malicious actor” is generally categorising someone or something as having nefarious intentions. Intention and motivation are attributes to consider but of limited value without skills and opportunity. Examples of this generalisation exist ubiquitously within the IT space. For example, if we suggest that high impact breaches are exclusively the responsibility of a state-sponsored aggressor, there is a feeling of placation: if all aggressors have unlimited funds, perseverance and skill, of course they are going to succeed. Whilst such adversaries exist, many data breaches are the result of poor security hygiene and the exploitation of vulnerabilities by lesser-skilled, opportunistic criminals. It is important we understand the threat landscape and the actors we are protecting our organisations from; when we talk about the risks of public cloud, it is important to ask: “risks from whom?” as much as it’s important to understand the organisational appetite for risk. If your actors are not motivated or capable, there is no risk.

The ISF [106, p. 27] has addressed the broad nature of threat actor / event categories through the introduction of threat attributes. These will form a key component of our assessment criteria for impact and likelihood of an exploit relating to resource isolation.

We pay special attention to the attributes associated with an adversarial / malicious actor:

- **Capability**: How proficient is the threat and how well resourced is it?
- **Commitment**: The resources the threat willing to expend (time, money)?
- **History**: Is there a history of this type of threat carrying out this type of attack?
- **Motivation**: To what extent is the actor motivated?

From these attributes, we derive two important risk factors [106, p. 29] (**Figure 4-16**):

**Likelihood of initiation** (LoI): The likelihood that a threat will initiate one or more threat events against the environment being assessed – assessed through **“**history” and “motivation”**.**

**Threat Strength** (TS): TS is critically important as all malicious actors should not be treated equally. Threat strength is established by assessing an actor / event’s “capability” and “commitment”.

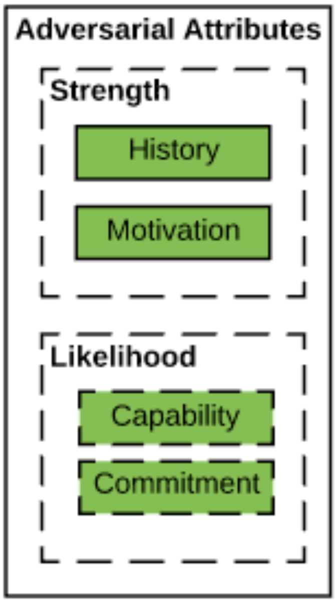

Figure 4-16 Adversarial Threat Attributes

Both LoI and TS are measured qualitatively across four ratings: Negligible, Low, Moderate or High (0-3). We take the aggregated score and this gives us our risk factor rating (**Figure 4-17**).

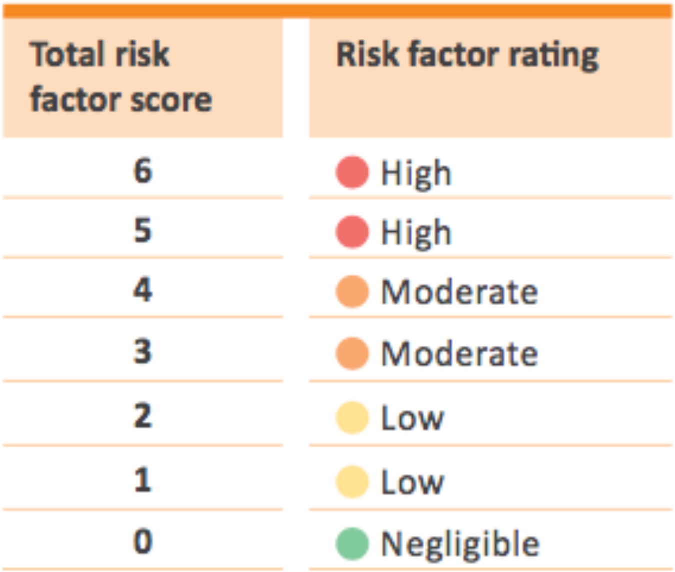

| Total risk factor score | Risk factor rating |
|---|---|
| 6 | 🔴 High |
| 5 | 🔴 High |
| 4 | 🟠 Moderate |
| 3 | 🟠 Moderate |
| 2 | 🟡 Low |
| 1 | 🟡 Low |
| 0 | 🟢 Negligible |

Figure 4-17 IRAM2 Threat Risk Factors [106]

In my research and analysis of resource isolation vulnerabilities, it is evident that likelihood of a successful attack depends on the attributes of the threat actor. In the following subsections, I will present the threat events, vulnerabilities and security controls associated with public cloud resource sharing. To contextualise these threats, I believe that it is necessary to select prominent threat actor categories and assess each threat in the context of the actors. For my study, I have decided to select three actor categories which I have documented in **Table 4-3**. This table includes a qualitative analysis of each actor's capability and commitment used to aggregate a "threat strength":

| Threat Actor | Description | Capability | Commitment | Threat Strength |
|---|---|---|---|---|
| Opportunistic | Script Kiddie, lone-wolf. | 1 | 2 | Low |
| Organised Criminal | Financially-motivated, skilled cyber adversary | 2 | 2 | Moderate |
| Nation State | Government supported and financed | 3 | 3 | High |

**Table 4-3 Threat Actor Capabilities**

Threat strength has been assessed using the IRAM2 model (capability + commitment). Threat strength will remain consistent across each type of attack scenario discussed below. For each case, I will consider individual "history" and "motivation" attributes. I perform this activity to allow organisations to make an informed decision regarding the threat actors they are willing to concede are capability of extracting information or infiltrating a network – this varies from industry to industry. Whilst a government agency may consider the need for protection from nation state actors, this may not be appropriate for a small to medium organisation; they may decide that the likelihood and / or impact of such an assailant compromising their network is too great a cost to mitigate. They therefore make a conscious decision to "accept" the risk (**Section 3.2**).

A process of qualitative risk analysis will be undertaken for each threat event by applying an attacker-centric view of exploitation. By decomposing threat actors, we can better assess risk. A common question from board executives in relation to a data breach or a cyber-attack is "could this happen to us?" Whilst several variables come into play, the security professional needs to consider the attributes of the assailant. I apply categorisations through the ISF's concept of "threat strength" [106].

Selecting which threat actors to include in my analysis was a difficult process. My selections of opportunistic, organised financial criminal and nation state were based on a desire to provide coverage across the broadest range of attackers possible. The rationale being that this would best contextualise findings and evidence that different actor types require different security controls. The category of "hacktivism" has been intentionally omitted from our classification. The motivations of a hacktivism focus on overt actions to further an ideological or political cause. Compromising the CIA of an enterprise through the exploitation of resource isolation vulnerabilities can be associated with the overt requirement of a successful hacktivism campaign although generally defacement of digital property and denial of service better fit the hacktivism modus operandi. As an MSc thesis is a time and scope-bound activity, threat actor categorisation has not been covered in detail. I would recommend that anyone wishing to obtain a broader understanding of threat actor categories should refer to the SANS Institute (Irwin) paper entitled "*Creating a Threat Profile for your Organisation*" [180].

I consider history and motivation of the actor to contextualise likelihood although this should be considered a measure of intent or an "inherent likelihood"; simply because an actor is motivated and there is a history of this event, this does not mean the threat event will be successful. I will subsequently assess the likelihood of success; this measurement takes the threat strength of an attacker initiating an event but also considers the maturity and coverage of selected controls. This is a pragmatic process which provides true visibility into feasibility of exploitation. A scoring table is included below which includes ratings from the IRAM2 paper:

| Control Strength | Threat Strength | Negligible | Low | Moderate | High |
|---|---|---|---|---|---|
| High | | Negligible | Negligible | Low | Moderate |
| Moderate | | Negligible | Low | Moderate | High |
| Low | | Low | Low | High | High |
| Negligible | | Low | Moderate | High | High |

Table 4-4 Likelihood of Success Formula

The final equation which needs to be performed is the residual likelihood. Through obtaining residual likelihood, the organisation can make a balanced risk decision. We understand the threat actor, the threat event but importantly, the efficacy of control(s).

| Likelihood of Initiation | Likelihood of Success | Negligible | Low | Moderate | High |
|---|---|---|---|---|---|
| High | | Moderate | Moderate | High | High |
| Moderate | | Low | Moderate | Moderate | High |
| Low | | Low | Low | Low | Moderate |
| Negligible | | Negligible | Low | Low | Low |

Table 4-5 Residual Likelihood Formula

The ISF IRAM2 documentation [106] demonstrates an end to end view of the risk assessment process and provides the reader with an understand of the constituent components (**Figure 4-18**). I include this diagrammatical view as the assessment criteria and metrics introduced above can appear daunting to anyone new to the IRAM2 process. **Figure 4-18** evidences not only that threat actors initiate threat events but also that the likelihood of success associated with an attempted data breach / cyber-attack takes not only the strength of the attack but also the strength of the controls within the victim environment. I assert that if a similar process had been applied to some of the seminal work around "cloud risk" [154, 18, 179], a more contextualised view of likelihood would have been concluded.

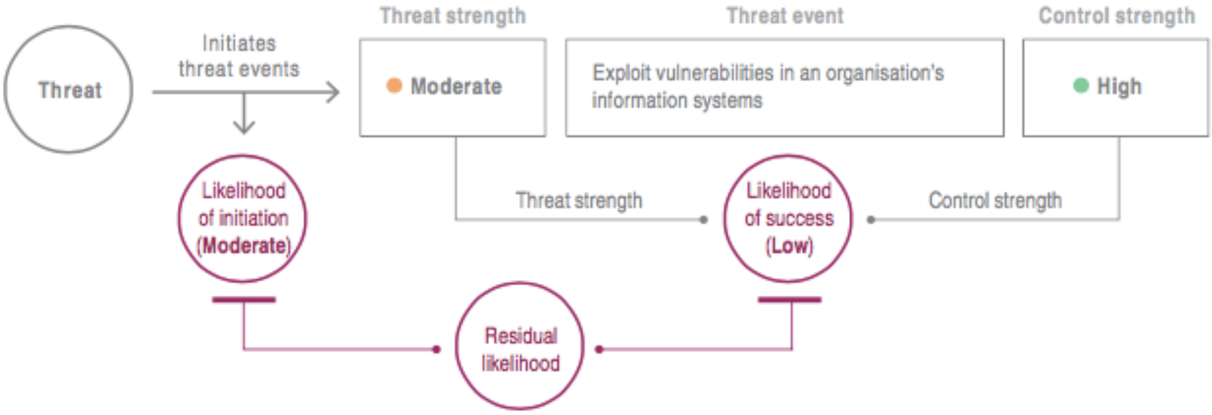

**Figure 4-18 IRAM2: Assessing Residual Likelihood [106]**

For each attack step outlined in the following section (and in **Figure 4-19**), I will include three tables. These contain the results of my qualitative analysis and will cover:

| **Table Name** | **Description** |
|---|---|
| Likelihood of Initiation | Assessing capability and commitment of the threat actor for a specific threat event |
| Likelihood of Success | Balancing the history and motivation of the actor with available strength and coverage of controls |
| Residual Likelihood | Combining Likelihood of Success with Likelihood of Initiation |

**Table 4-6 Categories of Table for Threat Analysis**

## 4.8. ATTACK CATEGORIES AND PREREQUISITES: RESOURCE ISOLATION

In **Section 4.7.1**, I reviewed the "who" in our attack scenario. The adversaries capable of mounting an attack on our cloud infrastructure. In this section, I am defining the "how". We understand the actors in our risk equation but how do they initiate an attack on our infrastructure?

Information security conferences and journals regularly contain features on the threat events that exploit vulnerabilities in a virtualisation stack which has afforded an attacker the ability to steal information and / or alter information in another tenant's space [181, 179]. Considerations must be made for the circumstances in which exploits are presented – several hypervisor and virtual machine attacks have been performed under proof-of-concept conditions without instances being publicly documented in the wild [154, 179].

Many academic papers [154, 18, 170] have been written extensively and comprehensively about threat events which have exploited vulnerabilities in virtualisation architecture. These threat events are real and require consideration. I have identified five core threats vectors which can be attributed to resource isolation and which would exploit the vulnerabilities outlined by OWASP [110] and ENISA [35]. In the interests of time and scope, I have decided against conducted my own threat modelling activity. The rationale being that many before me

[18, 166, 154] have defined the threat events applicable in this context. Whilst variations and subtleties exist, four prominent threats stand out for analysis in this paper:

- VM Escape
- Rogue Hypervisor
- Inter-VM Attacks
- Denial of Service / Resource Exhaustion

The following sections will explore each of the above threat events. The objective of this activity is to provide input into a qualitative risk equation and to understand the vulnerabilities unique to public cloud and the methods with which we can exploit them.

It is imperative to not only understand that a vulnerability exists and that a threat actor is incentivised to pursue a target but also to qualify what steps an attacker needs to go through to exploit a vulnerability. We will achieve this through the adoption of the ISF IRAM2 model and the inclusion of "Likelihood of Initiation". As discussed in **Section 4.7.1** threat strength can be static in each case as we are dealing with a closed set of consistent actors. I do however acknowledge that as exploits are developed by a skilled attacker, these are often transferred into repeatable packages that can be later leveraged by less-skilled assailants. In the context of this study, I am applying a static "threat strength" for each assailant and I am assuming that the "capability" and "commitment" of the attack remains consistent. I do appreciate that as malware is commoditised and vulnerability information is made available, an assailant's "capability" increases.

It is important that any risk management activity considers the strength and coverage of security controls. Simply because an attack is likely to take place, this does not mean that it will be successful. The focus of my study, for reasons of coverage and open access to the Xen hypervisor, is AWS. The security controls available to users of AWS provide organisations with access to carrier grade security capabilities which many would not be able to afford if they were sourcing and running them locally; at a high level, these are covered in **Section 2.3.2**. The service-based model of cloud not just being attractive for infrastructure and applications but also the improvement of security controls for organisations.

The sensationalism of cyber-attacks is common place in both technology and mainstream media; the laborious steps necessary to exploit a vulnerability can be overlooked in favour of a good headline. I have prepared an end-to-end attack tree with the goal of highlighting the complexity of hypervisor compromise and the steps that we will decompose in the following subsections (**Figure 4-19**):

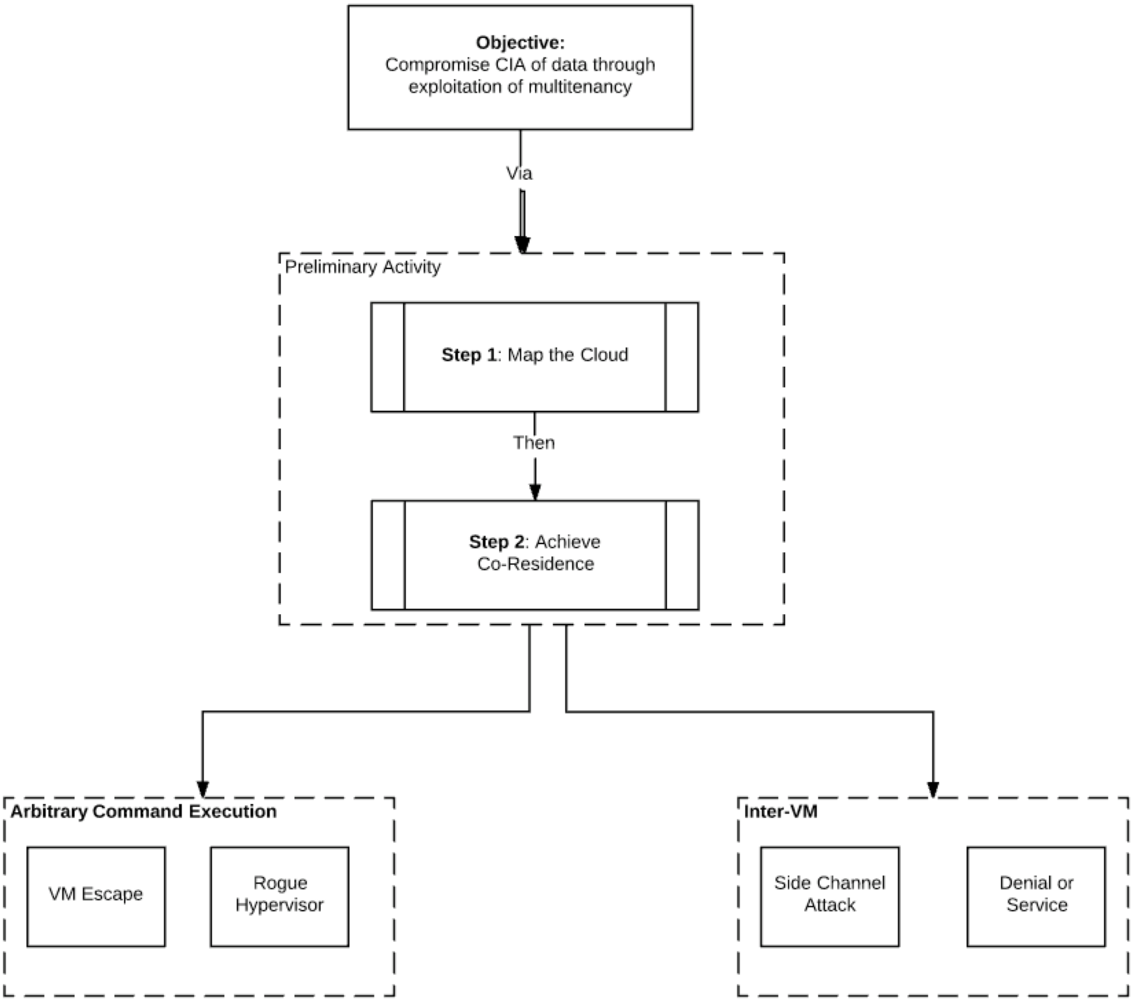
**Objective:**
Compromise CIA of data through exploitation of multitenancy
Via
Preliminary Activity
**Step 1**: Map the Cloud
Then
**Step 2**: Achieve Co-Residence
**Arbitrary Command Execution**
VM Escape
Rogue Hypervisor
**Inter-VM**
Side Channel Attack
Denial or Service

**Figure 4-19 Multitenancy Attack Tree**

I wish to prove a hypothesis that the threat events considered "in-scope" for public cloud environments require the completion of several preliminary tasks which in themselves would deter all but the most persistent, patient and well-funded of threat actor. There are two foot printing activities which are sometimes often overlooked in an analysis of resource isolation vulnerabilities.

### 4.8.1. STEP 1: MAPPING THE CLOUD

The first step in conducting an attack is to map the cloud environment. This aligns to the "Information Gathering" and "Network Mapping" phases of the Information Systems Security Assessment Framework (ISSAF) [182]. Public CSPs, by nature of being public, need to make their Internet Protocol (IP) address ranges readily available to customers.

Ristenpart et al [154] have researched the processes associated with foot printing cloud infrastructure based on Amazon's EC2 public cloud. Their findings identify a mechanism to the discover information regarding public Amazon AWS services and their study provides empirical evidence of co-residence based only on network probing / mapping. Their investigation identifies a strong association between a VM's IP address and other important creation parameters: size of VM, location, compute power, etc. A conclusion is reached that different Amazon Availability Zones leverage separate physical infrastructure [154]. Based on the limited sample sets of the study, there appears to be an association between IP subnets and resource types; again, understandable for reasons of management and operational efficiency.

At first-glance, the simplicity of network mapping could cause concern; however, this attack phase is predicated on the ability for an attacker to receive enumeration information from the hosts via tools such as NMAP [183] and Hping [184]. This is where control strength is an important consideration. It is a simple step to block the ability for hosts to respond to such interrogation. The methodology used in [154] leverages the use of a TCP Connect probe which attempts to complete a three-way TCP Handshake between source and target. Hping is used to send "TCP SYN" Traceroutes. Wget [185] is also used to return webpages over common HTTP ports: 80 and 443. Amazon AWS provides native "Security Group" functionality which allows the administrator to control the TCP and UDP ports, along with ICMP traffic and source addresses which are allowed ingress / egress access to your public cloud instances (**Figure 4-20**). The default configuration being a "deny all" inbound. The inherent impact is being assessed without an understanding of the controls and mitigations available.

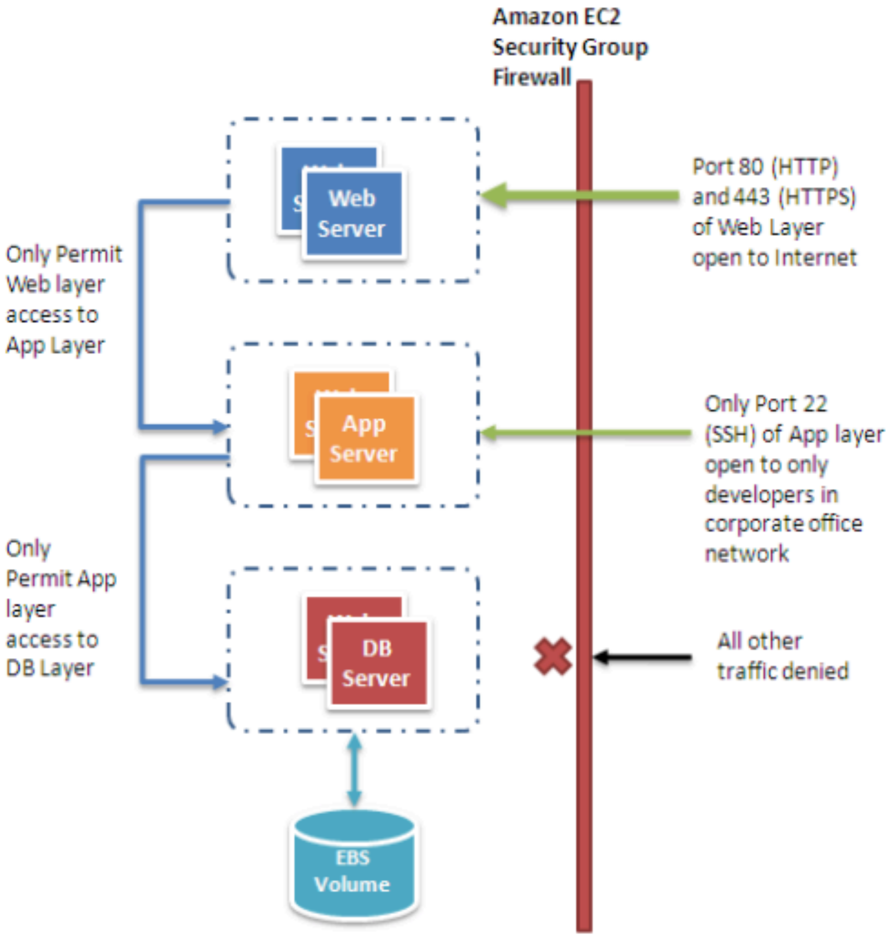
Amazon EC2 Security Group Firewall
Web Server
App Server
DB Server
EBS Volume
Only Permit Web layer access to App Layer
Only Permit App layer access to DB Layer
Port 80 (HTTP) and 443 (HTTPS) of Web Layer open to Internet
Only Port 22 (SSH) of App layer open to only developers in corporate office network
All other traffic denied

Figure 4-20 Amazon Security Groups [72]

There is reference in [154] to "obfuscating co-residence" where a CSP could render co-residence checks ineffective (through firewalls and access control lists (ACLs)) acknowledging that co-residence checks would subsequently need to rely on side-channel attacks which are proven to have a low success rate and will be dissected in **Section 4.9.3**.

**Figure 4-21** visualises Amazon's approach to "*architecting for multitenancy*" (**Section 4.5**). Applications should be designed considering a concept of "least privilege" in which only the access and configuration necessary for a component should be enabled. This is an important design configuration for multitenancy and salient to a discussion regarding mapping the cloud; the discovery of a web server should not infer the detection of databases and file stores which will house sensitive information and protected via appropriate logical and physical controls.

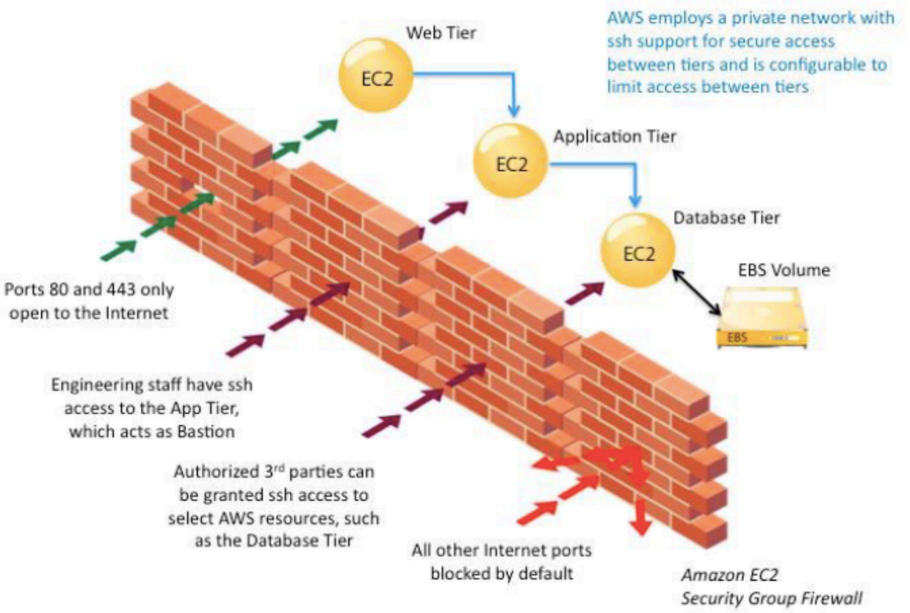
Web Tier
EC2
AWS employs a private network with ssh support for secure access between tiers and is configurable to limit access between tiers
Application Tier
EC2
Database Tier
EC2
EBS Volume
EBS
Ports 80 and 443 only open to the Internet
Engineering staff have ssh access to the App Tier, which acts as Bastion
Authorized 3rd parties can be granted ssh access to select AWS resources, such as the Database Tier
All other Internet ports blocked by default
Amazon EC2 Security Group Firewall

Figure 4-21 AWS EC2 Defence-in-Depth

Table 4-7, I perform a qualitative assessment of the likelihood of initiation for our three threat actors:

| Actor | History | Motivation | LoI Score |
|---|---|---|---|
| Opportunistic | 2: Moderate | 1: Low | Moderate |
| Organised Criminal | 2: Moderate | 3: High | High |
| Nation State | 2: Moderate | 3: High | High |

Table 4-7 Likelihood of Initiation: Mapping the Cloud

The threat event in this scenario (**Table 4-7**) is network scanning. Using AWS EC2 as our case example, the control strength must be "high". Controls exist within VM images and AWS provides a native cloud firewall capability within the hypervisor. Network scanning is a preliminary phase of most cyber-attacks and the skills needed to perform this activity are foundational and complex scanning tasks have been automated through penetration testing platforms such as Kali Linux [186]. Making this task easier within an AWS context is the recent release of Kali Linux within Amazon Marketplace [187]; actors can now paradoxically benefit from the service-based, cost-effective, elastic benefits of public cloud to potentially compromise data and system availability within public cloud (**Figure 4-22**).

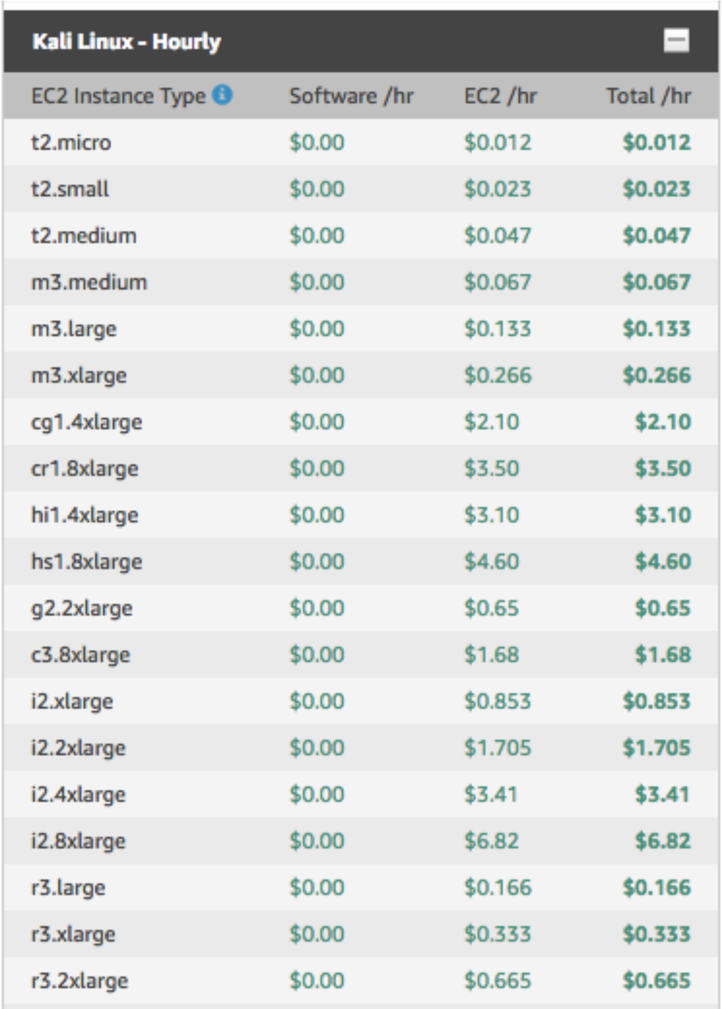

**Kali Linux - Hourly**

| EC2 Instance Type | Software /hr | EC2 /hr | Total /hr |
|---|---|---|---|
| t2.micro | $0.00 | $0.012 | $0.012 |
| t2.small | $0.00 | $0.023 | $0.023 |
| t2.medium | $0.00 | $0.047 | $0.047 |
| m3.medium | $0.00 | $0.067 | $0.067 |
| m3.large | $0.00 | $0.133 | $0.133 |
| m3.xlarge | $0.00 | $0.266 | $0.266 |
| cg1.4xlarge | $0.00 | $2.10 | $2.10 |
| cr1.8xlarge | $0.00 | $3.50 | $3.50 |
| hi1.4xlarge | $0.00 | $3.10 | $3.10 |
| hs1.8xlarge | $0.00 | $4.60 | $4.60 |
| g2.2xlarge | $0.00 | $0.65 | $0.65 |
| c3.8xlarge | $0.00 | $1.68 | $1.68 |
| i2.xlarge | $0.00 | $0.853 | $0.853 |
| i2.2xlarge | $0.00 | $1.705 | $1.705 |
| i2.4xlarge | $0.00 | $3.41 | $3.41 |
| i2.8xlarge | $0.00 | $6.82 | $6.82 |
| r3.large | $0.00 | $0.166 | $0.166 |
| r3.xlarge | $0.00 | $0.333 | $0.333 |
| r3.2xlarge | $0.00 | $0.665 | $0.665 |

Figure 4-22 Extract from AWS Pricing for Kali Linux Instances [187]

The likelihood of success (**Table 4-8)** for a threat actor to successfully map cloud infrastructure is negligible and low for opportunistic and financial criminals respectively. There is a moderate likelihood of success for a nation state actor based on the completeness of available controls. Whilst the threat is well-funded with technical skills and persistence, the controls are robust.

| **Actor** | **Control Strength** | **Threat Strength** | **Likelihood of Success** |
|---|---|---|---|
| Opportunistic | High | Low | Negligible |
| Organised Criminal | High | Moderate | Low |

| | | | |
|---|---|---|---|
| Nation State | High | High | Moderate |

Table 4-8 Likelihood of Success: Mapping Cloud

The residual likelihood (**Table 4-9**) for an opportunistic criminal is low. As our control strength is high, the adversary will need to apply evasion techniques to permeate the firewall-based controls which can be applied. These skills are out of the reach of the opportunistic criminal but a technically-skilled criminal could apply such methods.

| Actor | Likelihood of Initiation | Likelihood of Success | Residual Likelihood |
|---|---|---|---|
| Opportunistic | Moderate | Negligible | Low |
| Organised Criminal | High | Low | Moderate |
| Nation State | High | Moderate | High |

Table 4-9 Residual Likelihood: Mapping Cloud

Our findings here match industry expectations. If a nation state wishes to compromise an organisation they have the tools and funding to do so. My research identifies controls which can mitigate the inherent vulnerabilities outlined in [154].

### 4.8.2. STEP 2: ACHIEVING CO-RESIDENCE

In **Section 4.8.1**, I identified the steps necessary to identify a victim / target VM. If this was completed successfully, which is a complex activity, the attacker now needs to achieve co-residence on the same physical hardware. It is important to acknowledge that all the attacks identified in **Chapter 4** require physical co-residency.

For a VM escape or rogue hypervisor exploit, the attacker will need to ensure that she is resident on the same physical hardware as the victim machine. Any other form of attack would be opportunistic in nature. There are several means of achieving co-residence.

Xen hypervisor uses the Dom0 Host VM. When a guest VM communicates with Dom0, traffic is routed through the hypervisor layer. As the hypervisor is exposing virtual network resources, it is the only network hop between the two machines when performing a TraceRoute. In a situation where only a single hop exists, it can be determined that the same physical machine is being used.

Elisan et al [166, p. 172] explain how tools have been authored to cover co-residence and make the process of identification straight-forward. It is reported that "…*in some natural attack scenarios, just a few dollars invested in launching VMs can result in a 40% chance of placing a malicious VM on the same physical server as your target"* [154]. This process in known as "instance placement" and the figures here (40%) [154] relate to an m1.small EC2 instance type [188]. This is an important detail as the (now legacy) m1 Instance Type was a general purpose and relatively low-specification VM. It is unlikely that such images would be deployed for an enterprise customer thus rendering findings of limited value.

I assert that a 40% chance of success is only a 40% chance of achieving success across a single phase of a highly complex, multi-stage attack. This “high” likelihood isn’t representative of an overall chance of success in extracting sensitive information.

Other methods of confirming co-residence in Amazon EC2 include [154]:

- Identical Dom 0 IP Addresses
- Small IP Packet Round Trip Times (RTTs)
- Numerically close IP addresses (within 7)

In **Section 4.9.3**, I explore direct communication from VM-to-VM which avoid the hypervisor layer, essentially these are covert channels. The ability for successful VM-to-VM communication infers co-residence on the same physical hypervisor although this approach bears an extremely low level of likelihood.

A concept of brute-force placement of virtual machines has yielded some evidence of success. Brute forcing is the process of launching a high-volume of virtual machines over an extended period to have machines resident on the same physical target. By creating machines with similar creation parameters, it is believed that these machines will be created synchronously across physical machines. Results are scarce and the control set too small to indicate conclusive proof although Ristenpart et al [154] suggest that for an m1.small (instance type) VM, co-residence can be achieved some of the time. The methodology followed is included below:

1. Enumerate a set of victims
2. Assess which availability zone and resource type the IP belongs to
3. Now repeatedly spin up ‘probe instances’ in these regions and see if co residence is achieved
4. If not, terminate

Ristenpart et al [154] performed this work with limited success; in their target group of 1686 servers, over 18 days, an 8.4% coverage was achieved. This evidences a high concentration of effort, time and money for limited returns.

Instance Flooding is another form of placement strategy. It is analogous in structure to a buffer overflow attack [189]. The theory is that an attack would run as many virtual machines as possible in parallel in an availability zone and of the instance type required. This attack vector apparently exploits features of the placement algorithms in AWS. The attack is however predicated on knowing precisely when a target VM has been launched and being able to launch multiple VMs at this time. Amazon also limits a single account to launch 20 concurrent VMs thus further mitigating the vector.

Based on my research and the findings above, my qualitative assessment of the likelihood of initiation for co-residence is included below in **Table 4-10**:

| **Actor** | **History** | **Motivation** | **LoI Score** |
|---|---|---|---|
| Opportunistic | 1: Negligible | 1: Negligible | Low |
| Organised Criminal | 2: Low | 3: Low | Low |
| Nation State | 2: Moderate | 2: Moderate | Moderate |

Table 4-10 Likelihood of Initiation: Co-Residence

An opportunistic criminal is priced out of this attack through the effort in time and money needed to achieve success. It is possible that a financially motivated criminal has the skills and funding to achieve success some of the time although the return on investment isn't there to make brute-forcing of VMs or instance flooding viable business models. Nation state actors have the skills, funding, motivation and commitment to carry out such an attack if they feel that a target can be compromised through such means. Recent research would however suggest that more direct attacks leveraging social media campaigns would be easier to implement with a far higher likelihood of success [190].

| Actor | Control Strength | Threat Strength | Likelihood of Success |
|---|---|---|---|
| Opportunistic | Moderate | Low | Low |
| Organised Criminal | Moderate | Moderate | Moderate |
| Nation State | Moderate | High | High |

Table 4-11 Likelihood of Success: Co-Residence

I decided upon a control strength of "moderate" in relation to a co-residence attack (**Table 4-11**). Inherent controls exist to limit the numbers of machines that can be brought online in parallel [154]. AWS fraud protection mechanisms leverage machine learning and are mature in themselves [191] . To create an account with AWS, you must register an active credit card and payment address [192]. This essentially means that a criminal would have already have had to successfully harvested or purchased an identity and record a payment card on file which has not been reported lost or stolen.

| Actor | Likelihood of Initiation | Likelihood of Success | Residual Likelihood |
|---|---|---|---|
| Opportunistic | Low | Low | Low |
| Organised Criminal | Low | Moderate | Moderate |
| Nation State | Moderate | High | High |

Table 4-12 Residual Likelihood: Co-Residence

The barriers to entry for this phase are noticeably higher than those associated with network mapping. The likelihood of creating a virtual machine with a specific configuration and having that virtual machine coexist on the same physical hypervisor as a victim virtual machine are almost zero. The likelihood increases as investments are made in the time, funding and parallel processing.

## 4.9. ATTACKS ON THE HYPERVISOR

Context is critical across cyber and information security. **Sections 4.8.1 / 4.8.2** were included to stress the commitment and capability attributes that a threat actor must possess to perform a targeted attack leveraging a hypervisor as her point of entry. There are simply easier methods to 'pwn'[17] a victim. If preliminary footprinting and placement activity is successful, the attack now turns to compromising the hypervisor; the goal being the extract of or prevention of access to, or interference with data.

Attacks on the hypervisor are focused in two core areas:

**VM Escape:** An attacker can break from a VM and exploit arbitrary code with hypervisor permissions

**Rogue Hypervisor:** Attacker takes control of hypervisor.

Both attacks exploit the vulnerabilities called out by ENISA related to resource isolation. At face value both vulnerabilities appear critical and should be treated as a serious consideration in public cloud adoption. I will proceed to assess under what conditions these attacks could take place and if controls could be applied to mitigate or remove the threat (or vulnerability).

### 4.9.1. VM ESCAPE

The objective of a VM Escape attack is to conduct arbitrary actions with the privileges of the hypervisor layer as opposed to a guest VM [154]. The ability to conduct this form of attack violates the fundamental principle of virtualisation which is to ensure resource isolation between virtual machine instances across the physical hardware [151].

Hypervisors are required to run at the most privileged level of operation, by breaking out of a guest VM and running commands in a hypervisor context, the attacker is potentially able to take control of the underlying hardware and all VM images. *"VM escape attacks are catastrophic threats to the cloud platform as the hypervisor is a single point of failure"* [18]. My goal, when identifying VM Escape attacks as a credible threat to public cloud environments, is to assess the likelihood of a threat actor exploiting vulnerabilities of multitenancy through a VM Escape threat event.

I began my research with an assessment of likelihood and a review of publicly reported VM escape vulnerabilities:

In 2015, a vulnerability was discovered in the XEN Hypervisor which allowed administrators of paravirtualised hosts to exploit the mod_l2_entry function in arch/x86/mm.c file through a lack of level 2 page table entry validation. This elevation of privilege attack required the host to be running paravirtualisation and Xen versions 3.4 through 4.6.x. [193]. The discovery of the Virtualised Environment Neglected Operations Manipulation (VENOM) vulnerability [194] was the first widespread vulnerability which affected type 1 hypervisors of multiple vendors. VENOM exploited the Floppy Disk Controller (FDC) resident in Xen Hypervisor[18] (4.5.x and earlier) and had the potential to allow local (guest) users to cause a) denial of service attack or b) allow the escape from a

[17] PWN is a colloquial term and refers to the process of compromising an endpoint

[18] VENOM exploited the QEMU's virtual floppy disk controller. QEMU being used by multiple vendors: Xen, KVM and Virtualbox.

guest VM and execute arbitrary code on the Host VM (Dom0). VENOM was discovered (and named) by cyber security firm Crowdstrike [172]; responsible disclosure was followed and patches were developed for vendor (hypervisor) offerings. This collaborative approach avoided VENOM wreaking havoc in the wild. It is important to recognise that such an attacker would require local administrative-level rights to the guest VM in able to exploit the vulnerability.

I do not want to underplay the significance of VENOM or CVE-2015-7835 before it but it is important to contextualise the vulnerability. VENOM, in some circumstances, would allow the escape from a guest OS and allow the execution of code at a higher privilege level. VENOM could also endanger the confidentiality, integrity or availability of another guest VMs on the same physical server. Abuse of local administrative credentials are certainly not a new threat, exclusive to public cloud. Tools [195, 196] exist for multiple flavours of Microsoft Windows to essentially elevate from Local Administrator to Domain Administrator thus providing a network-wide level of administrative access and arbitrary code execution. The difference with VENOM or other conceptual VM Escape is the multitenancy position: other guest VMs will almost certainly belong to other customers. For VENOM to be leveraged in a targeted attack, the malicious actor would need to have performed network mapping and co-residence activities identified in this paper. Assuming the victim in question was housing sensitive information and / or the property of an enterprise organisation, it is a safe assumption to make that a network / host firewall or AWS security group would stop the assailant early in the kill chain[19].

Many proof of concept exploits have been discovered for type 2 hypervisors but even in these situations, where there is larger attack surface (compared with type 1), the impact of these attacks has been (perhaps unintentionally) exaggerated. In his 2014 Blackhat US talk [181], Wojtczuk talks about several hypervisor vulnerabilities which have been overplayed asserting, for example, that high-profile VM Escape attacks such as CVE-2014-7188 [197] allowed for the hypervisor to leak "some memory" but that the vulnerability has received an "unjust amount of interest" [181] (28:00 mins).

Denial-of-Service is another outcome of a successful VM Escape attack. Perez-Botero et al. [155] explain that VM Exit handling code does not possess the data structures necessary to invoke exploitable effects other than a host or a guest VM crash.

I expect that VM Escape vulnerabilities will continue to be discovered and analysed academically. The proliferation of bug bounty programmes widens the net of potential testers analysing a system. Conferences such as the Zero Day Initiative's "PWN2OWN" [157], incentivises hackers and security researchers to find vulnerabilities in mainstream virtualisation platforms. This research can be lucrative, their 2017 conference is offering $100,000 to anyone able launch an attack on VMware Workstation or Microsoft Hyper-V which results in a guest OS user, with admin privileges, being able to execute arbitrary code on the host operating system [157]. During these events, the subject of mitigations and context are rarely discussed.

[19] Phases of a cyber-attack outlined in [238]. By preventing an assaliant at a specific phase in the chain, the attack is thwarted.

In **Table 4-13**, I analyse the likelihood of a VM Escape attack:

| **Actor** | **History** | **Motivation** | **LoI Score** |
|---|---|---|---|
| Opportunistic | 1: Negligible | 1: Negligible | Low |
| Organised Criminal | 2: Low | 2:Low | Low |
| Nation State | **2: Moderate** | **3: Moderate** | **Moderate** |

Table 4-13 Likelihood of Initiation: VM Escape

A very positive result of research scrutiny is that vulnerabilities remain in-the-wild for a very short period, if at all. VENOM had the potential to cause a serious problem to organisations although it was disclosed appropriately and patches released in an expedient fashion. This makes the window of opportunity finite for any attacker.

**Table 4-14** outlines the likelihood of success for a VM Escape attack based on the context provided in this section:

| **Actor** | **Control Strength** | **Threat Strength** | **Likelihood of Success** |
|---|---|---|---|
| Opportunistic | Moderate | Low | Low |
| Organised Criminal | Moderate | Moderate | Moderate |
| Nation State | Moderate | High | High |

Table 4-14 Likelihood of Success: VM Escape

Given the reported history of government agencies discovering and “hoarding” software / hardware vulnerabilities [198], we cannot say with certainty that there is no history or motivation of nation states exploiting hypervisor vulnerabilities although history does suggest that attacks purportedly to be the work of ”nation states” would be focused higher up the software stack and targeted users [199, 190]; often seen as the weakest link. This brings us back to the motivation consideration; a VM escape attack carries a complex vector and a limited likelihood of success.

**Table 4-15** identifies that whilst VM Escape attacks are technically possible, the likelihood of such vulnerability being exploited is low in most cases.

| Actor | Likelihood of Initiation | Likelihood of Success | Residual Likelihood |
|---|---|---|---|
| Opportunistic | Low | Low | Negligible |
| Organised Criminal | Low | Moderate | Moderate |
| Nation State | Moderate | High | High |

Table 4-15 Residual Likelihood: VM Escape

The overwhelming majority of VM Escape attacks have focused on vulnerabilities in type 2 hypervisors. There is a high residual likelihood of a nation state carrying out this attack as there are limited controls available if a vulnerability is discovered in device emulation software (as with VENOM). The control is reactive and the vulnerability removed through patching.

### 4.9.2. ROGUE HYPERVISOR

As technology paradigms evolve, so does malware. Rutkowaska [179] defines four levels of sophistication in relation to rootkits and malware (**Table 4-16**). The most sophisticated of these being relevant for a discussion on virtualised rootkits; also, known as a "rogue hypervisor" or "hyperjacking":

| Malware Type | Description |
|---|---|
| Type 0 | Does not alter the operating system |
| Type 1 | Malware which modifies things which should never be modified outside of approved OS code (CPU registers, OS itself). |
| Type 2 | Malware that modifies things that **were** designed to be modified. Self-modifying code in data sections of binary files. |
| Type 3 | Virtual rootkits which can subvert a running operating system without modifying anything at all inside it. |

Table 4-16 Malware Types [179]

A virtual rootkit / rogue hypervisor moves from operating at the same level as an operating system (ring 0), to the level below it (ring -1) [166]. I will focus on hypervisor virtual machine (HVM) rootkits which leverage hardware virtualisation support to replace the underlying hypervisor completely with its own custom hypervisor and then envelope the currently-running operating systems (hosts and guests) on-the-fly [166]. Virtual Machine based rootkits (VMBRs) exist which rely on alteration to virtualisation software rather than hardware virtualisation support. Subvirt [200] is probably the most recognised VMBR and therefore warrants acknowledgement in this paper. SubVirt targets Intel x86 processor architecture and inserts itself underneath the host OS. This results in the creation of a new hypervisor layer (**Figure 4-23)**. VMBR have a significant performance overhead and require rebooting before operational. SubVirt required modification to the boot sector of a hard disk making it susceptible to detection. VMBRs are destructive although they are considered a traditional type of rootkit [166, p. 169] which by relative terms are easy to detect and lack persistence; they will therefore not be discussed further.

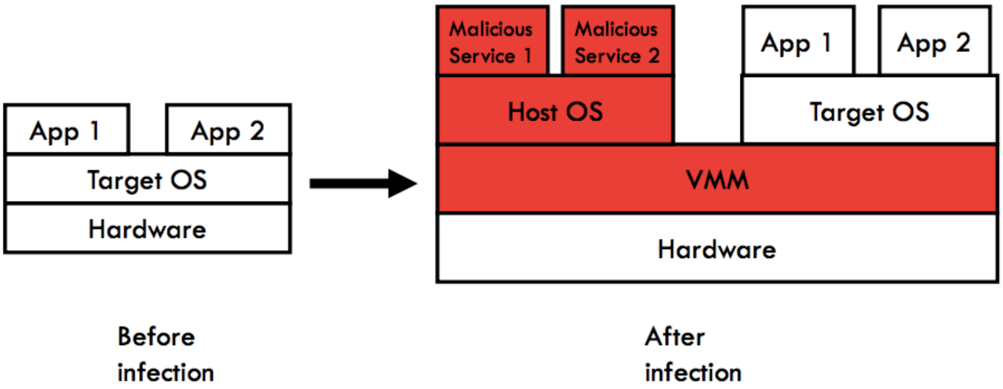

**Figure 4-23 VMBR Architecture [200]**

It is important to explicitly define the process necessary for a malicious threat actor to install and operating a HVM rootkit (HVMR). The consequences of such malware being resident within a physical machine are potentially disastrous, allowing the attacker full access to physical hardware and all virtualised instances on a machine. For a rogue hypervisor attack to take place, the installation of malicious rootkit code must occur – having researched the feasibility and complexity of this process, I felt it prudent to explicitly document the steps necessary [166, p. 178]:

1. Install a kernel driver in the guest OS
2. Find and initialise hardware virtualisation support
3. Load the malicious hypervisor code into memory from the driver
4. Create a new VM to place the host operating system inside
5. Bind the new VM to the rootkit's hypervisor
6. Launch the new VM – this will switch the host into guest mode.

Each of the above steps requires persistence and a skilled attacker. HVMRs do not require a reboot (as with VMRs). The Ring -1 layer essentially adds hardware support for virtualisation software. It is this layer which is exploited in a HVMR attack. Across AMD parlance, this technology is known as AMD-V Secure Virtual Machine (SVM) [165] and the Intel equivalent is called Virtualisation Technology extensions (VT-x) [164]. In the interests of impartiality and completeness, we will research a HVMR on both the Intel and the AMD architecture.

The purpose of this chapter is to understand the likelihood of a threat actor compromising the confidentiality, integrity or availability of information through an event which exploits a vulnerability. We will continue by exploring two Hypervisor Virtual Machine Rootkits to understand the likelihood of a successful compromise. BluePill is a HVMR introduced at Blackhat Europe 2006 by Joanna Rutkowska [179]. BluePill was originally based on the AMD SVM architecture. With BluePill, the Host OS is moved, without reboot, into a virtual machine via the AMD SVM extensions.

Vitriol is a second HVMR, released around the same time as BluePill. Vitriol is an Intel VT-x HVMR which leverages hardware support to raise and lower the execution level of the CPU, VMX Root for Ring 0 and Non-

VMX root "less privileged" mode. Guest OS are launched in non-VMX root but able to call a VM Exit instruction when they need to access privileged space [166].

Unfortunately, more recent examples of HVMR rootkits were unavailable to critically assess. This suggests that attacks exploiting inherent vulnerabilities in a hypervisor are either ineffective or laboriously difficult and time consuming. With so much written about hypervisor exploits and ENISA defining resource isolation risks as critically important, I was expecting to uncover a plethora of documented instances of virtualised rootkit malware "in-the-wild[20]". This wasn't the case.

Much like most areas of information and cyber security, vulnerabilities are discovered and defenders (vendors, researchers, ethical hackers) look to provide fixes, software updates and patches. A year after Rutkowska released BluePill, a group of researchers attended Blackhat US 2007 and presented the drolly titled "Don't Tell Joanna, the Virtualized Rootkit is Dead" [201] in which methods to detect the "undetectable" rootkit malware was put forward.

**Table 4-17** identifies that rouge hypervisor attacks are unlikely to occur in most cases:

| **Actor** | **History** | **Motivation** | **LoI Score** |
|---|---|---|---|
| Opportunistic | 1: Negligible | 1: Negligible | Low |
| Organised Criminal | 2: Low | 3: Low | Low |
| Nation State | 2: Moderate | 3: Moderate | Moderate |

**Table 4-17 Likelihood of Initiation: Rogue Hypervisor**

There have been no reported instances of HVMR in the wild on either the Intel or AMD instruction sets. The stealth and invisibility of kernel level malware is the desired result of all malware authors although misconfiguration and compatibility issues mean that there is a high likelihood of discover by a user. The Windows "Blue Screen of Death"[21] being an unfortunately consequence of kernel driver modification (nefarious or otherwise) for many years.

[20] Found within environments considered "production" in nature: enterprise networks, Internet components or consumer devices.

[21] Colloquial term, ubiquitous in the IT world, to describe a fatal system error across Microsoft Windows [239] technologies.

| Actor | Control Strength | Threat Strength | Likelihood of Success |
|---|---|---|---|
| Opportunistic | High | Low | Negligible |
| Organised Criminal | High | Moderate | Moderate |
| Nation State | High | High | Moderate |

Table 4-18 Likelihood of Success: Rogue Hypervisor

A skilled, persistent adversary is required to launch a rogue hypervisor attack. Success is predicated on the completion of several steps which we have outlined above. Having analysed the steps (**Table 4-18**), they are not out of the reach of an organised cyber group although the motivation would be questionable to invest the time and energy necessary to make this a financially rewarding venture.

| Actor | Likelihood of Initiation | Likelihood of Success | Residual Likelihood |
|---|---|---|---|
| Opportunistic | Low | Negligible | Low |
| Organised Criminal | Low | Moderate | Moderate |
| Nation State | Moderate | Moderate | Moderate |

Table 4-19 Residual Likelihood: Rogue Hypervisor

The prerequisite steps necessary to successfully locate a victim physical machine are outside the realistic reach of an opportunistic criminal or those with financial motivation. My assessment (**Table 4-19**) is that whilst the vector exists, it is unlikely to be exploited.

The creation of kernel-level rootkits is incredibly difficult; the steps for infection are complicated and compatibility issues are overtly obvious to a user. The research community [201] is also motivated to debunk theories of stealth as asserted by Rutkowaska. Much like the adversarial threat attributes defined for my isolation decomposition, the same can be applied to security researchers: Capability, Motivation, Commitment and History.

### 4.9.3. INTER-VM ATTACKS

Inter-VM attacks bypass the hypervisor entirely. The hypervisor is the brain of the virtual machine (VM) environment and as such, organisations understand the importance of hardening this layer. An inter-VM attack looks to launch an attack from one VM to another VM residing on the same physical hardware with the goal of information leakage or denial of service. Inter-VM attacks are harder to detect than those targeting the hypervisor as hypervisors are a focus of attention for security professionals.

It is important to acknowledge that direct VM-to-VM communication is leveraged for non-malicious traffic flows. Zhang [18] identifies several methods for inter-VM communication which avoid the hypervisor layer for purposes of performance and convenience; these direct communications introduce vulnerabilities which could be exploited by a malicious actor.

Upon completion of steps 1 & 2 outlined in **Section 4.8,** the attacker must now leverage a side channel attack to obtain information which is believed to be retained within the isolation perimeters of VM and hypervisor. Side channel attacks are commonly exploited in cryptographic attacks to uncover sensitive key information. Side channels can also be leveraged in virtualisation attacks across any sharable resource. Pertinent to this paper, side channel attacks for the extraction of cryptographic material (in a virtualisation context) present "serious practical challenge" [154]. Namely, coarser scheduling, double indirection of memory addresses and load from other machines.

Contemporary literature challenges the views of [154]. Apecechea et al. [202] suggesting that fine-grained side channel attacks are possible in bare-metal, virtualised hardware used in Xen. Their work outlines vulnerabilities in OpenSSL [203], PolarSSL [204] and Libgcrypt [205] which makes them susceptible to correlation attacks when run on a Xen or VMWare hypervisor. The attack proposed would result in the extraction of Advanced Encryption Scheme (AES) cryptographic key information. Prior to this study, it was believed that multicore processing mitigated the threat of side channel attacks although Apecechea et al. [202] assert that their attack is successful in AWS even across multiple cores on the same machine.

If such sensitive information could be extracted via this method, the impact of inter VM-based attacks would grow significantly; however, another primary objective was to understand if [202] introduces new attack vectors (via threat events) which needed including in our study. I conclude that whilst [202] is an interesting academic study and there is a theoretical avenue for the extraction of AES keys through side channel, there are several mitigations and compensating controls which make the attack complicated, namely [202]:

- The attack is considered infeasible when dealing with AES 256 bit keys.
- Only partial keys have ever been disclosed in an intra-VM attack[22].
- Obtaining co-residence is problematic.
- AES-NI hardware support completely mitigates the threat – XEN 4.0.1 supports usage in guest OS [206]

I wanted to include this attack vector as it challenges seminal work suggesting that side channel attacks are limited to coarse-grained information. I will not however be including it in our analysis inter-VM side channel attacks as AWS runs versions of the Xen hypervisor which are not vulnerable to the attack.

Having researched the available academic and industry information [18, 154, 202] I conclude that five attack channels could be leveraged for inter-VM data exfiltration. These areas are more coarse-grained than a cryptographic side-channel and comparatively less dangerous – the information leaked is not of the sensitivity of a secret or pre-shared key (or components thereof). In Error! Reference source not found. we will present e ach vector for side channel attack and supporting mitigations, where appropriate:

[22] I am not underplaying the significance of partial key discovery. The weakening of an AES 256 bit key makes it more susceptible to brute force attacks however partial key recover would still require extensive work on the part of the attacker.

| **Side Channel Attack** | **Explanation** | **Mitigations** |
|---|---|---|
| Cache usage | Measuring the CPU utilisation of a physical machine. Vector potentially for a denial of service attack through the placement of additional load on a server during peak periods | Critical systems are set to provide redundancy and load balancing. AWS provides auto-scaling and failover capabilities. AWS Cloudwatch can be automated to report on such metrics. |
| Load-based co-residence | Attacker induces computational load on a VM and measures timing differences when local (co-resident) and remote | Limited scientific testing: Many factors can influence response times and CPU performance. Even less conclusive that Step 2: Achieving Co-Residence although this would remove the efficacy of firewall controls for network probing. |
| Estimating Traffic Rates | Load measurements to estimate numbers of visitors to a VM. An attacker could estimate peak period of activity for a virtual machine | E-Commerce platforms are busy on public holidays, social media websites are busy in the evening. Passive analysis is easier and more likely to yield a positive result. |
| Keystroke timing attack | Measuring the time between keystrokes for, say, password entry and recovering a password. This is achieved in the cloud through cache-based load measurements. | The tests performed [154] were in a local environment specifically to overcome the complications / controls present in EC2, namely multicore architecture. This attack would only ever be feasible in EC2 if two VMs time-shared a core. There are simply easier ways (Phishing, Trojan malware, browser plugins) to steal passwords. |
| Sniffing attack | In [207] the author shows how to sniff traffic between virtual machines running on the same physical machine | "It is not possible for a virtual instance running in promiscuous mode to receive or sniff traffic that is intended for a different virtual instance. While customers can elect to place their interfaces into promiscuous mode, the hypervisor will not deliver any traffic to an instance that is not addressed to it. Even two virtual instances that are owned by the same customer located on the same physical host cannot listen to each other's traffic". [208] |

Table 4-20 Side Channel Attack Vectors

Having detailed the side-channel, inter-VM vectors available to an attacker, I will now perform a qualitative analysis covering the likelihood of such an attack being conducted (**Table 4-21**):

| Actor | History | Motivation | LoI Score |
|---|---|---|---|
| Opportunistic | Low | Low | Low |
| Organised Criminal | Low | Moderate | Moderate |
| Nation State | Moderate | High | High |

Table 4-21 Likelihood of Initiation: Inter VM Attack

If the objective is to extract passwords through keystroke timing attacks then cheaper, much more effective, routes exist through a combination of social engineering and the deployment of trojan malware such as Dyre [209]. **Table 4-22** details my qualitative likelihood of success assessment:

| Actor | Control Strength | Threat Strength | Likelihood of Success |
|---|---|---|---|
| Opportunistic | Moderate | Low | Low |
| Organised Criminal | Moderate | Moderate | Moderate |
| Nation State | Moderate | High | High |

Table 4-22 Likelihood of Success: Inter VM Attack

Successful attacks are predicated on transactions running on a single core architecture; something which is unlikely in AWS due to its multitenant, shared architecture. Sniffing attacks at a network level are not possible in AWS due to the restrictions on “promiscuous mode” Network Interface Card (NIC) configuration [29, p. 13] . This restriction makes the likelihood of success significantly lower on an AWS platform that an environment without network-level controls (**Table 4-23**).

| Actor | Likelihood of Initiation | Likelihood of Success | Residual Likelihood |
|---|---|---|---|
| Opportunistic | Low | Low | Negligible |
| Organised Criminal | Moderate | Moderate | Moderate |
| Nation State | High | High | High |

**Table 4-23 Residual Likelihood: Inter VM Attack**

### 4.9.4. DENIAL OF SERVICE ATTACKS

Denial of service (DoS) is the exhaustion of, or removal of access to, resources from a malicious or accidental party. Across 2016, distributed denial of service (DDoS) was in the headlines frequently after several high-profile attacks which brought many organisations to a standstill. In fact, in December 2016, a DDoS attack essentially "brought down" the Internet for a matter of hours [210]. Irrespective of the specific vector, DoS is about the exhaustion of resources, be those resources network, application, compute, disk or memory. Hyde [211] identifies DoS in a virtualisation context and I will therefore present a risk assessment to identify likelihood.

DoS in public cloud introduces a vulnerability which is not present in a single tenant solution. DoS attacks in public cloud, whilst potentially carrying a catastrophic impact, are easily mitigated through proper configuration of the hypervisor [211]. A vector presents itself when the attacker can exhaust, or remove access to, resources that are shared for all tenants on the physical hardware. The objective of the attack is to restrict resources to other users.

Our study focuses on public cloud through the Amazon EC2 Cloud. Controls exist within the platform designed specifically to mitigate and protect against resource exhaustion attacks. Amazon Cloudwatch is an event logging system designed to alert administrators to system and application events and *"delivers a near real-time stream of system events that describe changes in AWS resources"* [115]. Due to the maturity of controls in this space, a significant amount of time will not be spent on studying denial of service in the context of compute resource exhaustion.

From a network perspective, I assert that AWS improves an organisation's ability mitigate DoS and DDoS attacks. The AWS network is highly distributed and connected to multiple Internet Service Providers across each Availability Zone. Amazon are responsible for the global infrastructure elements of their platform and the customer must ensure that their VM instances are secured against application layer attacks [29]. AWS API endpoints (used for customer management) are a vector for DoS although Amazon have recognised this deploying (API) endpoints on the same Internet-scale infrastructure as Amazon's retail organisation where uptime and availability are of paramount importance [29].

For completeness, denial of service attacks have been included in this study. The exhaustion of resources on a physical machine hosting multiple customers presents a threat event exacerbated by cloud. **Table 4-24** includes my assessment of likelihood for DoS in a multitenant environment:

| **Actor** | **History** | **Motivation** | **LoI Score** |
|---|---|---|---|
| Opportunistic | Low | Low | Low |
| Organised Criminal | Low | Moderate | Moderate |
| Nation State | Low | Moderate | High |

**Table 4-24 Likelihood of Initiation: Denial of Service**

I was unable to find documented cases of a cloud-based resource exhaustion attack in a production environment. The motivation for such an event would not exist for an opportunistic criminal. It is possible that an organised financial criminal could consider using this vector to extort money from a customer by threatening to render services unavailable but as we have seen in this subsection, the architecture of AWS caters for flexibility and availability. Single points of failure (SPOF) are avoided and any resource exhaustion or DoS attack would have little to no impact on an enterprise that have designed a highly available system.

| **Actor** | **Control Strength** | **Threat Strength** | **Likelihood of Success** |
|---|---|---|---|
| Opportunistic | Moderate | Low | Low |
| Organised Criminal | Moderate | Moderate | Moderate |
| Nation State | Moderate | High | High |

Table 4-25 Likelihood of Success: Denial of Service

Likelihood of success in this context (**Table 4-25**) would require the attack to render resources unavailable to the victim. As I have documented in this section, the controls available natively in AWS make denial of service, at an infrastructure and / or network level, a low probability threat. It is important to reemphasise that Amazon is responsible for the security of the cloud, the customer is responsible for their security in the cloud [138]. As such, the customer is responsible for the deployment of security controls for their applications; this would include components to mitigation application-layer DoS such was Web Application Firewalls (WAF) / Intrusion Prevention Systems (IPS). AWS Marketplace does provide a comprehensive catalogue of security controls for customer consumption [56].

| **Actor** | **Likelihood of Initiation** | **Likelihood of Success** | **Residual Likelihood** |
|---|---|---|---|
| Opportunistic | Low | Low | Negligible |
| Organised Criminal | Moderate | Moderate | Moderate |
| Nation State | High | High | High |

Table 4-26 Residual Likelihood: Denial of Service

DoS attacks in public cloud carry a greater impact to those on a single tenant platform although the controls available to a customer are mature and robust.

## 4.10. COMPROMISE OF DATA IN THE PUBLIC CLOUD

The previous section reviewed the "tools, techniques and procedures" that a threat actor could adopt to alter the CIA of information in a public cloud environment. The objective was to apply qualitative reasoning to obtain indicative likelihood of an attack. Likelihood is also evaluated by reviewing the history of such attacks. Threat Capabilities are covered in **Section 3.4,** the history attribute is used extensively in risk management to gauge future likelihood.

For all that has been written about the risks of cloud computing, few reports were available to the author which covered the loss of confidentiality, integrity or availability of information through a threat event associated with multitenancy in a public cloud. In fact, public cloud breaches that are documented are a result of vulnerabilities which would be applicable to an on-premise model (physical or virtual). I will proceed to identity several high-profile examples of these in the following paragraphs.

Apple iCloud [212] is often labelled as a "hack" although their high-profile compromise of 2014 was the result of a poorly configured Application Programming Interface (API) which allowed for unlimited attempts at a password which facilitated brute-forcing of credentials [213].

In 2015, the anti-virus firm Bitdefender [214] suffered from a confirmed breach of customer data [215]. The data in question did reside in AWS Elastic Cloud Compute (Public Cloud). Upon analysis of the attack, it was confirmed that credentials were stored and processed in an unencrypted format. It was purported [216] that the attackers had access to the (Bitdefender) network and were taking the credentials as they traversed the network. The headlines read that attacks compromised data within public cloud which is technically true; it was however the responsibility of the customer to make sure that their data was encrypted at rest and in transit.

To retain scope and structure to this study, we will not review all breaches of recent times. The Breach Level Index (BLI) [217] is an excellent resource should anyone wish to review major breaches of the past four years. The BLI also attempts to assess compensating controls and context which is in line with my approach to risk management.

If we analyse the largest, most high profile data breaches of the past 12 months, attackers are leveraging tried-and-tested means of obtaining a foothold within an organisation for the purposes of obtaining information. The adult website "Adult Friend Finder" was hacked and reportedly [218] 400 million accounts were exposed, this was the largest breach of last year and will be reviewed as a result. This breach highlights that organisations are simply not doing the basics. The vector in this case was purportedly a local file include vulnerability (LFI) [219], something OWASP have provided guidance around for many years [220]. In our discussion regarding inherent impact, this breach identifies a lack of due diligence in the identification of threat events but also controls to mitigate impact; the records in the case of Adult Friend Finder were stored unencrypted [218] making exfiltration and reuse a trivial task.

While organisations are failing to perform good security hygiene, it is ill-advised to assert that the exploitation of multitenancy would be a viable vector for either an opportunist criminal or a financially-motivated adversary. The time, cost, and persistency required along with the high-likelihood of failure make resource isolation attacks a poor investment. A state-sponsored adversary would have the skills, time, motivation and funding to carry out attacks such as a rogue hypervisor exploit or achieve co-residency and launch a denial of service through an I/O exhaustion attack but there are quite simply other more readily available, easier, stealthier ways of exfiltration information [190].

## 4.11. COMPARISIONS WITH OTHER THREATS AND VULNERABIILTIES

An organisation is idiomatically only as strong as its weakest link. Whilst it is prudent to acknowledge the threats and vulnerabilities associated with public cloud computing, there are a myriad of risks to the confidentiality, integrity and availability which exist across enterprise environments and through my risk analysis, I assert that these are significantly more easily exploited.

Information security is about people, process and technology. I would therefore like to identify a vulnerability in each of these categories which carries with it a significantly higher likelihood to an enterprise than its adoption of public cloud computing or the technical considerations of resource isolation.

### 4.11.1. PEOPLE: CREDENTIAL MANAGEMENT: PASSWORDS

Human beings are not designed to storage multiple, complex passwords. Irrespective of an on-premise or a public cloud-based deployment, weak credential management is responsible for many high-profile data breaches; in a recent CSA study [221], 22 percent of respondents cited credential compromise as the threat event which resulted in the compromise of CIA.

Unlike the vulnerabilities I have identified for resource isolation, the mitigations associated with credential management issues are limited. Organisations can look to multi-factor authentication products and the introduction of out-of-band [23] credential delivery mechanisms although many people will inevitably fall for well-crafted social engineering attacks.

Phishing attacks continue to evolve and with the ubiquitous availability of digital certificates [222] contributes to the sophistication in appearance of sites aimed as tricking users into entering credentials.

Paradoxically, cloud computing and the benefits of elastic, service-based subscriptions is compounding the problem. Cyber criminals are hosting virtualised public-cloud infrastructure to attempt brute-force password cracking at scale. As graphics cards continue to be optimised for Bitcoin [223] mining, criminals can therefore leverage the same infrastructure [224] for cracking of poorly encrypted password credentials.

| **Actor** | **History** | **Motivation** | **LoI Score** |
|---|---|---|---|
| Opportunistic | High | High | High |
| Organised Criminal | High | High | High |
| Nation State | High | High | High |

Table 4-27 Likelihood of Initiation: Password Attacks

**Table 4-27** outlines the likelihood of initiation for a social engineering attack. The opportunistic criminal needs no specialised skills or technology to deliver a successful attack. Social engineering can take many forms: phone calls, SMS messages, phishing emails. It is a highly effective attack vector with low barriers to entry.

| **Actor** | **Control Strength** | **Threat Strength** | **Likelihood of Success** |
|---|---|---|---|
| Opportunistic | Moderate | Low | Low |
| Organised Criminal | Moderate | Moderate | Moderate |

Table 4-28 Likelihood of Success: Password Attacks

[23] Out of Band in this context covers the delivery of credentials

| | | | |
|---|---|---|---|
| Nation State | Moderate | High | High |

Table 4-28 (Cont.) Likelihood of Success: Password Attacks

Controls exist which significantly mitigate the impact of a successful social engineering attack. The implementation of multifactor authentication and out-of-band password delivery although these controls impact user experience and are therefore sacrificed in favour of a frictionless user experience.

| **Actor** | **Likelihood of Initiation** | **Likelihood of Success** | **Residual Likelihood** |
|---|---|---|---|
| Opportunistic | High | Low | Moderate |
| Organised Criminal | High | Moderate | High |
| Nation State | High | High | High |

Table 4-29 Residual Likelihood: Social Engineering

The likelihood of a compromise of CIA associated with social engineering is significantly higher than the likelihood of any resource isolation threat event.

### 4.11.2.PROCESS: PATCH MANAGEMENT: VULNERABILITIES

In 2015, Microsoft announced [225] that the most exploited vulnerability of that year was CVE-2010-2568; as the name suggests, a vulnerability identified (and patched) in 2010. With the proliferation of device types and software components across an enterprise, updating and managing the patch management process can be onerous and poorly implemented.

Criminals of all types are exploiting patch management vulnerabilities to deliver arbitrary exploit payloads into the enterprise. Open source vulnerability scanners and the evolution of automated tools such as Metaspolit [226] are making the process of scanning for and delivering malware straight-forward for all threat actors.

| **Actor** | **History** | **Motivation** | **LoI Score** |
|---|---|---|---|
| Opportunistic | High | High | High |
| Organised Criminal | High | High | High |
| Nation State | High | High | High |

Table 4-30 Likelihood of Initiation: Automated Attack Tooling

Controls in this category are detective in nature. Depending on the budget of an organisation, antimalware tools existing on endpoint. Organisations should maintain asset inventories of all applications, operating systems and middleware components although maintenance of such a list is onerous.

| Actor | Control Strength | Threat Strength | Likelihood of Success |
|---|---|---|---|
| Opportunistic | Moderate | Low | Low |
| Organised Criminal | Moderate | Moderate | Moderate |
| Nation State | Low | High | High |

Table 4-31 Likelihood of Success: Automated Attack Tooling

The likelihood of success for an attacker looking to find and exploit a vulnerability on an endpoint device or server are significantly higher than those associated with resource isolation. Automated tools required to exploit OS and applications are readily available and have been used by opportunistic actors to devastating effect [227].

| Actor | Likelihood of Initiation | Likelihood of Success | Residual Likelihood |
|---|---|---|---|
| Opportunistic | High | Low | Moderate |
| Organised Criminal | High | Moderate | High |
| Nation State | High | High | High |

Table 4-32 Residual Likelihood: Automated Attack Tooling

The vulnerabilities associated with patch management are a problem which looks to increase in severity as our consumption of applications becomes more broad and diverse.

### 4.11.3.TECHNOLOGY: DISTRIBUTED DENIAL OF SERVICE

**Section 4.9.4** discussed the specifics of DoS in a multitenant environment. I assert that Distributed Denial of Service (DDoS) attacks are a serious consideration for any organisation although not exacerbated by cloud. DD0S attacks are growing in terms of impact. Seven of the 10 DDoS attacks greater than 300+ Gbps ever tracked by Akamai occurred in 2016, including three in Q4 [228].

A DDoS attack is launched when an attacker uses the resources of a victim (computer). Last year's attack [210] on the DYN Domain Name Server (DNS) infrastructure reportedly emanated from a myriad of compromised devices. These victims were not computers in a traditional sense; video recorders and CCTV devices were exploited through the Mirai Botnet and programmed to send malicious traffic to Dyn's infrastructure. Dyn estimated [210] that 100,000 devices were compromised and involved in sending traffic to their servers which resulted in sites such as Reddit and Twitter being offline for several hours.

The purpose of analysing DDoS is to understand the barriers to entry and potential likelihood of attack in comparison to exploitation of multitenancy vulnerabilities attributed to public cloud. Unlike public cloud resource isolation vulnerabilities, we have a myriad of attacks which have taken place and for which we can critically analyse.

Cloudflare [229] are an organisation providing Content Delivery Network (CDN) and web optimisation capabilities. In February 2016, Cloudflare reported having to consume DDoS traffic which peaked at over 400 Gbps (gigabits per second) and worked largely through layer 3 (OSI) and volumetric means [230]. The objective being to send more traffic that the target infrastructure can consume and taking a service offline. Invariably, amplification of protocols such as Domain Name System (DNS) and Network Time Protocol (NTP) are used to maximise the volume of traffic which can be sent from a small number of attackers. The point here is that DDoS attacks are on the rise and there are few available controls available to reduce or mitigate the obvious availability concerns. Cloudflare stating: *"L3 attacks are dangerous because most of the time the only solution is to acquire large network capacity and buy beefy networking hardware, which is simply not an option for most independent website operators"* [230].

| **Actor** | **History** | **Motivation** | **LoI Score** |
|---|---|---|---|
| Opportunistic | Moderate | High | High |
| Organised Criminal | Moderate | High | High |
| Nation State | Moderate | High | High |

Table 4-33 Likelihood of Initiation: Distributed Denial of Service

Unlike the threat events we have outlined for public cloud, DDoS attacks can be launched with leased infrastructure and carried out by essentially unskilled attackers. In the world of cyber security, rhetoric suggests we are not dealing anymore with a "kid in their bedroom". As with the threat events associated with patch management, opportunistic criminals can leverage prebuilt "stress testing" environments on the dark web and affect the availability of online services with little technical knowledge or funding.

| **Actor** | **Control Strength** | **Threat Strength** | **Likelihood of Success** |
|---|---|---|---|
| Opportunistic | Moderate | Low | Low |
| Organised Criminal | Moderate | Moderate | Moderate |
| Nation State | Moderate | High | High |

Table 4-34 Likelihood of Success: Distributed Denial of Service

As cited [230], control strength for DDoS prevention is limited. Interestingly, public cloud significantly mitigates the availability risks associated with distributed denial of service.

| Actor | Likelihood of Initiation | Likelihood of Success | Residual Likelihood |
|---|---|---|---|
| Opportunistic | High | Low | Moderate |
| Organised Criminal | High | Moderate | High |
| Nation State | High | High | High |

Table 4-35 Residual Likelihood: Automated Attack Tooling

DDoS attacks are an increasingly prevalent threat to organisations, with actors adopting increasingly sophisticated means of restricting access to resources [228]. DDoS services are available for hire on the dark web making the barriers to entry low. The controls available to organisations are only partially effective: at a network layer, an attack with more bandwidth than a defender is likely to achieve success. An application layer attack relies on the configuration of layer 7 firewalls to mitigate attacks targeted at the application. Unlike the hypervisor-based vulnerabilities outlined in this chapter, DDoS attacks are frequently reported "in the wild" [230, 210].

## 4.12. CONCLUSION

Regardless of NIST's classification [41], my research has led me to believe that multitenancy is a fundamental component of public cloud computing; as is virtualisation. Gartner [144] acknowledge that as soon as tenants are shared at any level from Infrastructure through to application, virtualisation must be utilised to leverage cost and performance benefits for both the consumer and the CSP.

Across this chapter, I have decomposed the components of multitenancy and draw the conclusion that threats and vulnerabilities to resource isolation exist in all forms of multitenancy (IaaS, PaaS and SaaS). Of the vulnerabilities identified, a significant number of these are a result of information leakage which could be prevented through strong application security lifecycle processes.

Two discoveries stood out particularly to the author. Firstly, the ubiquitous reliance on virtualisation and multitenancy for the delivery of contemporary IT solutions and secondly, the sparsity of publicly exploited hypervisor vulnerabilities or inter-VM attacks. There is a disproportionately large amount of focus on hypervisor and virtualisation vulnerabilities across our industry. It is accurate to identify the fact that vulnerabilities do exist in virtualisation technologies; however, I believe that the effort necessary to successfully exploit these vulnerabilities are simply not worth the effort. I initially assumed that many of the academic papers in the field of VM and hypervisor vulnerabilities [154, 18] predate contemporary controls, such as AWS Security Groups, which render attacks ineffective or impossible; except for [202], this was not the case.

Through my research, it has become clear that whilst our ability to innovate isn't in question, our ability to secure existing solutions remains ineffective. Breaches are being associated with cloud computing [216] where cloud is the deployment model, although not the source of the vulnerability. This is analogous to circumstantial evidence. Cloud adoption is growing and with global organisations adopting cloud-first strategies [231] it is imperative that our risk decisions are based on quantitative and qualitative assessment as opposed to visceral and poorly-founded assumptions based on a desire for physical control.

In the efforts of impartiality, I have applied the same risk assessment processes to a common vulnerability and threat event with the likelihood of a successful attack being substantially higher in each case and for each category of actor.

# 5. RECOMMENDATIONS: APPLYING RISK MANAGEMENT FOR PUBLIC CLOUD

There are risks which are associated with cloud computing. In this thesis, I have identified that cloud computing requires a shift in mind-set as to the way we provide security controls but also a pragmatic approach to the impact and likelihood of a threat event. In the world of cloud, controls need to be a combination of people, process and technology and appropriate for the threats and vulnerabilities which manifest themselves in any given situation.

I have several recommendations for organisations in their inevitable adoption of public cloud. There is nothing revolutionary in my findings but I believe that the impact of poor security hygiene can be exacerbated when dealing with public cloud services. My findings are covered in detail across the following subsections although, in summary:

1. Risk Management is a holistic set of processes. Risk is contextual to the environment and requires an understanding of threat actors, vulnerabilities and events. Details are provided in **Section 5.1**.

2. For organisations to achieve “security conservation”, a comprehensive understanding of data flows is required. These flows need to be supported by an understanding of the logical and physical security controls availability from the selected CSP. Details provided in **Section 5.2**.

## 5.1. CLOUD RISK METAMODEL

Having analysed the threat actors, events and vulnerabilities with public cloud computing, I wanted to assimilate information and provide a single blueprint for cloud risk analysis. Only through a thorough assessment of the components involved, can organisations make informed, pragmatic risk decisions.

My public cloud risk metamodel includes the actors, events, vulnerabilities and processes which should be considered when reviewing a public cloud architecture. The model allows security professionals to suitably assess the in-scope actors, events and vulnerabilities for an environment. In **Section 2.8**, I explain that multiple reference architectures should be reviewed to provide comprehensive risk coverage and requirements traceability. Whilst I do not expect this artefact to replace phases of a risk management lifecycle, it should be used in a risk planning phase to ensure that controls and impact are considered rather than simply untreated or inherent risk.

As new threat events are discovered, these should be added. Risk management must be an iterative activity and consequentially, the artefacts and frameworks adopted will evolve.

I assert that more time should be spent understanding the attributes of the threat actor and the controls available to mitigate vulnerabilities. An abundance of literature [18, 126, 179, 154] covers the vulnerabilities (or “risks”) of cloud, virtualisation and multitenancy; far less academic research exists which details mitigations and countermeasures.

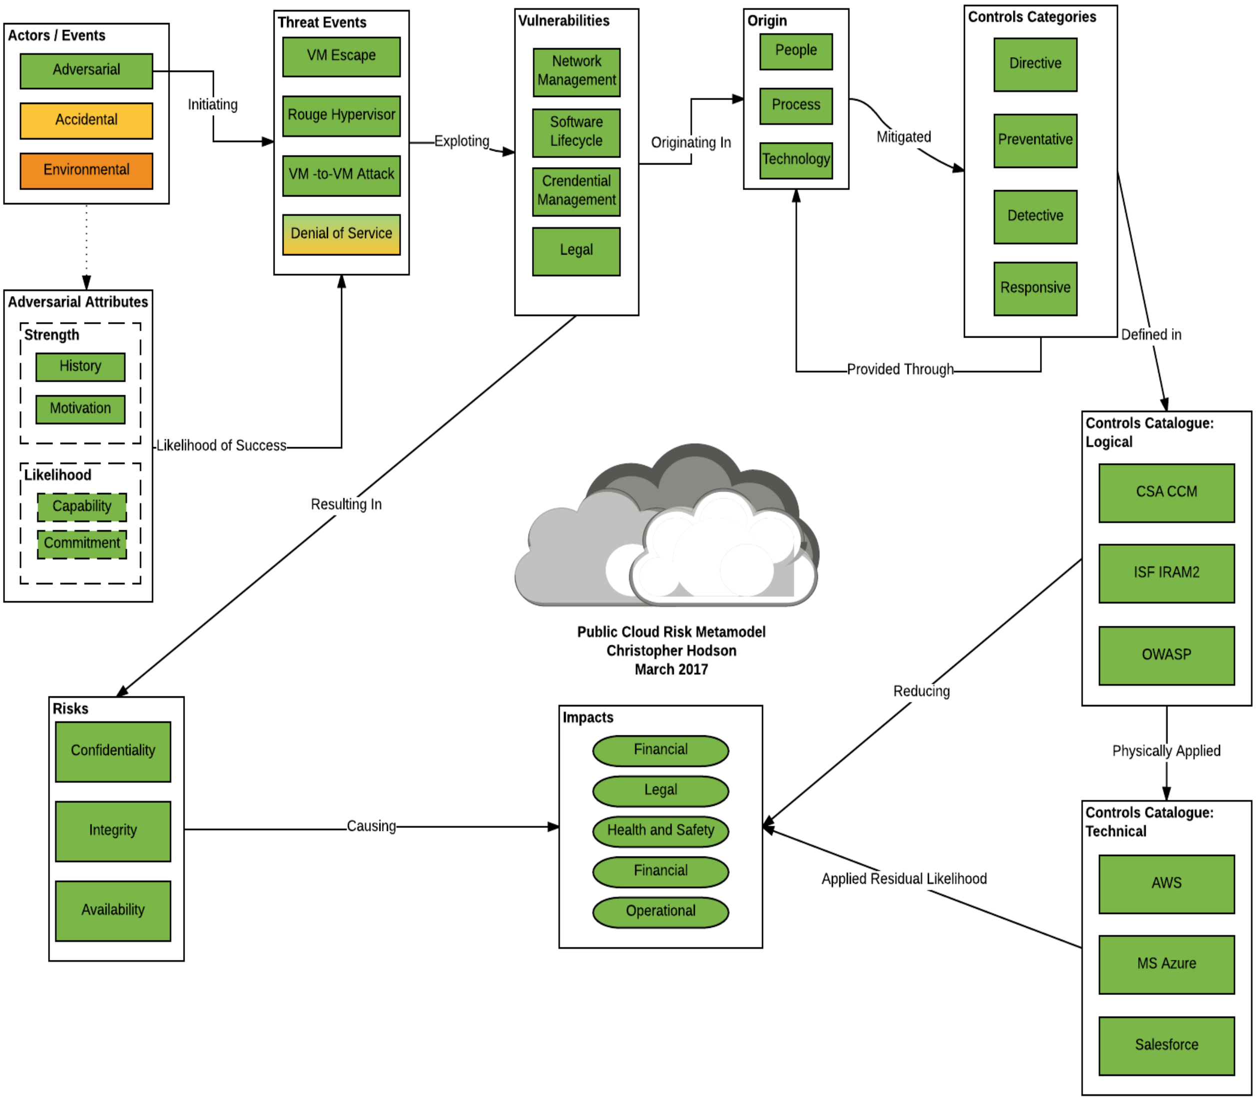

Table 5-1 Public Cloud Risk Metamodel

The metamodel helps visualise the intrinsic relationship between risks and impacts. I conclude that a better understanding of this relationship will assist organisations understand risk context and apply an appropriate information classification framework. Calculating impact should be a business-driven activity. Security teams are responsible for protecting the confidentiality, integrity and availability of information; it is the risks to these three tenants which are the very core of information security irrespective of data location (cloud or on-premise).

I produced the model with a view that similar activity had not been undertaken for public cloud deployments. I still believe this to be the case although I subsequently identified a similar logical model for threat actors and their associations with vulnerabilities and countermeasures [180]. I feel that this validates the importance of such models and I hope that the additions I have made help the reader to better contextualise any unique public cloud components. The most interesting discovery I made in constructing the model was the importance of ensuring that an organisation's security controls need to be a combination of people, process and technology If all avenues of risk are to be treated.

## 5.2. SECURITY CONSERVATION AND HOLISTIC SECURITY ASSURANCE

NIST [83] defines the principle of "security conservation". This principle identifies that as a service is migrated to the cloud, it should retain the security capabilities and services that were applied in its previous location. The responsibility for the application of these controls may change (from customer to CSP); however, it is critically important that the controls are still applied.

Returning to **Chapter 3,** cloud security often requires a shift in the methods and mechanisms for validating the security posture of a solution. Legacy computing models have relied largely on logical and physical security controls deployed on-premise by the organisation's staff. Security compliance has been provided through programmes and testing exercises which were commissioned by employees of the organisation. Often, this process of point-in-time, offensive testing is not practical in a multitenant environment as it would affect the operations of other customers.

Organisations should not be blindly trusting cloud providers to provide secure applications and infrastructure. What is needed is a translation of the assurance that was previously provided on-premise, in the cloud. Quite often this requires the engagement of stakeholders outside of IT, something that has not previously been required. It is important to establish roles and responsibilities between the customer and the cloud service provider. Amazon have taken the need for role delineation very seriously and constructed a "Shared Responsibilities Model" [138]. Figure 5-1 outlines the responsibilities of the customer and those of the CSP.

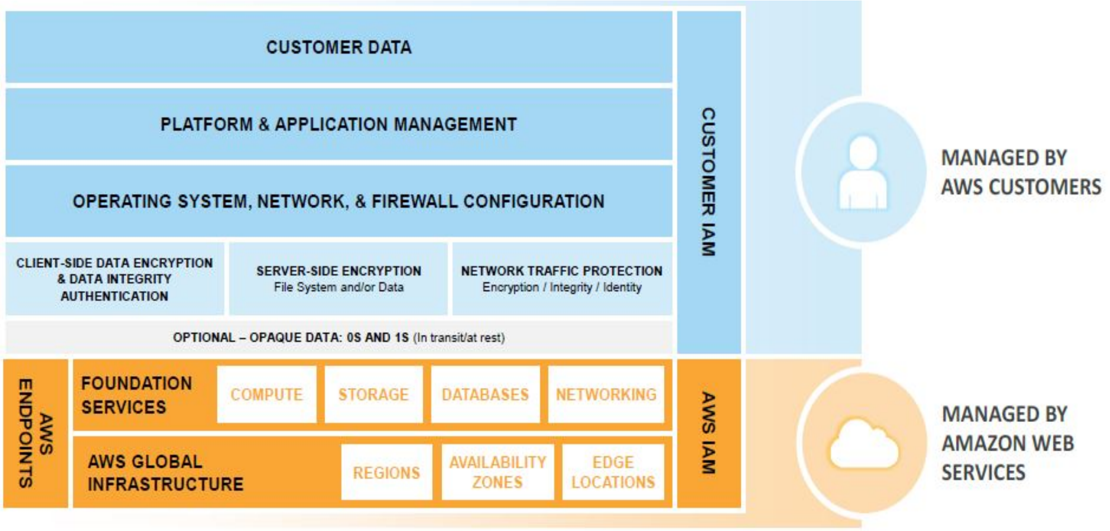
CUSTOMER DATA
PLATFORM & APPLICATION MANAGEMENT
OPERATING SYSTEM, NETWORK, & FIREWALL CONFIGURATION
CLIENT-SIDE DATA ENCRYPTION & DATA INTEGRITY AUTHENTICATION
SERVER-SIDE ENCRYPTION File System and/or Data
NETWORK TRAFFIC PROTECTION Encryption / Integrity / Identity
OPTIONAL – OPAQUE DATA: 0S AND 1S (In transit/at rest)
CUSTOMER IAM
MANAGED BY AWS CUSTOMERS
AWS ENDPOINTS
FOUNDATION SERVICES
COMPUTE
STORAGE
DATABASES
NETWORKING
AWS GLOBAL INFRASTRUCTURE
REGIONS
AVAILABILITY ZONES
EDGE LOCATIONS
AWS IAM
MANAGED BY AMAZON WEB SERVICES

Figure 5-1 AWS Shared Responsibilities Model for EC2 [115]

I have included explicit reference to this model as a thorough understand on of the concepts is necessary if organisations are going to achieve security conservation. Failure to implement controls at each of the layers outlined by Amazon will introduce avoidable vulnerabilities. Much like **Figure 2-7,** responsibilities between CSP and customer vary depending on service model; as organisations move from IaaS through to SaaS, the responsibilities on Amazon increase for areas such as platform management and server-side encryption. What is important to recognise is that AWS provides guidance at a component level as to what is the responsibility of the customer and where Amazon are responsible; Amazon are responsible for the security of the cloud, you (customer) are responsible for security in the cloud – <u>you cannot outsource accountability</u>.

An architectural framework worthy of further reference is the Sherwood Applied Business Security Architecture (SABSA) [82]. SABSA allows an enterprise to define security capabilities from the view of the business. SABSA is included as a core framework under the CSA Enterprise Architecture [84] and **Figure 5-2** details a logical set of architectural building blocks applicable to a public cloud deployment. I would recommend anyone considering a public cloud deployment to use this process in understanding who (CSP or Customer) is responsible for specific security capabilities. Having analysed the services outlined in the SABSA model, almost all these capabilities remain the responsibility of the customer. I was encouraged by this finding as it supports my assertion that operational security will remain the responsibility of the customer in a public cloud deployment. The alignment of capabilities to controls will be discussed in **Section 5.2.2**.

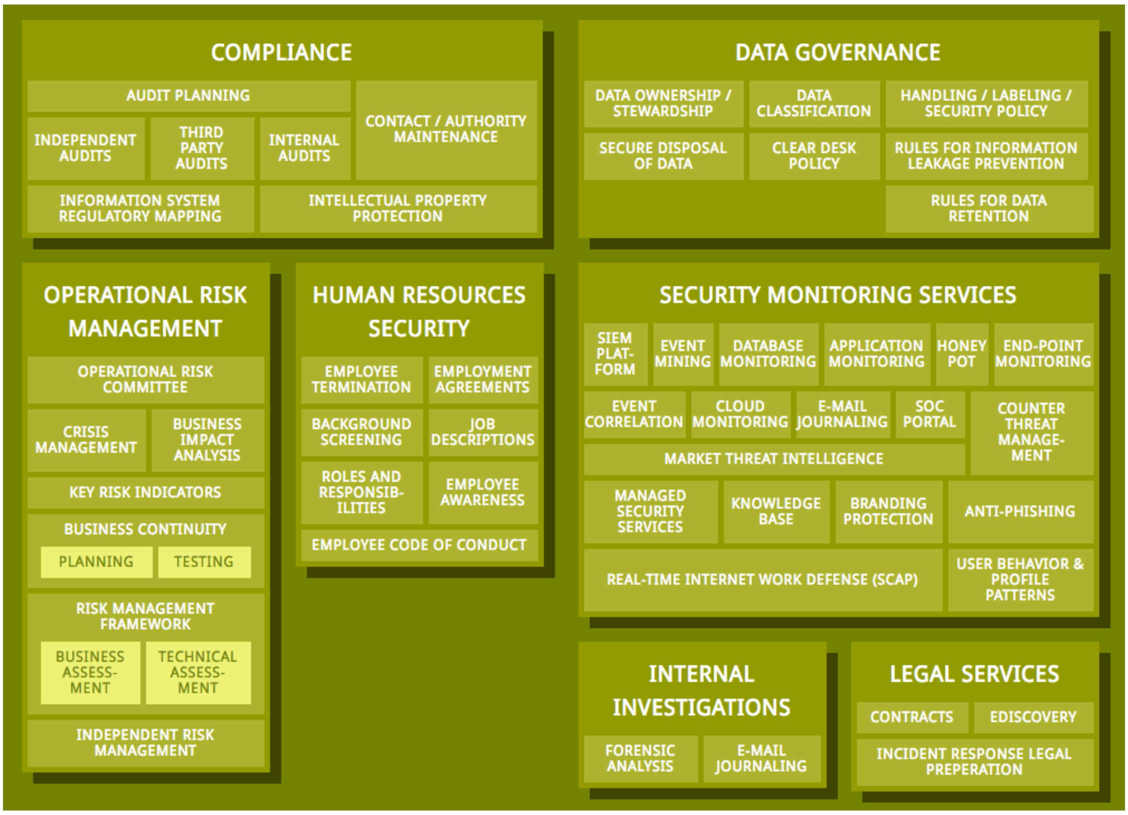

Figure 5-2 SABSA Application to Cloud Services [84]

In summary, I believe that security conservation brings confidence and assurance to the enterprise. A two-phased process is recommended to ensure that cloud services are adopted securely through the retention of security conservation (**Figure 5-3**):

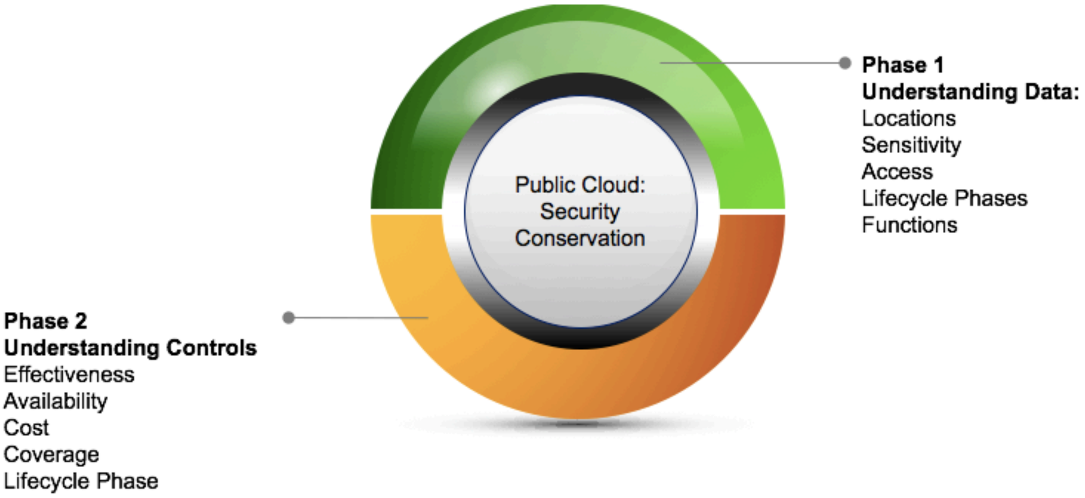

Figure 5-3 Public Cloud: Security Conservation

### 5.2.1. PHASE 1: UNDERSTANDING DATA

Organisations cannot protect what they do not know about; we need to understand where our data is, always. As discussed in **Section 3.5.5,** understanding where enterprise data resides, and the path it takes between systems, is imperative to comply with legal and regulatory requirements but also to ensure that appropriate controls are being deployed.

Across my research, I discovered a data lifecycle model which was straight-forward and non-technical. Two attributes necessary if productive business stakeholder engagement is to be achieved. The "*Data Security Lifecycle*" (DSL) was created by Securosis and version 2.0 is available online [232]. The DSL is comprised of six phases which cover data from inception to destruction.

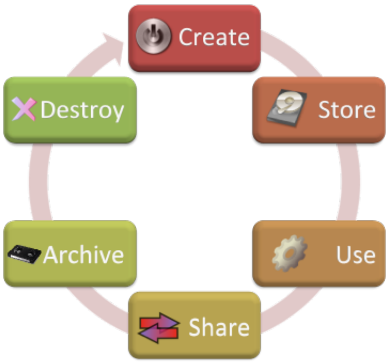

Figure 5-4 Data Security Lifecycle [232]

Considering the explosion of public cloud adoption, the DSL has been updated to include "Locations" and "Access". These additions are a critical component to aid organisations in cloud adoption journey.

The addition of locations was made as organisations are commonly leveraging cloud services for backup, innovation, DR and testing facilities. The original model assumed an on-premise architecture in all scenarios. Any organisation can benefit from asking the questions which are outlined in the revised DSL [232]:

1. Where are the potential locations for my data?
2. What are the lifecycles (Figure 5-4) and controls in each of those locations?

3. Where in each lifecycle can data move between locations?
4. How does data move between locations (via what channel)?

Once data residency has been established, it is important the we understand who is accessing the data and from where. Completing this activity successfully will satisfy many of legal and regulatory requirements which are frequently cited at as being an inhibitor to public cloud adoption. Two important questions to ask are:

1. Who accesses the data?
2. How can the data be accessed (device and channel)?

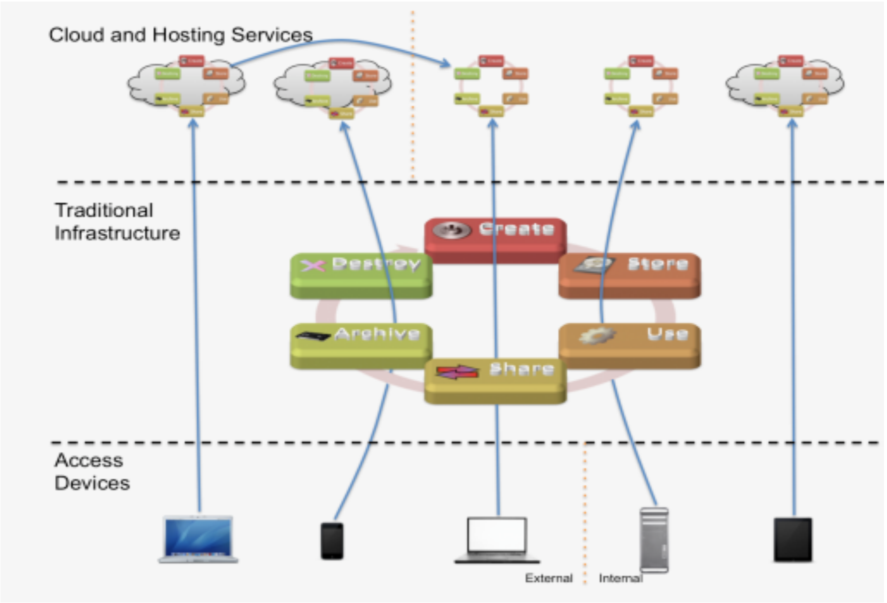

Figure 5-5 Data Lifecycle in the Cloud [232, p. 2]

The challenge of data governance for public cloud adoption is visually represented **Figure 5-5**. As an on-premise architecture, there was generally a single data lifecycle process; a means of creating, storing and sharing information. With public cloud, each CSP has their own unique, sometimes proprietary means of performing similar tasks. Management of data can become exponentially more complicated.

The Securosis model concludes by suggesting that a series of functions are necessary to provide the necessary visibility of who is accessing what, from where. Functions on data take three forms [232] **Figure 5-6**:

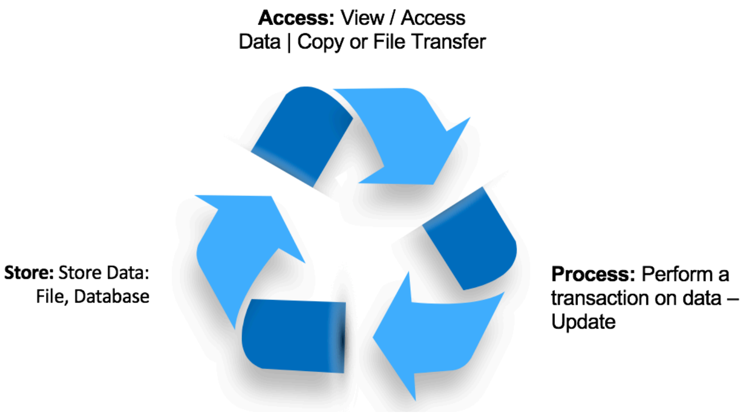
**Access:** View / Access Data | Copy or File Transfer
**Store:** Store Data: File, Database
**Process:** Perform a transaction on data – Update

Figure 5-6 Data Functions

If organisations document who has access to what data, at each public cloud provider, they're able to achieve phase 1 of security conservation: knowing where your data resides.

I recommend that the activity outlined in this section should be completed at the time at which organisations complete a Business Impact Assessment (BIA). A BIA framework is common component of an organisation's business continuity planning activity. The BIA defines the Recovery Time Objective (RTO) (maximum amount time an organisation can be without their information) and Recovery Point Objective (RPO) (maximum amount of data a company can afford to lose). This activity further identifies that public cloud adoption cannot be treated as an "IT Problem". It is critical that we engagement business owners who understand the sensitivity requirements of their information prior to approaching the application of controls.

General Data Protection Regulation (GDPR) becomes law in 2018 and contains clauses which are particularly relevant for public cloud although these too would fall into our "exacerbated by cloud" category identified in **Section 3.6**. Organisations need to understand the type, volume and location of all personally identifiable information. Whilst explicit in GDPR, we have seen in this section that good security practices would see this activity as recommended across any information security programme.

### 5.2.2. PHASE 2: UNDERSTANDING CONTROLS

Understanding the available controls can be provided through the application of a security capabilities model. As we explored in **Chapter 3**, good risk management assesses the maturity, coverage and strength of security controls available for the protection of information. Earlier in this section, I introduced SABSA's approach to identify logical controls. I assert that understanding controls from a technology agnostic perspective ensures architectural traceability and removes the propensity to assume a technology is required without understanding why. This way of working is colloquially referred to as "bottom up" – the loose-coupling of technology without necessarily possessing an understanding of the risks (being mitigated).

Pre-cloud, the deployment and procurement of technology systems was almost exclusively the reserve of the IT department. Systems and infrastructure couldn't be brought online without engineers "racking and stacking" servers, while the desktop technology teams deployed applications to user workstations. With public cloud, business teams can deploy production systems in public cloud almost instantaneously. Whilst this is not a vulnerability introduced by cloud, it is certainly exacerbated and made considerably easier by cloud. The same agility and flexibility is now available in public cloud for the application of security services.

Whilst many academic papers exist which identify the risks of public cloud [155, 154], I assert that these predate the advancements that have been made by mature CSPs to incorporate security into their core platform offerings. I put this down to two organisational trends:

1. Security was not a board-room agenda item. C-level executives were largely disinterested in the cyber posture of their organisations.
2. Security assurance was not a top demand of the consumer.

As the threat landscape has evolved and data breaches are front-page headlines daily, the demand for cyber security is pervasive in the enterprise and the consumer market. This inertia has incentivised CSPs to prioritise the security of their clouds but also the availability of controls for customers within the cloud. Amazon Marketplace [56] provides a holistic portfolio of virtual security appliances for the enterprise. In fact, "security" is a homepage category returning over 500 different solutions (**Figure 5-7**).

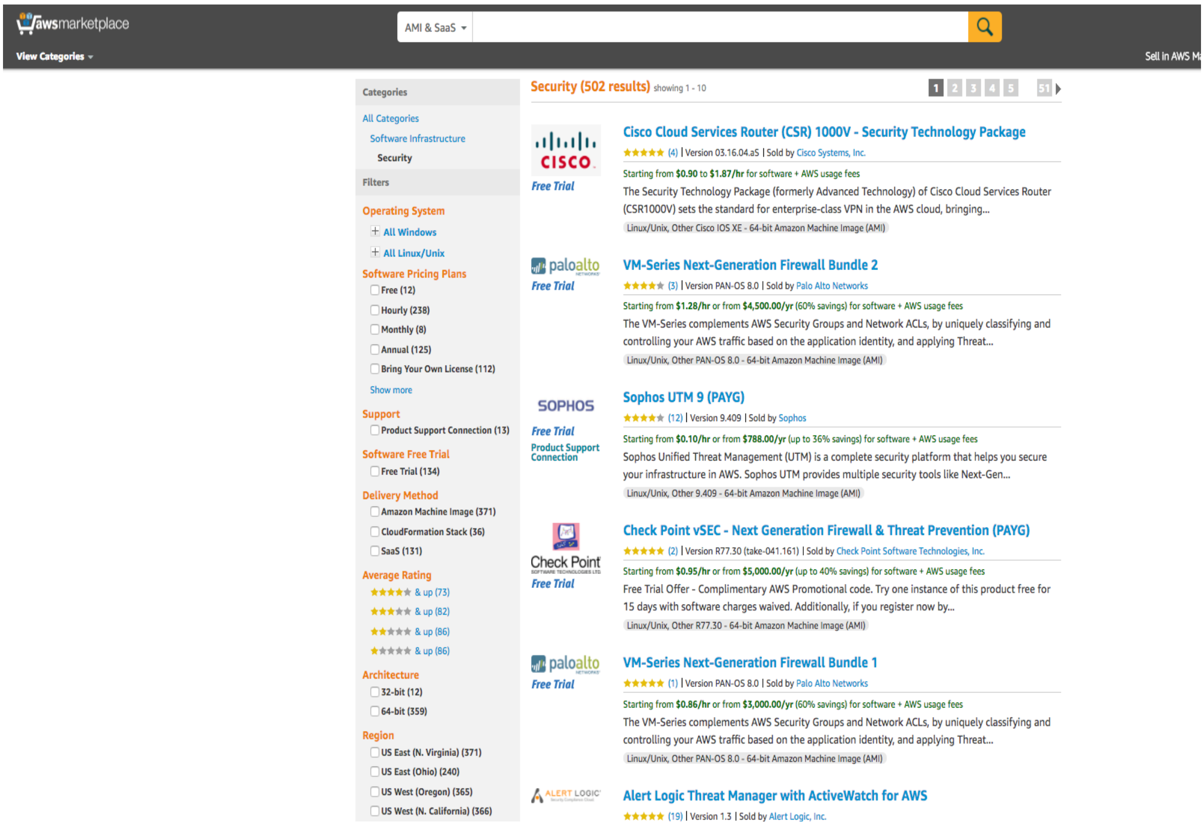

Categories

All Categories

Software Infrastructure

Security

Filters

Operating System

- All Windows
- All Linux/Unix

Software Pricing Plans

- □ Free (12)
- □ Hourly (238)
- □ Monthly (8)
- □ Annual (125)
- □ Bring Your Own License (112)

Show more

Support

- □ Product Support Connection (13)

Software Free Trial

- □ Free Trial (134)

Delivery Method

- □ Amazon Machine Image (371)
- □ CloudFormation Stack (36)
- □ SaaS (131)

Average Rating

- ★★★★☆ & up (73)
- ★★★☆☆ & up (82)
- ★★☆☆☆ & up (86)
- ★☆☆☆☆ & up (86)

Architecture

- □ 32-bit (12)
- □ 64-bit (359)

Region

- □ US East (N. Virginia) (371)
- □ US East (Ohio) (240)
- □ US West (Oregon) (365)
- □ US West (N. California) (366)

Security (502 results) showing 1 - 10 1 2 3 4 5 51

Free Trial

Cisco Cloud Services Router (CSR) 1000V - Security Technology Package

★★★★★ (4) | Version 03.16.04.aS | Sold by Cisco Systems, Inc.

Starting from $0.90 to $1.87/hr for software + AWS usage fees

The Security Technology Package (formerly Advanced Technology) of Cisco Cloud Services Router (CSR1000V) sets the standard for enterprise-class VPN in the AWS cloud, bringing...

Linux/Unix, Other Cisco IOS XE - 64-bit Amazon Machine Image (AMI)

Free Trial

VM-Series Next-Generation Firewall Bundle 2

★★★★☆ (3) | Version PAN-OS 8.0 | Sold by Palo Alto Networks

Starting from $1.28/hr or from $4,500.00/yr (60% savings) for software + AWS usage fees

The VM-Series complements AWS Security Groups and Network ACLs, by uniquely classifying and controlling your AWS traffic based on the application identity, and applying Threat...

Linux/Unix, Other PAN-OS 8.0 - 64-bit Amazon Machine Image (AMI)

Free Trial

Product Support Connection

Sophos UTM 9 (PAYG)

★★★★☆ (12) | Version 9.409 | Sold by Sophos

Starting from $0.10/hr or from $788.00/yr (up to 36% savings) for software + AWS usage fees

Sophos Unified Threat Management (UTM) is a complete security platform that helps you secure your infrastructure in AWS. Sophos UTM provides multiple security tools like Next-Gen...

Linux/Unix, Other 9.409 - 64-bit Amazon Machine Image (AMI)

Free Trial

Check Point vSEC - Next Generation Firewall & Threat Prevention (PAYG)

★★★★★ (2) | Version R77.30 (take-041.161) | Sold by Check Point Software Technologies, Inc.

Starting from $0.95/hr or from $5,000.00/yr (up to 40% savings) for software + AWS usage fees

Free Trial Offer - Complimentary AWS Promotional code. Try one instance of this product free for 15 days with software charges waived. Additionally, if you register now by...

Linux/Unix, Other R77.30 - 64-bit Amazon Machine Image (AMI)

Free Trial

VM-Series Next-Generation Firewall Bundle 1

★★★★★ (1) | Version PAN-OS 8.0 | Sold by Palo Alto Networks

Starting from $0.86/hr or from $3,000.00/yr (60% savings) for software + AWS usage fees

The VM-Series complements AWS Security Groups and Network ACLs, by uniquely classifying and controlling your AWS traffic based on the application identity, and applying Threat...

Linux/Unix, Other PAN-OS 8.0 - 64-bit Amazon Machine Image (AMI)

Alert Logic Threat Manager with ActiveWatch for AWS

★★★★★ (19) | Version 1.3 | Sold by Alert Logic, Inc.

Figure 5-7 Amazon Marketplace: Security Solutions [56]

I wanted to understand if capability gaps existed between the available solutions for an entirely on-premise architecture and one residing exclusively in public cloud. I have identified "security conservation" as critically important in public cloud adoption. If it is not possible to apply commensurate security to your public cloud instance (as with your on-premise solution) then a weakening of the organisation's security posture is taking place. Is it possible to take the requirements of **Figure 5-2** and evidence people, process and technology controls for public cloud?

To delivery this analysis, I will return to concepts introduced in **Chapter 2**, where I detailed the OSA's Cloud Security Pattern [68], a technology agnostic pattern which identifies the actors and controls commonly appropriate for public cloud architectures. After much deliberation, I decided that the OSA model presented a digestible view of my analysis; it is also specifically created for public cloud.

Taking the OSA's pattern, I have overlaid a legend based on Amazon's Shared Responsibilities Model which evidences responsibilities for the controls associated with a generic public cloud deployment. **Figure 5-8** clearly delineates customer and CSP responsibilities. What was interesting in my research was the obvious separation of technology from people and process. Even when an organisation provisions resources within an AWS public cloud, the operational security processes still sit with the customer. AWS provide the technical capabilities for the customer: provisioning, cryptographic key management, logging solutions but the management responsibly still remain with the customer.

Whilst unintentional in its design, **Figure 5-8** shows that the CSP (AWS) is responsible for the foundations of a public cloud environment, a "spine": the technology, the network interconnects and physical security although process and people are still largely the responsibly of the customer.

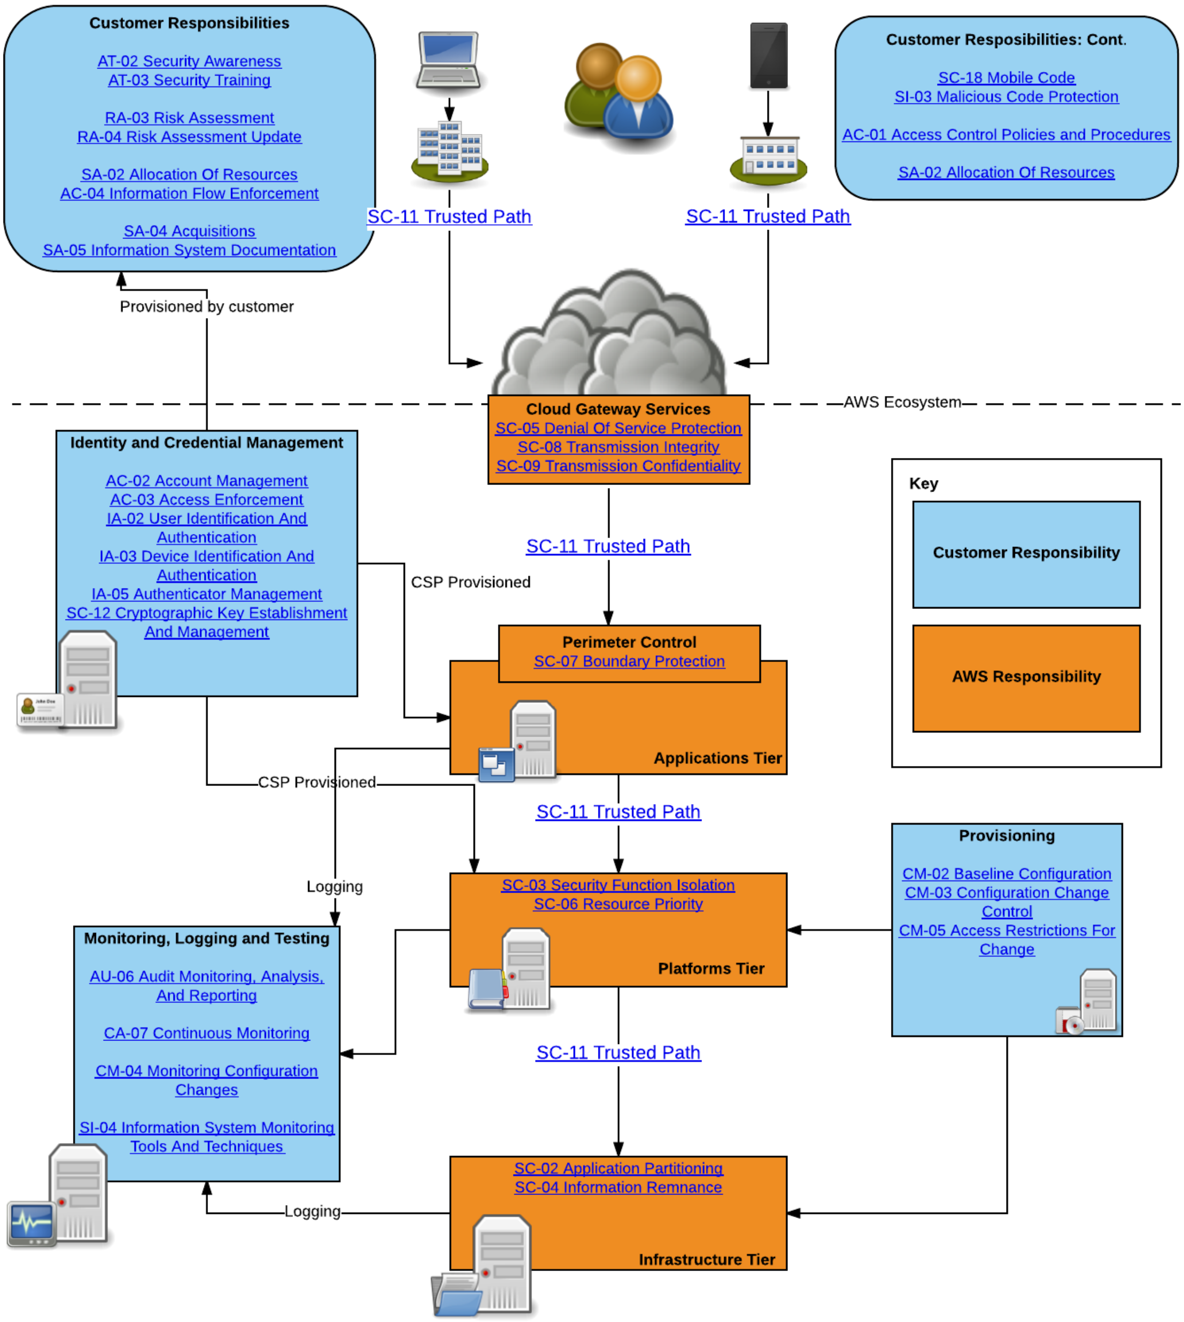
Customer Responsibilities
AT-02 Security Awareness
AT-03 Security Training
RA-03 Risk Assessment
RA-04 Risk Assessment Update
SA-02 Allocation Of Resources
AC-04 Information Flow Enforcement
SA-04 Acquisitions
SA-05 Information System Documentation
SC-11 Trusted Path
SC-11 Trusted Path
Customer Resposibilities: Cont.
SC-18 Mobile Code
SI-03 Malicious Code Protection
AC-01 Access Control Policies and Procedures
SA-02 Allocation Of Resources
Provisioned by customer
AWS Ecosystem
Cloud Gateway Services
SC-05 Denial Of Service Protection
SC-08 Transmission Integrity
SC-09 Transmission Confidentiality
Identity and Credential Management
AC-02 Account Management
AC-03 Access Enforcement
IA-02 User Identification And Authentication
IA-03 Device Identification And Authentication
IA-05 Authenticator Management
SC-12 Cryptographic Key Establishment And Management
SC-11 Trusted Path
CSP Provisioned
Key
Customer Responsibility
AWS Responsibility
Perimeter Control
SC-07 Boundary Protection
Applications Tier
CSP Provisioned
SC-11 Trusted Path
Provisioning
CM-02 Baseline Configuration
CM-03 Configuration Change Control
CM-05 Access Restrictions For Change
Logging
SC-03 Security Function Isolation
SC-06 Resource Priority
Platforms Tier
Monitoring, Logging and Testing
AU-06 Audit Monitoring, Analysis, And Reporting
CA-07 Continuous Monitoring
CM-04 Monitoring Configuration Changes
SI-04 Information System Monitoring Tools And Techniques
SC-11 Trusted Path
SC-02 Application Partitioning
SC-04 Information Remnance
Logging
Infrastructure Tier

**Figure 5-8 Architecture: Customer / CSP Responsibilities**

It is important to identify that **Figure 5-8** details responsibility of services. There are components of the figure where the customer is responsible yet Amazon are providing capability. A good example is "provisioning"; Amazon provides services such as Elastic Beanstalk with which a developer uploads application code and Elastic Beanstalk automatically handles resource provisioning, load balancing, auto scaling and monitoring [115].

Amazon are providing the service although the customer remains responsible for the code and the development lifecycle that goes with it. Similar could, and should, be said for cryptographic service, logging and monitoring.

### 5.2.3. SECURITY CONSERVATION: PROCESS FLOW

To complete my analysis, I propose a series of steps which should be incorporated into any risk management activity. Following these sequential steps will allow an organisation to make balanced information risk decisions and achieve security conservation.

Having drafted a mind-mapping activity (**Figure 5-9**), I have identified that a simple set of steps can be followed:

- ✓ Work with business stakeholders and conduct a BIA.
- ✓ Ensure that data flows are fully understood and documented.
- ✓ Understand the sensitivity of all data elements being stored or processed.
- ✓ Conduct threat modelling activity: understand the actors, threat events and vulnerabilities that are in scope for the environment.
- ✓ Apply controls to mitigate the impact and/or likelihood of a vulnerability being exploited.
- ✓ Ensuring that the customer and the CSP understand their responsibilities for the operation and maintenance of technology services.
- ✓ Validate that the mitigated risk is within risk tolerance thresholds of the organisation.

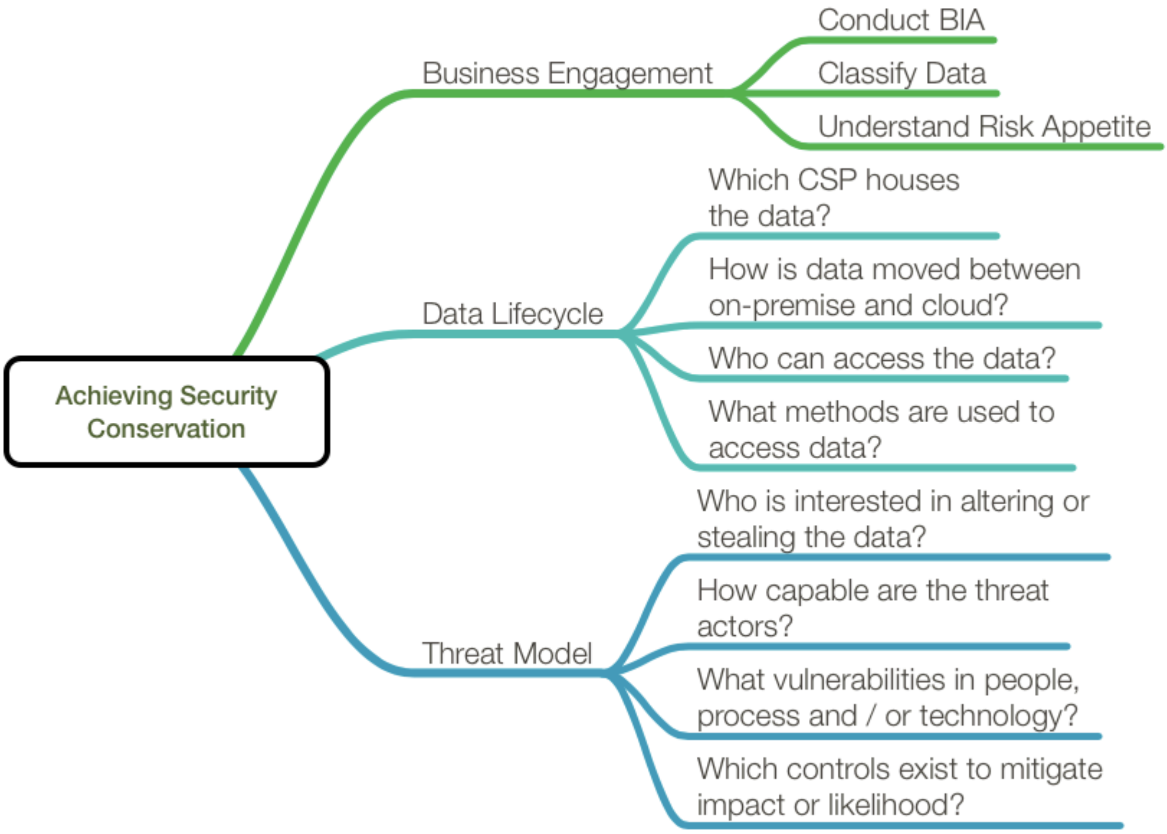
Achieving Security Conservation
Business Engagement
Conduct BIA
Classify Data
Understand Risk Appetite
Data Lifecycle
Which CSP houses the data?
How is data moved between on-premise and cloud?
Who can access the data?
What methods are used to access data?
Threat Model
Who is interested in altering or stealing the data?
How capable are the threat actors?
What vulnerabilities in people, process and / or technology?
Which controls exist to mitigate impact or likelihood?

Figure 5-9 Mind Map: Security Conservation

# 6. CONCLUSIONS: PUBLIC CLOUD IS NOT A TECHNOLOGY PROBLEM

*"Failing to analyze cloud risk will result in missed opportunity and/or unacceptable risk to the business. Risk management is a mature discipline that can determine how much cloud risk is acceptable." [111]*

GARTNER 2015

To conclude this thesis, I will return to the objectives aims in **Section 1.2** and ensure that I can evidence appropriate coverage of the defined goals.

My **first objective** was to understand cloud. Public cloud adoption is growing year-on-year [26, 28] and organisations need to consider how to consume cloud services securely. By way of preliminary assessment, I asserted that public cloud challenges convention. I validated this statement through an analysis of data centre evolution. In 2017, virtualisation is ubiquitous; datacentres, irrespective of their deployment model, are filled with converged infrastructure which takes advantage of multitenant architecture. Due to this converging of technology, I assert that most modern technology paradigms are "cloud like" in nature making it imperative that we take a data-centric and attack based view for information protection: what are we trying to protect and who is trying to compromise the data? I also wanted to understand the impact of public cloud on our traditional network architectures. I conclude that as public cloud adoption grows, the organisational need for a WAN reduces.

As introduced in **Section 2.3**, public cloud computing is bringing security services to customers that they would not be able to afford in a traditional, on-premise, appliance-based model. The breadth of capabilities available within AWS Marketplace [56] ensure that customers can benefit from a service-based, pay-as-you-go model for security. IT services are being deployed via a new paradigm and our businesses need to understand not only the technical IT changes that cloud brings but also the changes to user education and business process.

**Objective two** set out an analysis into the security risks associated with public cloud. It was primarily my own experience in the industry that drove me to set this aim. All too often, terms associated with risk management are used interchangeably despite having specific, and very different, meanings. In **Chapter 3** we introduced the concept of risk. An important discovery for me was that many different forms of risk (financial, information, health and safety) all have common attributes. In **Section 3.1**, I presented the DHS Risk Lexicon [103] and was surprised to see almost all attributes application in an information risk context but also a financial, defence and personal safety lens.

An area I struggled to reach a conclusion on was the legal implications of public cloud computing. The legal and regulatory considerations of public cloud require organisations to fully understand their data flows and the volume / types of information stored in public cloud locations. I have identified in this thesis that such requirements are not only required by regulation but necessary for any organisation serious about protecting the confidentiality, integrity and availability of their sensitive data assets. A vulnerability in people and process *could* result in significant legal and regulatory repercussions but in **Section 2.3.2.4** I present the comprehensive set of regulations that AWS complies with by default. This is further evidence that the technological construction of public cloud can broadly cope with legal and regulatory compliance requirements.

I have reached the conclusion that it is not only the threats, vulnerabilities and controls which are important in assessing risk but also the actors responsible for compromising data; be that intentionally or accidentally. The likelihood of an attack occurring varies significantly depending on many "threat attributes" which I validate in **Section 4.7.1**. In summary, several of the activities required to exploit vulnerabilities in a hypervisor with any meaningful degree of likelihood require a level of investment in time, effort and often funding which precludes an opportunistic attack. The poor returns and high likelihood of being unsuccessful would also deter all but the most targeted of attacks from highly motivated, capable actors. As actor attributes are so important, they form a key component of my "Public Cloud Risk Metamodel" outlined in **Section 5.1**.

In the conclusion to **Chapter 3**, I assert that our understanding of information risk is immature because we apply components of a risk model without a pragmatic or holistic understanding of context. Across my research, the most important takeaway I have gathered is that establishing a structured risk management methodology is critical if we are to understand and contextualise the risks which public cloud computing may present. Approaching "inherent risk" or the worst-case scenario without considering the motivation of the actor or the controls which might be available provides little opportunity for risk prioritisation as everything is considered a risk.

I have identified that technical vulnerabilities do exist within architecture which is intrinsically-linked to public cloud computing. It is however important to consider that the same technological components are present in today's on-premise deployments.

Having analysed industry and academic literature covering the risks and vulnerabilities of public cloud, I have reached the conclusion that the technological construction of public cloud mirrors that of contemporary datacentre architectures irrespective of location. Two significant differences arise with public cloud:

1) The impact of a cyber-attack is exacerbated. Shared infrastructure often means shared impact.

2) Data could be stored in geographical regions with additional legal and data privacy considerations.

My **third objective** was to conduct a risk analysis regarding the technical construction of virtualisation and multitenancy. This step could only be undertaken once the threat actors, events and vulnerabilities had been documented. To satisfy this aim, a thorough decomposition of multitenancy and virtualisation was conducted across **Chapter 4**. I discuss the different forms of multitenancy and contrast implementations across the service models of IaaS, PaaS and SaaS. I discovered that as we moved from IaaS, to PaaS, through to SaaS more of our technology stack was shared tenant. Such a model would infer that there is more risk with a SaaS platform than that of an IaaS implementation although as the CSA asserts [34] (Figure 4-8, Figure 4-10), SaaS deployments leave less responsibility to the customer. Solutions have fewer opportunities for configuration and standardisation lowers the system footprint. I assert in **Section 2.3.2.1** that CSPs must adopt "paranoia by default" [52], this is made significantly easier in situations where the CSP retains control of operational processes.

The risk analysis activity in **Chapter 4** provides the reader with an analysis of the in-scope threat events associated with hypervisor architecture in a shared services model. My hypothesis being that hypervisor architecture presented many unique attack vectors which did not present themselves in a traditional, appliance based model where organisations retained organisationally isolated infrastructure. To satisfy **objective 3**, I could not simply identify that threat events and vulnerabilities existed. I needed to investigate if controls (people, process or technology) were available in the environment to mitigate the vulnerability or deter the threat actor.

To fully fulfil **objective 3,** I reviewed vulnerabilities in people, process and technology which present themselves in modern computing. As attacks are increasingly focusing on users and applications [233], I considered it prudent to focus my people and process examples in this space. My hypothesis being that whilst vulnerabilities in hypervisors cannot be avoided, there are vulnerabilities affecting our users and their data which are more straight-forward to carry out and considerably more likely to succeed.

I have analysed the threat events associated with public cloud and assert that the complexity of achieving co-residence of VMs on a regular basis precludes technical cloud security vulnerabilities warranting further specific attention. For any form of threat actor, there are many easier, cheaper and higher likelihood threat events and vulnerabilities that can be leveraged to compromise the confidentiality, integrity or availability of data.

**Objective 4** was concerned with understanding if existing risk management methodologies suitable cater for public cloud. Having reached the conclusion that location of data should not be the defining factor in the application of security controls, I assert that preeminent risk frameworks are, to a greater extent, applicable across public cloud. In **Chapter 5**, I present a series of recommendations to improve and augment established risk frameworks. Cloud computing does exacerbate the requirement to fully understand data flows; the volumes of data stored in public cloud and associated sensitivity can have an impact on an organisation's regulatory and legal compliance position. In **Chapter 5** I propose a two-stage process for better understanding data flows and assessing the availability of controls in an AWS deployment. By following this approach, organisations can pragmatically ensure that the security controls deployed are commensurate with the sensitivity of the information being stored or processed.

The overarching objective of this thesis was to demystify the risks of public cloud computing. I have complied a "Top Ten Findings" below which summarise my research and provide the reader with important considerations in any public cloud project:

1. Information risk deals with the compromise of the confidentiality, integrity and/or availability of data. Cloud introduces new vulnerabilities which are exploited by existing threat actors and variations of existing threat events.

2. Technical security controls exist across mature CSPs [56] to provide "Security Conservation".

3. Public cloud computing offers organisations a comprehensive suite of logical and physical security controls for the protection of sensitive enterprise data.

4. The technical vulnerabilities associated with public cloud are difficult to exploit.

5. Preeminent regulatory guidelines [140, 141] support the use of, and provide guidance around, the use of public cloud.

6. In 2017, most cyber-attacks are focused at the user or application level [233, p. 2] – a network-centric defence strategy, focusing solely on network vulnerabilities, is doomed to fail in a cloud-first world.

7. In many situations, public cloud adoption can improve an organisation's security posture.

8. Sensitive information is increasingly being stored in public cloud [17].

9. Public cloud consumption is growing rapidly [25, 26], securing public cloud is a discipline the security team needs to become comfortable with.

10. Business processes inside and outside of IT are altered because of public cloud. We are not dealing with a technology problem.

# BIBLIOGRAPHY

[1] M. G. Alex Burmaster, “Britons to spend more on online video than DVDs for the first time,” 6 May 2016. [Online]. Available: https://www.strategyanalytics.com/strategy-analytics/news/strategy-analytics-press-releases/strategy-analytics-press-release/2016/05/06/britons-to-spend-more-on-online-video-than-dvds-for-first-time#.WChHNeGLSV4. [Accessed 13 November 2016].

[2] OFCOM, “The Communications Market Report 2016,” OFCOM, 2016.

[3] K. Leswing, “The 5 most valuable public companies are all tech companies,” 1 August 2016. [Online]. Available: http://uk.businessinsider.com/4-most-valuable-public-companies-all-tech-companies-2016-8. [Accessed 13 November 2016].

[4] L. Chen, “How Uber Surpasses Ford and GM in Valuation in 5 years,” 4 December 2015. [Online]. Available: http://www.forbes.com/sites/liyanchen/2015/12/04/at-68-billion-valuation-uber-will-be-bigger-than-gm-ford-and-honda/#7f698ae15858. [Accessed 13 November 2016].

[5] airbnb, “About Us,” 2016. [Online]. Available: https://www.airbnb.co.uk/about/about-us.

[6] T. Mather, S. Kumaraswamy and S. Latif, Cloud Security and Privacy: An Enterprise Perspective on Risks and Compliance, First ed., vol. 1, Sebastopol, California: O'Reilly, 2009, p. 26.

[7] NASCIO, “State CIO Top Ten Policy and Technology Priorities for 2016,” 2015.

[8] IDG Enterprise, “Cloud Continues to Transform Buisness Landscape as CIOs Explore New Areas for Hosting,” 17 November 2015. [Online]. Available: http://www.idgenterprise.com/news/press-release/cloud-continues-to-transform-business-landscape-as-cios-explore-new-areas-for-hosting/. [Accessed 13 November 2016].

[9] European Network and Information Security Agency, “Cloud Computing: Risk, Benefits and recommendations for information security,” ENIA, 2009.

[10] Gartner, “Gartner Says By 2020, a Corporate "No-Cloud" Policy Will Be as Rare as a "No-Internet" Policy Is Today,” 22 June 2016. [Online]. Available: http://www.gartner.com/newsroom/id/3354117. [Accessed 29 January 2017].

[11] A. Ukil, “A Security Framework in Cloud Computing Infrastructure,” *International Journal of Network Security & Its Applications (IJNSA),* vol. 5, no. 5, pp. 13-22, September 2013.

[12] Gartner, “Gartner Says Worldwide Public Cloud Services Market to Grow 17 Percent in 2016,” Gartner, 15 September 2016. [Online]. Available: http://www.gartner.com/newsroom/id/3443517. [Accessed 2 February 2017].

[13] A. Andress, Surviving security: how to integrate people, process, and technology, vol. 2, CRC Press, 2003.

[14] Gartner, “Gartner Says By 2020, a Corporate "No-Cloud" Policy Will Be as Rare as a "No-Internet" Policy Is Today,” 22 June 2016. [Online]. Available: http://www.gartner.com/newsroom/id/3354117. [Accessed 15 January 2017].

[15] H.-Y. |. H. P.-Y. Lin, “A Framework to Evalutate n-Tier Architecture ERP Systems,” Electronic Commerce Studies, 2005.

[16] Netwrix, “2016 Cloud Security Report,” Netwrix, Irvine, 2017.

[17] Gemalto, “The 2016 Global Cloud Data Security Study (Ponemon Institute LLC),” Gemalto, 2016.

[18] S. Zhang, “Deep-diving into an easily-overlooked threat: Inter-VM attacks,” Kansas State University, Manhattan, 2012.

[19] D. Ruest, “Virtualization hypervisor comparison: Type 1 vs. Type 2 hypervisors,” Search Server Virtualisation, September 2010. [Online]. Available: http://searchservervirtualization.techtarget.com/tip/Virtualization-hypervisor-comparison-Type-1-vs-Type-2-hypervisors. [Accessed 5 February 2017].

[20] Amazon, “Start Building on AWS Today,” Amazon, 2017. [Online]. Available: https://aws.amazon.com/. [Accessed 5 February 2017].

[21] L. Leong, G. Petri, B. Gill and M. Dorosh, “Magic Quadrant for Cloud Infrastructure as a Service, Worldwide,” Gartner, 2016.

[22] The Xen Project, “Xen Project Software Overview,” 21 March 2016. [Online]. Available: https://wiki.xen.org/wiki/Xen_Project_Software_Overview. [Accessed 13 March 2017].

[23] K. Martin, “MSc in Information Security Project Handbook,” Royal Holloway University London, 2015.

[24] McAfee Inc, “Building Trust in a Cloudy Sky: The state of cloud adoption and security,” 2017.

[25] Rightscale, “State of Cloud Report 2017,” 2017. [Online]. Available: https://www.rightscale.com/lp/state-of-the-cloud. [Accessed 3 March 2017].

[26] Salesforce, "Salesforce FY 2016 Annual Report," Salesforce, Delaware, 2017.

[27] The Economist Intelligence Unit, "Ascending cloud The adoption of cloud computing in five industries," The Economist, 2015.

[28] C. Suh, S. Nadella and A. Hood, "Microsoft Earnings Conference Call," 2016.

[29] Amazon Web Services, "Overview of Security Processes," Amazon, 2016.

[30] S. D. Harris, "2009 Q&A: Marc Benioff, CEO of Salesforce.com," The Mercury News, 23 October 2009. [Online]. Available: http://www.mercurynews.com/2009/10/23/2009-qa-marc-benioff-ceo-of-salesforce-com/. [Accessed 13 February 2017].

[31] C. J. Hodson, *Myths of Cloud Computing Debunked,* London: CIO Inspired, 2016.

[32] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," National Institute of Standards and Technology, Gaithersburg, 2011.

[33] Cloud Security Alliance, "About," CSA, [Online]. Available: https://cloudsecurityalliance.org/about/. [Accessed 4 December 2016].

[34] Cloud Security Alliance, "Security Guidance for the Critical Areas of Focus in Cloud Computing V3.0," Cloud Security Alliance, 2011.

[35] ENISA, "Cloud Computing: Benefits, Risks and Recommendations for Information Security," ENISA, 2009.

[36] L. Goasduff, "The Financial Case for Moving to the Cloud," Gartner, 20 August 2015. [Online]. Available: http://www.gartner.com/smarterwithgartner/the-financial-case-for-moving-to-the-cloud/. [Accessed 1 December 2016].

[37] Rightscale , "2016 State of Cloud Report," Rightscale.

[38] J. Mullich, "16 Ways The Cloud Will Change Our Lives," The Wall Street Journal, [Online]. Available: http://online.wsj.com/ad/article/cloudcomputing-changelives. [Accessed 3 February 2017].

[39] J. J. Mark, "Heraclitus of Ephesus," Ancient History Encyclopedia, 14 July 2010. [Online]. Available: https://www.ancient.eu/Heraclitus_of_Ephesos/. [Accessed 22 January 2017].

[40] J. Woods, “The evolution of the data center: Timeline from the Mainframe to the Cloud,” SiliconANGLE, 05 March 2014. [Online]. Available: http://siliconangle.com/blog/2014/03/05/the-evolution-of-the-data-center-timeline-from-the-mainframe-to-the-cloud-tc0114/. [Accessed 22 January 2017].

[41] F. Lui, J. Tong, J. Mao, R. Bohn, J. Messina, L. Badger and D. Leaf, “NIST Cloud Computing Reference Architecture: Recommendations of the National Institute of Standards and Technology,” National Institute of Standards and Technology, 2011.

[42] Gartner, “Gartner Says Smart Organizations Will Embrace Fast and Frequent Project Failure in Their Quest for Agility,” Gartner, 8 May 2013. [Online]. Available: http://www.gartner.com/newsroom/id/2477816.

[43] Amazon, “Amazon Homepage,” Amazon, 16 March 2017. [Online]. Available: https://www.amazon.com. [Accessed 16 March 2017].

[44] Spotify, “Spotify Homepage,” 16 March 2017. [Online]. Available: https://www.spotify.com. [Accessed 16 March 2017].

[45] Netflix, “Netflix Homepage,” 16 March 2017. [Online]. Available: https://www.netflix.com. [Accessed 16 March 2017].

[46] M. E. Porter and J. E. Heppelmann, “How Smart, Connected Products Are Transforming Competition,” November 2014. [Online]. Available: https://hbr.org/2014/11/how-smart-connected-products-are-transforming-competition. [Accessed 3 March 2017].

[47] Cloud Standards Customer Council, “Impact of Cloud Computing on Healthcare,” 2017.

[48] Supply Chain 247, “How the Internet of Things Is Improving Transportation and Logistics,” 9 September 2015. [Online]. Available: http://www.supplychain247.com/article/how_the_internet_of_things_is_improving_transportation_and_logistics. [Accessed 3 March 2017].

[49] D. P. Thomond, “THE ENABLING TECHNOLOGIES OF A LOW-CARBON ECONOMY: A Focus on Cloud Computing,” Enabling Technologies 2020, 2013.

[50] V. Kundra, “Tight Budget? Look to the "Cloud",” The New York Times, 30 August 2011. [Online]. Available: http://www.nytimes.com/2011/08/31/opinion/tight-budget-look-to-the-cloud.html. [Accessed 21 March 2017].

[51] Salesforce, “Business Benefits of Cloud,” [Online]. Available: https://www.salesforce.com/uk/blog/2015/11/why-move-to-the-cloud-10-benefits-of-cloud-computing.html.

[52] D. Linthicum, “The public cloud is more secure than your data center,” IDG, 1 December 2015. [Online]. Available: http://www.infoworld.com/article/3010006/data-security/sorry-it-the-public-cloud-is-more-secure-than-your-data-center.html. [Accessed 18 February 2017].

[53] L. Rittenburg and F. Martens, “Understanding and Communicating Risk Appetite,” Committee of Sponsoring Organisation, 2012.

[54] J. Huang and D. M. Nicol, “Trust mechanisms for cloud computing,” *Journal of Cloud Computing: Advances, Systems and Applications,* vol. 2, no. 9, 2013.

[55] Marketwatch, “Two months after damaging data breach, Target stock has its best day in 5 years,” 26 February 2014. [Online]. Available: http://blogs.marketwatch.com/behindthestorefront/2014/02/26/two-months-after-damaging-data-breach-target-stock-has-its-best-day-in-5-years/. [Accessed 16 March 2017].

[56] Amazon Web Services, “Security,” Amazon, 23 February 2017. [Online]. Available: https://aws.amazon.com/marketplace/b/2649363011?ref_=hmpg_categories_2649363011. [Accessed 23 February 2017].

[57] J. McLaughlin, “Spy Chief Complains That Edward Snowden Sped Up Spread of Encryption by 7 Years,” The Intercept, 25 April 2016. [Online]. Available: https://theintercept.com/2016/04/25/spy-chief-complains-that-edward-snowden-sped-up-spread-of-encryption-by-7-years/. [Accessed 4 March 2017].

[58] Network Working Group, “The Transport Layer Security (TLS) Protocol Version 1.2,” IETF: RFC 5246, August 2008. [Online]. Available: https://tools.ietf.org/html/rfc5246. [Accessed 6 March 2017].

[59] Sandvine, “2016 Global Internet Phenomena - Spotlight: Encrypted Internet Traffic,” Sandvine, 2016.

[60] K. Finley, “Half the Web Is Now Encrypted. That Makes Everyone Safer,” Wired, 30 January 2017. [Online]. Available: https://www.wired.com/2017/01/half-web-now-encrypted-makes-everyone-safer/. [Accessed 4 March 2017].

[61] D. Gooley, “The Rise in SSL-based Threats,” Zscaler, 16 February 2017. [Online]. Available: https://www.zscaler.com/blogs/research/rise-ssl-based-threats-1. [Accessed 4 March 2017].

[62] Zscaler, “Eliminate the appliance mess, for better security at a lower cost,” 2017. [Online]. Available: https://www.zscaler.com/why-zscaler/simplify-it. [Accessed 18 February 2017].

[63] NSS Labs, “NSS Labs Predicts 75% of Web Traffic Will Be Encrypted by 2019,” NSS Labs, 9 November 2016. [Online]. Available: https://www.nsslabs.com/company/news/press-releases/nss-labs-predicts-75-of-web-traffic-will-be-encrypted-by-2019/. [Accessed 4 March 2017].

[64] P. Belcher, “Hash Factory: New Cerber Ransomware Morphs Every 15 Seconds,” 2 June 2016. [Online]. Available: https://www.invincea.com/2016/06/hash-factory-new-cerber-ransomware-morphs-every-15-seconds/. [Accessed 18 February 2017].

[65] A. Venkatraman, “Most cloud services pose security and compliance risks to European businesses,” Computer Weekly, 14 April 2014. [Online]. Available: http://www.computerweekly.com/news/2240218798/Most-cloud-services-pose-security-and-compliance-risks-to-European-businesses. [Accessed 3 March 2017].

[66] Amazon Web Services, “AWS Cloud Compliance,” 2017. [Online]. Available: https://aws.amazon.com/compliance/. [Accessed 3 March 2017].

[67] Gartner Research, “Cloud Access Security Brokers (CASBs),” Gartner, 2017. [Online]. Available: http://www.gartner.com/it-glossary/cloud-access-security-brokers-casbs/. [Accessed 13 March 2017].

[68] Open Security Architecture, “Cloud Security Pattern,” [Online]. Available: http://www.opensecurityarchitecture.org/cms/library/patternlandscape/251-pattern-cloud-computing . [Accessed 27 December 2016].

[69] Gartner, “Gartner Says 6.4 Billion Connected "Things" Will Be in Use in 2016, Up 30 Percent From 2015,” 10 November 2015. [Online]. Available: http://www.gartner.com/newsroom/id/3165317. [Accessed 4 December 2016].

[70] M. B. Barcena and C. Wueest, “Security Response: Insecurity in the Internet of Things,” Symantec, 2015.

[71] D. Kaminsky, “Keynote: The Hidden Architecture of our Time: Why This Internet Worked, How We Could Lose It...,” in *Blackhat / UDM Media*, Las Vegas, 2016.

[72] Amazon Web Services, “AWS Cloud Best Practices,” Amazon, 2011.

[73] N. R. Herbst, S. Kounev and R. Reussner, “Elasticity in Cloud Computing: What It Is, and What It Is Not,” in *10th International Conference on Autonomic Computing*, San Jose, 2013.

[74] Microsoft, “Get the most secure Office for your business,” Microsoft, 2017. [Online]. Available: https://products.office.com/en-gb/business/get-latest-office-365-for-your-business-with-2016-apps?&wt.srch=1&wt.mc_id=AID522516_SEM_6Jp48WIm. [Accessed 4 March 2017].

[75] Salesforce, “Salesforce Homepage,” 2017. [Online]. Available: https://www.salesforce.com/uk/?ir=1. [Accessed 4 March 2017].

[76] Chef, “Achieve speed, scale, and consistency by automating your infrastructure with Chef,” 2017. [Online]. Available: https://www.chef.io/chef/. [Accessed 3 March 2017].

[77] Puppet, “Solutions: Automate DevOps, security & cloud practices,” Puppet, 2017. [Online]. Available: https://puppet.com/solutions. [Accessed 3 March 2017].

[78] Microsoft Azure, “What is PaaS,” Microsoft, 2017. [Online]. Available: https://azure.microsoft.com/en-gb/overview/what-is-paas/?&wt.mc_id=AID529440_SEM_. [Accessed 3 March 2017].

[79] D. Clark, “Commodity Computing,” 31 March 2003. [Online]. Available: http://www.vmware.com/company/news/articles/wsj_2.html. [Accessed 16 March 2017].

[80] W. Stallings, Foundations of Modern Networking: SDN, NFV, QoE, IoT, and Cloud, 1 ed., vol. 1, Addison-Wesley Professional, 2015, p. 357.

[81] British Standards Intitute, “ISO42010 Systems and software engineering — Architecture description,” British Standards Institute, 2011.

[82] SABSA, “SABSA,” www.SABSA.org, 2011.

[83] N. C. C. S. W. Group, “NIST Cloud Computing Security Reference Architecture,” National Institute of Standards and Technology.

[84] Cloud Security Alliance, “Cloud Security Alliance: Enterprise Architecture,” 2017. [Online]. Available: https://research.cloudsecurityalliance.org/tci/. [Accessed 13 March 2017].

[85] National Institute of Standards and Technology, “Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach,” NIST, 2010.

[86] I. Dobson and J. Hietala, “Jericho Forum declare "success" and sunsets,” 29 October 2013. [Online]. Available: https://blog.opengroup.org/2013/10/29/jericho-forum-declares-success-and-sunsets/. [Accessed 3 February 2017].

[87] Cloud Security Alliance, “Cloud Controls Matrix Working Group,” [Online]. Available: https://cloudsecurityalliance.org/group/cloud-controls-matrix/. [Accessed 4 December 2016].

[88] Telecommunication Standardisation Sector of ITU, “Information technology – Cloud computing – Reference architecture,” ITU, 2014.

[89] Merriam-Webster, “Taxonomy,” [Online]. Available: http://www.merriam-webster.com/dictionary/taxonomy. [Accessed 28 11 2016].

[90] C. N. Hofer and G. Karagiannis, “Cloud Computing Services: Taxonomy and Comparison,” *Journal of Internet Services and Applications,* vol. 2, no. 2, pp. 81-94, September 2011.

[91] C. Clayton, “Standard Cloud Taxonomies and Windows Azure,” Microsoft, 7 June 2011. [Online]. Available: https://blogs.msdn.microsoft.com/cclayton/2011/06/07/standard-cloud-taxonomies-and-windows-azure/. [Accessed 4 December 2016].

[92] Cisco, “MPLS FAQ For Beginners,” Cisco, 2 May 2016. [Online]. Available: http://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html. [Accessed 9 March 2017].

[93] A. Sharma, “MPLS - Hub and Spoke Topology,” 26 July 2014. [Online]. Available: http://ip-mpls.com/mpls/mpls-hub-and-spoke-topology/. [Accessed 10 March 2017].

[94] J. Scarpati, “MPLS networks not obsolete, but Internet as WAN catches up,” TechTarget, June 2014. [Online]. Available: http://searchnetworking.techtarget.com/feature/MPLS-networks-not-obsolete-but-Internet-as-WAN-catches-up. [Accessed 9 March 2017].

[95] D. Bubley, “Is SD-WAN a Quasi-QoS overlay for enterprise, independent of telcos & NFV?,” 22 March 2016. [Online]. Available: http://disruptivewireless.blogspot.co.uk/2016/03/is-sd-wan-quasi-qos-overlay-for.html. [Accessed 9 March 2017].

[96] Cloud Standards, “Cloud Standards WiKi,” 15 November 2016. [Online]. Available: http://cloud-standards.org/wiki/index.php?title=Main_Page. [Accessed 29 January 2017].

[97] A. Gordon, The Official ISC2 Guide to the CCSP CBK, 2nd ed., Sybes, 2016, p. 62.

[98] D. Gardner, Risk: The Science and Politics of Fear, Virgin Books, 2008.

[99] Wisegate, “A CIO's Cloud Decision and 7 Lessons Learned From Peers,” Wisegate, Austin, 2012.

[100] A.-M. Dinu, “General Concepts Regarding Risk Appetite,” *Knowledge Horizons - Economics ,* vol. 6, no. 2, pp. 157-159, 2014.

[101] D. Pritchard, “RISK,” *Metaphilosophy,* vol. 46, no. 3, pp. 436-461, 22 July 2015.

[102] US Department of Homeland Security, “DHS Risk Lexicon 2010,” 2010.

[103] Risk Steering Committee, “DHS Risk Lexicon - 2008 Edition,” US Department of Homeland Security, 2008.

[104] C. Hood and D. K. C. Jones, Accident and Design: Contemporary Debates in Risk Management, First ed., London: Routledge, 1996, pp. 2-3.

[105] International Organisation for Standardisation, “ISO Guide 73:2009 - Risk Management Vocabulary,” ISO, 2009.

[106] Information Security Forum, “Information Risk Assessment Methodology 2,” Information Security Forum, 2014.

[107] National Institute of Standards and Technology, “Guide for Conducting Assessments - 800-30,” NIST, 2012.

[108] D. Vohradsky, “Cloud Risk - 10 Principles and a Framework for Assessment,” ISACA, 2012. [Online]. Available: http://www.isaca.org/Journal/archives/2012/Volume-5/Pages/Cloud-Risk-10-Principles-and-a-Framework-for-Assessment.aspx. [Accessed 29 January 2017].

[109] International Organisation for Standardisation, “Information technology -- Security techniques -- Code of practice for information security controls,” International Organisation for Standardisation, 2013.

[110] Open Web Application Security Project, “Cloud-10 Multi Tenancy and Physical Security,” OWAP, [Online]. Available: https://www.owasp.org/index.php/Cloud-10_Multi_Tenancy_and_Physical_Security. [Accessed 19 December 2016].

[111] P. E. Proctor, D. C. Plummer and J. Heiser, “A Public Cloud Risk Model: Accepting Cloud Risk is OK, Ignoring Cloud Risk Is Tragic,” Gartner, 2015.

[112] Joint Task Force Transformation Initative, “Managing Information Security Risk Organization, Mission, and Information System View,” National Institute of Standards and Technology, 2011.

[113] S. Sims, “Qualitative vs. Quantitative Risk Assessment,” SANS, 3 May 2012. [Online]. Available: https://www.sans.edu/cyber-research/leadership-laboratory/article/risk-assessment. [Accessed 2 February 2017].

[114] P. Institute, “2016 Cost of Data Breach Study: Global Analysis,” Ponemon Institute LLC, 2016.

[115] Amazon Web Services, “AWS Security Best Practices,” 2016.

[116] E. U. A. f. N. a. I. Security, “About ENISA,” 2015. [Online]. Available: https://www.enisa.europa.eu/about-enisa. [Accessed 4 February 2017].

[117] T. Wilhelm, Professional Penetration Testing: Creating and Learning in a Hacking Lab, vol. 2, Syngress, 2013, p. 94.

[118] J. Brodkin, "Amazon S3 outage takes out large parts of the Internet," Ars Technica, 28 February 2017. [Online]. Available: https://arstechnica.co.uk/information-technology/2017/02/amazon-s3-cloud-outage/. [Accessed 21 March 2017].

[119] Amazon Web Services, "Summary of the Amazon S3 Service Disruption in the Northern Virginia (US-EAST-1) Region," 2017. [Online]. Available: https://aws.amazon.com/message/41926/. [Accessed 21 March 2017].

[120] G. Gigerenzer, "Dread Risk, September 11, and Fatal Traffic Accidents," *Psychological Science,* vol. 4, no. 15, pp. 286-287, 2004.

[121] O. f. N. Statistics, "Crime in England and Wales: Year ending Sept 2016," Office for National Statistics, 2017.

[122] M. Evans, "Fraud and cyber crime are now the country's most common offences," 19 January 2017. [Online]. Available: http://www.telegraph.co.uk/news/2017/01/19/fraud-cyber-crime-now-countrys-common-offences/. [Accessed 4 February 2017].

[123] D. Pauli, "Hackers pop 6000 sites on active 18-month carding bonanza," The Register, 13 October 2016. [Online]. Available: https://www.theregister.co.uk/2016/10/13/hackers_pop_6000_sites_on_active_18month_carding_bonanza/. [Accessed 5 March 2017].

[124] The History Channel, "9/11: Timeline of Events," 2017. [Online]. Available: http://www.history.com/topics/9-11-timeline. [Accessed 23 February 2017].

[125] B. Krebs, "Espionage Hackers Target 'Watering Hole' Sites," Krebs on Security, 25 September 2012. [Online]. Available: https://krebsonsecurity.com/tag/watering-hole-attack/. [Accessed 5 March 2017].

[126] A. Shostack, Threat Modelling: Designing for Security, Wiley, 2014, p. 246.

[127] J. Kindervag, "No More Chewy Centers: Introducing The Zero Trust Model Of Information Security," Forrester Research, 2010.

[128] S. E. Schmidt, "AWS Security," in *ASW Summits 2014*, 2014.

[129] "Notorious Nine: Cloud Computing Top Threats in 2013," 2013.

[130] T. T. W. Group, "The Treacherous 12: Cloud Computing Top Threats in 2016," Cloud Security Alliance, 2016.

[131] The European Parliment and the Council of the European Union, “REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL,” The European Parliment and the Council of the European Union, 2016.

[132] T. Hyman, “Is cloud the biggest risk to GDPR compliance?,” LinkedIn, 13 November 2016. [Online]. Available: https://www.linkedin.com/pulse/cloud-biggest-risk-gdpr-compliance-tim-hyman. [Accessed 19 February 2017].

[133] W. Ashford, “Most cloud applications not GDPR-ready, report reveals,” 28 July 2016. [Online]. Available: http://www.computerweekly.com/news/450301241/Most-cloud-applications-not-GDPR-ready-report-reveals. [Accessed 19 February 2017].

[134] The Open Group, The Open Group Architecture Framework (TOGAF) Version 9, V. H. Publishing, Ed., Basharat Hussain, 2009, p. 430.

[135] The Open Group, “Architectural Artifacts,” 2011. [Online]. Available: http://pubs.opengroup.org/architecture/togaf9-doc/arch/chap35.html. [Accessed 19 February 2017].

[136] D. Gabel and T. Hickman, “Chapter 10: Obligations of controllers – Unlocking the EU General Data Protection Regulation,” White & Case, 2016. [Online]. Available: https://www.whitecase.com/publications/article/chapter-10-obligations-controllers-unlocking-eu-general-data-protection. [Accessed 17 March 2017].

[137] A. Litan, “How PCI failed Target and U.S. Consumers,” Gartner, 20 January 2014. [Online]. Available: http://blogs.gartner.com/avivah-litan/2014/01/20/how-pci-failed-target-and-u-s-consumers/.

[138] Amazon Web Services, “Shared Responsibiilty Model,” 2017. [Online]. Available: https://aws.amazon.com/compliance/shared-responsibility-model/. [Accessed 5 February 2017].

[139] Microsoft, “The STRIDE Threat Model,” Microsoft, 2005. [Online]. Available: https://msdn.microsoft.com/en-us/library/ee823878(v=cs.20).aspx. [Accessed 7 February 2017].

[140] Financial Conduct Authority, “FG 16/5 - Guidance for firms outsourcing to the ‘cloud’ and other third-party IT services,” Financial Conduct Authority, 2016.

[141] PCI Security Standards Council, “Document Library,” 2017. [Online]. Available: https://www.pcisecuritystandards.org/document_library. [Accessed 17 March 2017].

[142] Cloud Special Interest Group: PCI Security Standards Council, “Information Supplement: PCI DSS Cloud Computing Guidelines,” PCI Security Standards Council., 2013.

[143] Amazon Web Services, “AWS Case Studies (All),” 2017. [Online]. Available: https://aws.amazon.com/solutions/case-studies/all/. [Accessed 11 March 2017].

[144] Y. V. Natis and E. Knipp, “Gartner Reference Architecture for Multitenancy,” Gartner, 2010.

[145] W. J. Brown, V. Anderson and Q. Tan, “Multitenancy - Security Risks and Countermeasures,” Athabasca University.

[146] R. L. Villars and R. Perry, “The Business Value of VCE Vblock Systems: Leveraging Convergence to Drive Business Agility,” IDC, 2015.

[147] Dell EMC, “VxBlock and VBlock Systems,” Dell, 2017. [Online]. Available: https://www.emc.com/en-us/converged-infrastructure/converged-systems.htm#collapse=. [Accessed 19 February 2017].

[148] Openstack, “Open source software for creating private and public clouds,” 2017. [Online]. Available: https://www.openstack.org/. [Accessed 19 February 2017].

[149] L. MacVittie, “Multi-Tenancy Requires More Than Just Isolating Customers,” F5 Networks, 9 August 2010. [Online]. Available: https://devcentral.f5.com/articles/multi-tenancy-requires-more-than-just-isolating-customers. [Accessed 26 January 2017].

[150] vFabric Team, “Putting the ‘Single’ Back in Single Sign-On (SSO),” VMWare, 14 March 2013. [Online]. Available: https://blogs.vmware.com/vfabric/2013/03/putting-the-single-back-in-single-sign-on-sso.html. [Accessed 7 March 2017].

[151] J. Kabbedijk, C.-P. Bezemer, S. Jansen and A. Zaidman, “Defining Multi-Tenancy: A Systematic Mapping Study on the Academic and the Industrial Perspective,” Utrecht, 2014.

[152] International Organisation for Standardisation, “Information Technology -- Open Systems Interconnection - Basic Reference Model,” International Organisation for Standardisation, 1994.

[153] Microsoft Azure, “Design patterns for multitenant SaaS applications and Azure SQL Database,” Microsoft, [Online]. Available: https://docs.microsoft.com/en-us/azure/sql-database/sql-database-design-patterns-multi-tenancy-saas-applications . [Accessed 4 January 2017].

[154] T. Ristenpart, E. Tromer, H. Shacham and S. Savage, “Hey, You. Get off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds,” ACM, San Diego, 2009.

[155] D. Perez-Botero, J. Szefer and R. B. Lee, “Characterizing Hypervisor Vulnerabilities in Cloud Computing Servers,” in *Workship on Security in Cloud Computing*, May, 2013.

[156] A. Howard, “Increase in Discovered Hypervisor Vulnerabilities,” VMWare, 21 November 2016. [Online]. Available: http://vmblog.com/archive/2016/11/21/kudelski-security-2017-predictions-increase-in-discovered-hypervisor-vulnerabilities.aspx#.WI3vCbaLRTY. [Accessed 29 January 2017].

[157] Zero Day Initative, “PWN2OWN Contest,” Trend Micro, 2017. [Online]. Available: http://zerodayinitiative.com/Pwn2Own2017Rules.html. [Accessed 8 February 2017].

[158] G. Grispos, T. Storer and W. B. Glisson, “Calm Before the Storm: The Challenges of Cloud Computing in Digital Forensics,” *The International Journal of Digital Crime and Forensics ,* vol. 4, no. 2, pp. 28-48, 2012.

[159] A. Kamath, “View and analyze Azure Audit Logs in Power BI and more,” 30 September 2015. [Online]. Available: https://azure.microsoft.com/en-us/blog/analyze-azure-audit-logs-in-powerbi-more/. [Accessed 29 January 2017].

[160] Zscaler, “Zscaler Nanolog Streaming Service,” Zscaler, San Jose, 2016.

[161] A. Amini, N. Jamil, A. R. Ahmad and M. R. Z'aba, “Threat Modelling Approaches for Securing Cloud Computing,” *Journal of Applied Sciences,* vol. 15, no. 7, pp. 953-965, 2015.

[162] M. Fawzi, “Virtualization and Protection Rings (Welcome to Ring -1) Part I,” 24 May 2009. [Online]. Available: https://fawzi.wordpress.com/2009/05/24/virtualization-and-protection-rings-welcome-to-ring-1-part-i/. [Accessed 16 January 2017].

[163] Intel, “Intel® 64 and IA-32 Architectures Software Developer Manuals,” Intel, 29 December 2016. [Online]. Available: https://software.intel.com/en-us/articles/intel-sdm#combined. [Accessed 29 January 2017].

[164] Intel, “Intel Virtualisation Technology (Intel VT),” [Online]. Available: http://www.intel.com/content/www/us/en/virtualization/virtualization-technology/intel-virtualization-technology.html. [Accessed 29 January 2017].

[165] AMD, “Virtualisation Solutions,” 2017. [Online]. Available: http://www.amd.com/en-gb/solutions/servers/virtualization#. [Accessed 28 January 2017].

[166] C. C. Elisan, M. A. Davis, S. M. Bodmer and A. LeMasters, Hacking Exposed: Malware and Rootkits, vol. 2nd, New York: McGraw Hill Education, 2017.

[167] B. Milewski, “Virtual Machines: Thin Hypervisor,” 3 January 2012. [Online]. Available: https://corensic.wordpress.com/2012/01/03/virtual-machines-thin-hypervisor/. [Accessed 5 February 2017].

[168] Microsoft, “Hyper-V,” Microsoft, 4 April 2016. [Online]. Available: https://technet.microsoft.com/en-us/library/mt169373(v=ws.11).aspx. [Accessed 5 February 2017].

[169] Terendo, “File: Hyper-V.png,” WikiMedia Commons, 9 February 2010. [Online]. Available: https://commons.wikimedia.org/wiki/File:Hyper-V.png. [Accessed 5 February 2017].

[170] J. Szefer, E. Keller, R. B. Lee and J. Rexford, “Eliminating the Hypervisor Attack Surface for a More Secure Cloud,” Princeton University, 2011.

[171] QEMU, “What is QEMU,” 2016. [Online]. Available: http://www.qemu-project.org/. [Accessed 9 February 2017].

[172] Crowdstrike, “Venom,” Crowdstrike, 21 May 2015. [Online]. Available: http://venom.crowdstrike.com/. [Accessed 28 January 2017].

[173] Rackspace, “Virtual Cloud Servers Powered by OpenStack,” Rackspace, 2017. [Online]. Available: https://www.rackspace.com/en-gb/cloud/servers. [Accessed 5 February 2017].

[174] CVE Details, “VM VirtualBox: Security Vulnerabiities,” MITRE, 2017. [Online]. Available: https://www.cvedetails.com/vulnerability-list/vendor_id-93/product_id-20406/Oracle-Vm-Virtualbox.html. [Accessed 7 March 2017].

[175] CVE Details, “VMWare Workstation: Security Vulnerabilities,” MITRE, 2016. [Online]. Available: https://www.cvedetails.com/vulnerability-list/vendor_id-252/product_id-436/Vmware-Workstation.html. [Accessed 7 March 2017].

[176] VMWare, “Workstation for Windows,” VMWare, 2017. [Online]. Available: https://www.vmware.com/uk/products/workstation.html. [Accessed 5 February 2017].

[177] Oracle, “Welcome to VirtualBox.org,” Oracle, 2017. [Online]. Available: https://www.virtualbox.org/. [Accessed 5 February 2017].

[178] G. Henningsen, “My New Favorite Tool: Oracle VM Virtualbox,” November 2011. [Online]. Available: http://www.oracle.com/technetwork/articles/servers-storage-admin/vmlove-1368887.html. [Accessed 9 February 2017].

[179] J. Rutkowska, “Subverting the Vista Kernel for Fun and Profit,” in *Blackhat Europe 2006*.

[180] S. Irwin, “Creating a Threat Profile for Your Organization,” SANS Institute, 2014.

[181] R. Wojtczuk, “Poacher Turned Gatekeeper: Lessons Learned From Eight Years of Breaking Hypervisors,” in *Blackhat USA 2014*.

[182] Sourceforge, “ISSAF,” Sourceforge, 05 05 2015. [Online]. Available: https://sourceforge.net/projects/isstf/. [Accessed 5 February 2017].

[183] NMAP, “NMAP,” [Online]. Available: https://nmap.org/. [Accessed 5 February 2017].

[184] S. Sanfilippo, 2006. [Online]. Available: http://www.hping.org/. [Accessed 5 February 2017].

[185] G. O. System, “GNU WGET,” 25 May 2016. [Online]. Available: https://www.gnu.org/software/wget/. [Accessed 5 February 2017].

[186] Offensive Security, “Kali Linux,” Offensive Security, 2017. [Online]. Available: https://www.kali.org/. [Accessed 20 February 2017].

[187] Amazon Web Services, “Kali Linux,” 2017. [Online]. Available: https://aws.amazon.com/marketplace/pp/B01M26MMTT. [Accessed 20 February 2017].

[188] Amazon Web Services, “Instance Types,” 2017. [Online]. Available: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-types.html. [Accessed 18 March 2017].

[189] O. W. A. S. Project, “Buffer Overflow Attack,” 3 September 2014. [Online]. Available: https://www.owasp.org/index.php/Buffer_overflow_attack. [Accessed 11 February 2017].

[190] Department of Homeland Security / Federal Bureau of Investigation, “GRIZZLY STEPPE – Russian Malicious Cyber Activity,” 2016.

[191] P. Agnihotri, *Fraud Detection and Machine Learning on AWS,* YouTube, 2016.

[192] Amazon Web Services, “How do I create and activate a new Amazon Web Services account?,” 30 December 2015. [Online]. Available: https://aws.amazon.com/premiumsupport/knowledge-center/create-and-activate-aws-account/. [Accessed 18 March 2017].

[193] Common Vulnerabiilties and Exposures, “CVE-2015-7835,” 2015. [Online]. Available: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2015-7835. [Accessed 8 February 2017].

[194] National Institute of Standards and Technology, “Vulnerability Summary for CVE-2015-3456,” 13 May 2015. [Online]. Available: https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2015-3456. [Accessed 21 January 2017].

[195] Sourceforge, “Incognito,” Sourceforge, 16 March 2016. [Online]. Available: https://sourceforge.net/projects/incognito/. [Accessed 21 January 2017].

[196] Nirsoft, "LSA Secrets Dump," Nirsoft, [Online]. Available: http://www.nirsoft.net/utils/lsa_secrets_dump.html. [Accessed 21 January 2017].

[197] The Standard for Information Security Vulnerbility Names, "CVE-2014-7188," 2014. [Online]. Available: http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2014-7188. [Accessed 8 February 2017].

[198] A. Greenberg, "The Shadow Brokers Mess Is What Happens When the NSA Hoards Zero-Days," Wired, 08 July 2016. [Online]. Available: https://www.wired.com/2016/08/shadow-brokers-mess-happens-nsa-hoards-zero-days/. [Accessed 21 February 2017].

[199] World Anti-Doping Agency, "Cyber Hack Update: Data leak concerning 41 athletes from 13 countries and 17 sports," 23 September 2016. [Online]. Available: https://www.wada-ama.org/en/media/news/2016-09/cyber-hack-update-data-leak-concerning-41-athletes-from-13-countries-and-17. [Accessed 21 February 2017].

[200] S. T. King, P. M. Chen, Y.-M. Wang, C. Verbowski, H. J. Wang and J. R. Lorch, "Subvirt: Implementing Malware with Virtual Machines".

[201] T. Ptacek, N. Lawson and P. Ferrie, "Don't Tell Joanna, The Virtualized Rootkit Is Dead," in *Blackhat USA 2007*.

[202] M. S. I. Apecechea, T. Eisenbarth and B. Sunar, "Fine grain Cross-VM Attacks on Xen and VMware are possible!," in *2014 IEEE Fourth International Conference on Big Data and Cloud Computing,* , Sydney, 2014.

[203] OpenSSL, "OpenSSL: Cryptography and SSL/TLS Toolkit," OpenSSL Software Foundation, 2016. [Online]. Available: https://www.openssl.org/. [Accessed 13 March 2017].

[204] BlackDuck, "PolarSSL," BlackDuck, 2017. [Online]. Available: https://www.openhub.net/p/PolarSSL. [Accessed 13 March 2017].

[205] Free Software Foundation Inc, "Libgcrypt," Free Software Foundation Inc, 27 August 2015. [Online]. Available: https://www.gnu.org/software/libgcrypt/. [Accessed 13 March 2017].

[206] "Intel Advanced Encryption Standard New Instructions Ecosystem," Intel, 2013.

[207] Spirovskib, "Hacking Virtualisation: Sniffing Traffic," Youtube, 2010.

[208] Amazon Web Services, "VPC Security Capabilites," Amazon, 2017. [Online]. Available: https://aws.amazon.com/answers/networking/vpc-security-capabilities/. [Accessed 21 February 2017].

[209] "Dyre times ahead: Zeus-style trojan slurps your banking login creds," The Register, 8 July 2015. [Online]. Available: https://www.theregister.co.uk/2015/07/08/dyre_banking_trojan_spam_surge/. [Accessed 23 February 2017].

[210] N. Woolf, "DDoS attack that disrupted the internet was largest of its kind in history, experts say," 26 October 2016. [Online]. Available: https://www.theguardian.com/technology/2016/oct/26/ddos-attack-dyn-mirai-botnet. [Accessed 27 January 2017].

[211] D. Hyde, "A Survey on the Security of Virtual Machines," 2009.

[212] Apple, "iCloud," Apple Inc., 5 February 2017. [Online]. Available: https://www.icloud.com. [Accessed 5 February 2017].

[213] M. Landi, "Stars' nude photo attack may have been down to password codes," Independent Ireland, 9 September 2014. [Online]. Available: http://www.independent.ie/business/technology/news/stars-nude-photo-attack-may-have-been-down-to-password-codes-30552629.html. [Accessed 5 February 2017].

[214] Bitdefender, "Bitdefender," 2017. [Online]. Available: https://www.bitdefender.co.uk/. [Accessed 5 February 2017].

[215] T. Fox-Brewster, "Anti-Virus Firm BitDefender Admits Breach, Hacker Claims Stolen Passwords Are Unencrypted," Forbes, 31 July 2015. [Online]. Available: http://www.forbes.com/sites/thomasbrewster/2015/07/31/bitdefender-hacked/#2bbb060af4ec. [Accessed 5 February 2017].

[216] Hacker Film Blog, "Antivirus Maker Bitdefender Hacked, Customer Data Being Sold In Shady Black Market Deals," 29 July 2015. [Online]. Available: http://www.hackerfilm.com/2015/07/antivirus-maker-bitdefender-hacked.html. [Accessed 5 February 2017].

[217] Gemalto, "Breach Level Index," Gemalto, 2017. [Online]. Available: http://breachlevelindex.com/. [Accessed 5 February 2017].

[218] T. Mendelsohn, "AdultFriendFinder hacked: 400 million accounts exposed," Ars Technica, 14 November 2016. [Online]. Available: https://arstechnica.com/security/2016/11/adultfriendfinder-hacked-exposes-400-million-hookup-users/. [Accessed 5 February 2017].

[219] S. Ragan, "Researcher says Adult Friend Finder vulnerable to file inclusion vulnerabilities," CSO Online, 18 October 2016. [Online]. Available: http://www.csoonline.com/article/3132533/security/researcher-says-adult-friend-finder-vulnerable-to-file-inclusion-vulnerabilities.html. [Accessed 5 February 2017].

[220] OWASP, "Testing for Local File Inclusion," OWASP, 8 August 2014. [Online]. Available: https://www.owasp.org/index.php/Testing_for_Local_File_Inclusion. [Accessed 5 February 2017].

[221] J. Yeoh and H. Baron, “Identity Solutions: Security Beyond the Perimeter,” Cloud Security Alliance, 2016.

[222] Linux Foundation Collaborative Projects, “Let's Encrypt,” Linux Foundation Collaborative Projects, 2017. [Online]. Available: https://letsencrypt.org/. [Accessed 23 February 2017].

[223] Bitcoin, “Bitcoin Homepage,” Bitcoin, 2017. [Online]. Available: https://bitcoin.org/en/. [Accessed 16 March 2017].

[224] Coindesk, “How to Set Up a Bitcoin Miner,” 26 November 2016. [Online]. Available: http://www.coindesk.com/information/how-to-set-up-a-miner/. [Accessed March 2017].

[225] Microsoft, “Security Intelligence Report,” Microsoft, 2016.

[226] Rapid 7, “Metasploit,” 2017. [Online]. Available: https://www.metasploit.com/. [Accessed 9 March 2017].

[227] B. Quinn and M. Brignall, “Boy, 15, arrested in Northern Ireland over TalkTalk cyber-attack,” The Guardian, 27 October 2015. [Online]. Available: https://www.theguardian.com/business/2015/oct/26/talktalk-cyber-attack-boy-15-arrested-in-northern-ireland. [Accessed 23 February 2017].

[228] Akamai, “State of the Internet: Security - 2016 Executive Summary,” Akamai, 2017.

[229] Cloudflare, Cloudflare, 2017. [Online]. Available: https://www.cloudflare.com/. [Accessed 23 February 2017].

[230] M. Majkowski, “400Gbps: Winter of Whopping Weekend DDoS Attacks,” 3 March 2016. [Online]. Available: https://blog.cloudflare.com/a-winter-of-400gbps-weekend-ddos-attacks/. [Accessed 23 February 2017].

[231] C. Donnelly, “BP to go all-in on public cloud to cut datacentre upkeep costs,” Computer Weekly, 30 November 2016. [Online]. Available: http://www.computerweekly.com/news/450403781/BP-to-go-all-in-on-public-cloud-to-cut-datacentre-upkeep-costs. [Accessed 9 March 2017].

[232] Rich, “Introducing the Data Security Lifecycle,” Securosis, 09 August 2011. [Online]. Available: https://securosis.com/blog/introducing-the-data-security-lifecycle-2.0. [Accessed 22 February 2017].

[233] F5 Networks, “Inside the head of a hacker,” F5 Networks, 2017.

[234] B. Schneier, “Homomorphic Encryption Breakthrough,” 9 July 2009. [Online]. Available: https://www.schneier.com/blog/archives/2009/07/homomorphic_enc.html. [Accessed 18 March 2017].

[235] CNBC, “Uber’s books still top secret, but its biggest weakness isn't,” 8 June 2016. [Online]. Available: http://www.cnbc.com/2016/06/08/ubers-66-billion-valuation-may-ride-on-shaky-foundation.html. [Accessed 2 December 2016].

[236] J. Verhage, “An Expert in Valuation Says Uber Is Only Worth $28 Billion, Not $62.5 Billion,” 17 August 2016. [Online]. Available: https://www.bloomberg.com/news/articles/2016-08-17/an-expert-in-valuation-says-uber-may-have-already-peaked. [Accessed 2 December 2016].

[237] A. Gray, “Cyber risks too big to cover, says Lloyd’s insurer,” Financial Times, 5 February 2015. [Online]. Available: https://www.ft.com/content/94243f5a-ad38-11e4-bfcf-00144feab7de. [Accessed 2 February 2017].

[238] E. M. Hutchins, M. J. Cloppert and R. M. Amin, “Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains,” *Leading Issues in Information Warfare & Security Research ,* vol. 1, p. 80, 2011.

[239] Microsoft, “Windows Homepage,” 17 March 2017. [Online]. Available: https://www.microsoft.com/en-gb/windows. [Accessed 17 March 2017].

# APPENDIX A: ENISA CLOUD VULNERABILITY ASSESSMENT

| Vulnerabilities Specific to Cloud | Description | Originating \| Exacerbated \| General | People \| Process \| Technology | Context |
|---|---|---|---|---|
| AAA | Poor systems for authentication, authorisation and accounting could facilitate unauthorised access to resources. Poor AAA could result in a limited or non-existent ability to audit or forensically analyse a breach | General | People, Process and Technology | Latency and / or synchronisation issues could be introduced with a cloud-based AAA provider. Good practice recommendations would recommend the use of SAML which would leverage the customer's existing identity provider. Directory services latency is not the exclusive reserve of public cloud computing. Organisations that operate in multiple geographies have network architectures which unavoidably introduce network latency |
| User Provisioning | Concerns exist regarding a customer's ability to provision credentials for cloud services. | Exacerbated | Process | Management of multiple application credentials introduces operational overheads and vectors for compromise. |
| User De-provisioning | Users remain provisioned once expired or revoked due to synchronisation delays | Exacerbated | Process | As with above entry |
| Remote access to administrative interface | Vulnerabilities in an endpoint could compromise cloud infrastructure. | General | Technology | Interface could be compromised in any deployment scenario |

| | | | | |
|---|---|---|---|---|
| Hypervisor | The hypervisor in fact fully controls the physical resources and the VMs running on top of it. Exploiting the hypervisor potentially means exploiting every VM. | Exacerbated | Technology | Hypervisor technology is at the core of virtualisation. Without virtualisation, cloud computing efficiencies and performance would not be possible. A mitigating factor being that datacentre technologies, irrespective of location, are moving to virtualised stacks and multitenancy. This is discussed in Chapters 3 and 4. |
| Lack of Resource Isolation | Resources of one organisation can affect resource use by another customer | Originating | People \| Process \| Technology | Public Cloud specific vulnerability which forms the basis of chapter 4 |
| Lack of Reputational Isolation | Activities of one customer impact the reputation of another customer | Exacerbated | Technology | Exacerbated by cloud although the vulnerability has existed for as long as companies have retained partnerships; digital or physical |
| Communication Encryption | These vulnerabilities concern the possibility of reading data in transit via, for example, MITM attacks, poor authentication, acceptance of self-signed certificates, etc. | General | Technology | Issues for all and any communication of sensitive information. |
| Lack of or weak encryption of archives and data in transit | Reading of data in transit, at rest in archives, file shares or disk images. | General | Technology | As with above entry |
| Impossibility of processing data in encrypted form | Encrypting data at rest is not difficult, but despite recent advances in homomorphic encryption, there is little prospect of any commercial system being able to maintain this encryption during processing. In one article, Bruce Schneier estimates that performing a web search with encrypted keywords -- a perfectly | Exacerbated | Technology | Mature players in the IaaS, PaaS and SaaS markets provide mechanisms for at rest data encryption. Often improving on the security posture of the customer organisation. |

| | | | | |
|---|---|---|---|---|
| | reasonable simple application of this algorithm -- would increase the amount of computing time by about a trillion [234]. This means that for a long time to come, cloud customers doing anything other than storing data in the cloud must trust the cloud provider. | | | |
| Poor Key Management Procedures | Cloud computing infrastructures require the management and storage of many kinds of keys; examples include session keys to protect data in transit (e.g., SSL keys), file encryption keys, key pairs identifying cloud providers, key pairs identifying customers, authorisation tokens and revocation | General | Process | Key Management is a customer issue. |
| Key Generation: Low Entropy for Random Number Generation | The combination of standard system images, virtualisation technologies and a lack of input devices means that systems have much less entropy than physical RNGs; see Cloud Computing Security (33). This means that an attacker on one virtual machine may be able to guess encryption keys generated on other virtual machines because the sources of entropy used to generate random numbers might be similar. This is | General | Process | Key Management is a customer issue. |

|  |  |  |  |  |
|---|---|---|---|---|
|  | not a difficult problem to solve, but if it is not considered during system design, it can have important consequences. |  |  |  |
| Lack of standard technologies and solutions | A lack of standards means that data may be ‘locked-in’ to a provider. This is a big risk should the provider cease operation. | Exacerbated | Process | Standards are rapidly evolving across Cloud Computing. |
| No source escrow agreement | Lack of source escrow means that if a PaaS or SaaS provider goes into bankruptcy, its customers are not protected. | General | Process | As with any partner relationship |
| Inaccurate model of resource usage | resource exhaustion because they are provisioned statistically. Although many providers allow customers to reserve resources in advance, resource provisioning algorithms can fail | General | Technology | Mature CSPs have the scale and capacity to support customer resource demand requirements. Considerably more so than the constraints with an appliance-based model |
| No control or vulnerability management process | Restrictions on port scanning and vulnerability testing are an important vulnerability which, combined with a ToU which places responsibility on the customer for securing elements of the infrastructure, is a serious security problem. | General | Process, Technology | Cloud providers often provide services for the assessment of vulnerabilities in cloud solutions |
| Possibility that internal (cloud) network probing will occur | Cloud customers port scanning and identifying other customers / hosts | General | Technology | Any resource which is publicly addressable from the Internet is |
| Possibility that co-residence checks will be performed | Side channel attacks exploiting a lack of resource isolation | Originating | Technology | Explored in chapter 4. Attacks require significant investment in time and result in a disproportionate amount of effort for the returns. |

| | | | | |
|---|---|---|---|---|
| Lack of forensic readiness | Restricted access to important information in the event of a breach: IP information, etc. | Exacerbated | Process, Technology | Discussed in Chapter 3. Can be argued both ways. Public cloud offers improvements in capability for offline machine images and |
| Sensitive media sanitisation | Shared tenancy of physical storage means sensitive data may leak because data destruction policy could be problematic to implement | Exacerbated | Process, Technology | Data destruction issues not reserved for cloud. Mature CSPs providing capabilities for secure destruction aligned to industry standards - NIST, etc. |
| Synchronising responsibilities or contractual obligations external to cloud | Cloud customers unaware of their responsibilities in the customer / CSP relationship | Exacerbated | Process | Ambiguity exacerbated by cloud but endemic across customer <> partner relationships. Awareness and visibility of amended working practices is needed. |
| Cross-cloud applications creating hidden dependency | Hidden dependencies in supply chain and cloud provider architecture | Exacerbated | Process | Cloud or otherwise, loosley-coupled solutions invariably introduce process dependencies |
| SLA clauses with conflicting promises to different stakeholders | Self explanatory | Exacerbated | Process | Customer must ensure that the SLAs are suitable for their use cases. |
| SLA clauses containing excessive business risk | Self explanatory | General | Process | As with any partner relationship |
| Audit or certificating not available to customers | No assurances over audit and certification | General | Process | AWS actually improves an organisation's ability to comply with regulatory compliance frameworks. This is covered within Chapters 2 and 3 |
| Certification schemes not adapted to cloud infrastructure | No cloud-specific control | Exacerbated | Process | As covered in chapter 2. Maturity is growing in this space. |
| Inadequate resource provisioning and investments in infrastructure | Assumption cloud provider inadequately provisioned | General | Technology | Public cloud provides scale and flexibility to mitigate these vulnerabilities. |
| No policies for resource capping | Volatility of resource allocation | General | Process | As above |
| Storage of data in multiple jurisdictions and lack of transparency about this | Opaque agreements as to where data is stored | Originating | Process | This is a cloud-first problem. Visibility and education needed at the customer |

| | | | | |
|---|---|---|---|---|
| Lack of information on jurisdictions | As above | Originating | Process | As above |
| Lack of completeness and transparency in terms of use | As above | Exacerbated | Process | A lack of transparency in terms of use |
| **Non-Cloud Vulnerabilities** | | | | |
| Lack of security awareness | Cloud customers are not aware of the risks they could face when migrating into the cloud | Exacerbated | People, Process | Credible vulnerability. Users are not familiar with cloud technologies nor the processes for secure enablement. |
| Lack of vetting process | Since there may be very high privilege roles within cloud providers, due to the scale involved, the lack or inadequate vetting of the risk profile of staff with such roles is an important vulnerability. | Exacerbated | Process | Mature CSPs provide comprehensive staff vetting as I cover in Chapter 3 although cloud does introduce an element of the unknown if appropriate due dilligence is not followed. |
| Unclear roles and responsibilities | These vulnerabilities regard the inadequate attribution of roles and responsibilities in the cloud provider organization. | Exacerbated | Process | Ambiguity over "who does what". Identified by the CSA in their "Critical Areas of Security for Cloud Computing" which is extensively referenced in my thesis body. |
| Poor enforcement of role definitions | Within the cloud provider, a failure to segregate roles may lead to excessively privileged roles which can make extremely large systems vulnerable. For example, no single person should be given access privileges to the entire cloud. | General | Process \| Technology | Role enforcement issues arise irrespective of data location. |
| Need-to-know principle not applied | This is a special case of a vulnerability regarding roles and responsibilities. Parties should not be given unnecessary access | General | Process \| Technology | This principle is a challenge in all organisations irrespective of the location of service / data. |

| | | | | |
|---|---|---|---|---|
| | to data. If they are then this constitutes an unnecessary risk. | | | |
| Inadequate physical security procedures | Lack of physical perimeter controls / lack of electromagnetic shielding | General | Process \| Technology | Physical security controls invariably stronger at a CSP. |
| Misconfiguration | This class of vulnerabilities include: inadequate application of security baseline and hardening procedures, human error and untrained administrator. | General | Technology | Affects all deployments irrespective of location. |
| System or OS vulnerabilities | Out of data software components present vectors for malware exploitation. | General | Technology | Affects all deployments irrespective of location. |
| Untrusted software | Unapproved, untrusted software | General | Technology | Affects all deployments irrespective of location. |
| Lack of a BCP / DR plan | No business continuity arrangements or DR planning. | General | Process | Affects all deployments irrespective of location. |
| Incomplete asset inventory | No formal register of hardware and software components | General | Process | Affects all deployments irrespective of location. |
| Lack of or poor asset classification | Lack of consistent categorisation and terminology for assets | General | Process | Affects all deployments irrespective of location. |
| Unclear asset ownership | Lack of asset ownership - either a framework or assigned owners. | General | Process | Affects all deployments irrespective of location. |
| Poor identification of project requirements | These include a lack of consideration of security and legal compliance requirements, no systems and applications user involvement, unclear or inadequate business requirements, etc. | General | Process | Affects all deployments irrespective of location. |

|  |  |  |  |  |
|---|---|---|---|---|
| Poor provider selection | Inappropriate providers selected for cloud services | General | Process \| Technology | Poor provider selection is a systemic issue for companies - it's not reserved for public cloud nor does public cloud really introduce any unique considerations |
| Lack of supplier redundancy | Lack of redundancy in terms of service availablity and data portability. | General | Process \| Technology | A consideration for any supplier relationship. An outsourced service would present the same issues regardless of cloud usage. |
| Application vulnerabilities or poor patch management | This class of vulnerabilities include: bugs in the application code, conflicting patching procedures between provider and customer, application of untested patches, vulnerabilities in browsers, etc. | General | Technology | Affects all deployments irrespective of location. |
| Resource consumption vulnerabilities | Vectors for resource exhaustion and denial of service | General | Process \| Technology | The scale and breadth of public cloud improves likelihood of being able to withstand these forms of vulnerability. |
| Breach of NDA by provider | Discloure of sensitive customer information either accidentially or for competitive advantage. | General | People \| Process | Poor provider selection is a systemic issue for companies - it's not reserved for public cloud nor does public cloud really introduce any unique considerations |
| Liability from data loss | Unclear ownership for loss of data: repercussions - fines, etc. | Exacerbated | Process | Another process issue related to awareness and public cloud introducing amended working practices. |
| Lack of policy or poor procedures for log collection and retention | Policy and technology for log collection and retention. | General | Process \| Technology | Affects all deployments irrespective of location. |
| Inadequate or misconfigured filtering resources | Leveraged as a vector for denial of service: network and application vulnerability | General | Technology | Affects all deployments irrespective of location. |